

The Dynamics of Adapting Neurons

THÈSE N° 5217 (2011)

PRÉSENTÉE LE 15 DÉCEMBRE 2011
À LA FACULTÉ SCIENCES DE LA VIE
LABORATOIRE DE CALCUL NEUROMIMÉTIQUE (IC/SV)
PROGRAMME DOCTORAL EN NEUROSCIENCES

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Richard NAUD

acceptée sur proposition du jury:

Prof. C. Petersen, président du jury
Prof. W. Gerstner, directeur de thèse
Prof. L. F. Abbott, rapporteur
Prof. N. Brunel, rapporteur
Prof. M. C. Gastpar, rapporteur



ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Suisse
2011

J'ai fait de plus loin que moi un voyage abracadabrant
— Gaston Miron

MACBETH - Canst thou not raze the written troubles of the brain?
DOCTOR - Therein the patient must minister to himself.
MACBETH - (laugh) Throw the physics to the dog, I'll have none of it.
– William Shakespeare

Acknowledgements

Wulfram Gerstner has played a central part in shaping the present thesis. He is the first to have believed in my work and since then was always very trusting. He guided the footpath of the young scientist remotely and gave an obligingly high quantity of wise advices.

All the members of the lab were highly supportive, especially Skander Mensi, Christian Pozzorini, Felipe Gerhard, Guillaume Hennequin, Tim Vogels, Friedemann Zenke, Henning Sprekeler, Nicolas Frémaux and Brice Bathellier.

I am grateful to external collaborators for their contributions and ideas, in particular Carl Petersen, Michael Avermann, Rainer Gregor, Shigeru Shinomoto, Ryota Kobayashi, Thomas K. Berger, Brice Bathellier and Matthew Larkum.

Many discussions in conferences or with external speakers in Lausanne helped to shape my ideas, I would like to thank Matteo Carandini, Eugene Izhikevich, André Longtin, Idan Segev, Srdjan Ostojic, Walter Senn, Magnus Richardson, Giancarlo La Camera and Renaud Jolivet. Similarly I would like to acknowledge the anonymous reviewers for their much valued criticisms.

My parents and wife were always very encouraging, even in the worse of times. Without them none of this would have been possible.

Lausanne, 6th of September 2011

R. N.

Abstract

How do neurons dynamically encode and treat information? Each neuron communicates with its distinctive language made of long silences intermitted by occasional spikes. The spikes are prompted by the pooled effect of a population of pre-synaptic neurons. To understand the operation made by single neurons is to create a quantitative description of their dynamics. The results presented in this thesis describe the necessary elements for a quantitative description of single neurons. Almost all chapters can be unified under the theme of adaptation. Neuronal adaptation plays an important role in the transduction of a given stimulation into a spike train. The work described here shows how adaptation is brought by every spike in a stereotypical fashion. The spike-triggered adaptation is then measured in three main types of cortical neurons. I analyze in detail how the different adaptation profiles can reproduce the diversity of firing patterns observed in real neurons. I also summarize the most recent results concerning the spike-time prediction in real neurons, resulting in a well-founded single-neuron model. This model is then analyzed to understand how populations can encode time-dependent signals and how time-dependent signals can be decoded from the activity of populations. Finally, two lines of investigation in progress are described, the first expands the study of spike-triggered adaptation on longer time scales and the second extends the quantitative neuron models to models with active dendrites.

Keywords Spiking Neuron Models - Model Characterization - Spike-Train Metrics - Neural Coding - Spike-Frequency Adaptation

Résumé

Comment une information donnée peut-elle être encodée et traitée par le système nerveux? Les neurones émettent d'occasionnelles impulsions espacées de longs silences. Ce "code morse" constitue une réponse à l'activité groupée d'une population de neurones pre-synaptiques. On peut comprendre le fonctionnement des neurones si on arrive à une description précise de leurs dynamiques. Les résultats décrits dans la présente thèse établissent les éléments nécessaires à une description quantitative des neurones. L'ensemble du travail peut être rassemblé sous le thème de l'adaptation. L'adaptation neuronale joue un rôle important dans la traduction du stimulus en une série d'impulsions. Les travaux réunis ici décrivent comment l'adaptation est provoquée par chaque impulsion neuronale de manière stéréotypée. Ce schéma d'adaptation est ensuite mesuré expérimentalement dans les trois principales classes de neurones du cortex. La manière dont différentes structures d'adaptation contribuent à la diversité neuronale est analysée en détail. Des modèles simples capables d'adaptation peuvent prédire avec grande précision la réponse de neurones soumis à des stimuli complexes. Je résume les plus récents résultats portant sur les méthodes permettant de discriminer les modèles selon leur pertinence et ainsi d'obtenir un modèle valide et testé. Ce modèle est ensuite analysé pour comprendre comment des populations peuvent encoder un signal variable et comment ce signal peut être décodé de l'activité d'une population. Finalement, deux travaux en cours sont discutés : l'adjonction d'un compartiment dendritique et l'adaptation sur de longues échelles de temps.

Mots Clefs Modèles de neurones à impulsions - Estimation de paramètres - Métrique de séries d'impulsions - Codage neuronal - Adaptation neuronale

Foreword

This thesis summarizes the work I have done in Ecole Polytechnique Federale de Lausanne from March 2007 to August 2011. I have compiled texts from book chapters and peer-reviewed articles that are – at the time of writing this foreword – either published, accepted and awaiting publication, under review, or close to submission.

Although the present thesis compiles articles, there are three factors that allows it to stand alone. First, I have included here only the publications for which I have written most of the text and contributed substantially to the conceptual and technical work. Second, the references of each article is collected in a single bibliography section at the end of the thesis. Last, each chapter can be connected to the central theme of spike-frequency adaptation in single neuron models. One important difference with a proper stand-alone thesis is the significant amount of repetition in the different articles. For instance, the sections describing the methods are almost the same in Chapter 4 and Chapter 5. I have tried to indicate the repetitions in footnotes in order to relieve the reader. Another difference with the stand-alone thesis consists in having cross-references between chapters that are replaced by a reference to the associated journal publication and a footnote to indicate the corresponding chapter.

The first chapter is intended to be a general introduction to simple neuron models with a special emphasis on refractory and adaptation processes. This chapter will be part of a textbook called *Computational Systems Biology* edited by Nicolas Le Novère [Naud and Gerstner, 2012a]. Therein I introduce two models which will be of central importance throughout the thesis: the Adaptive Exponential Integrate and Fire (AdEx) and the Spike-Response Model (SRM).

The second chapter of the thesis explores the versatility of the AdEx – or aEIF as it was called at the time – model. This article appeared in *Biological Cybernetics* [Naud et al., 2008] and shows how spike-triggered and sub-threshold adaptation contribute to generate multiple firing patterns.

At the time I started my thesis, the AdEx was considered state-of-the-art in simple neuron models. However the reasons for its popularity were mainly theoretical and experimental evidences were lacking to prove its quantitative validity. The need for quantitative assessment of neuron models was the main motivation for organizing the spike-time prediction competition in 2007, 2008 and 2009. Chapter 3 of the thesis is a summary of what we have learned from the

Foreword

spike-timing prediction competition. This chapter is also going to be published as a chapter in a book on spike timing edited by Jonathan Victor and Patricia Di Lorenzo [Naud and Gerstner, 2012b]. This chapter summarizes the content of three previous publications that I have not included in the thesis [Jolivet et al., 2008a,b, Gerstner and Naud, 2009].

In the last edition of the competition (2009) we noticed that the participants using stochastic models were systematically low-ranking. As one question led to another, we were forced to study in greater details the way we evaluated the spike-timing prediction, and found that our previous method was bearing a bias detrimental to stochastic models. The details concerning the similarity measures are contained in Chapter 4 and corresponds to a publication in *Neural Computation* [Naud et al., To appear; 2012]. This chapter does not directly relate to adaptation. It is, nevertheless, an important technical aspect needed to find the optimal single-neuron model.

In the chapters that follow, we use the SRM. This model was both more easily fitted to real data and more easily treated analytically. Katja Miller and myself have shown that the SRM can reproduce all the firing patterns considered in Chapter 2 (not published). Mensi et al. [2011] not included in the thesis describes how the SRM relates with the AdEx with noise. This article makes the bridge between the AdEx of Chapter 2 and the SRM of Chapter 5 and Chapter 6.

Chapter 5 analyzes the adaptation profiles found in three main cortical neuron types. This article was submitted to the *Journal of Neurophysiology* and is currently under review.

Chapter 6 is a theoretical investigation on the role of adaptation in populations of SRM neurons. This article is not finished but it is soon to be submitted to *Plos Computational Biology*.

Finally, Chapter 7 condenses two lines of investigation in progress. In the first part of Chapter 7, I briefly report on the extension of adapting dynamics to longer time scales. In the second part, I report an extension of the simple neuron models to non-linear dendritic integration. This is based on a poster presented at SAND5 [Naud et al., 2010] and should be transformed into an article.

All the work reported here was made in collaboration with Wulfram Gerstner and colleagues from the EPFL. My specific contribution to each of the chapters is outlined in the Contribution section at the end of the thesis.

Lausanne, the 6th of September 2011

R. N.

Contents

Acknowledgements	v
Abstract (English/Français)	vii
Foreword	ix
List of figures	xiv
List of tables	xvii
1 Introduction: The Performance (and limits) of Simple Neuron Models	1
1.1 Basic Threshold Models	1
1.1.1 The Perfect Integrate-and-Fire Model	2
1.1.2 The Leaky Integrate-and-Fire Model	3
1.2 Refractoriness and Adaptation	5
1.2.1 Spike-Triggered Current	7
1.2.2 Moving Threshold	8
1.3 Linearized Subthreshold Currents	9
1.4 Nonlinear Integrate-and-Fire Models	11
1.4.1 The Exponential Integrate-and-Fire Model	12
1.5 Unifying Perspectives	13
1.5.1 The Adaptive Exponential Integrate-and-Fire Model	14
1.5.2 Integrated Models	15
1.6 Noise	16
1.6.1 Synaptic Noise	17
1.6.2 Electrical Noise	18
1.6.3 Generalized Linear Models	18
1.7 Advantages of Simple Models	20
1.7.1 Variety of Firing Patterns	20
1.7.2 Spike-Time Prediction	22
1.7.3 Ease of Mathematical Analysis	23
1.8 Limits	25

2	Firing patterns in the adaptive exponential integrate-and-fire model	27
2.1	Introduction	27
2.2	Adaptive Exponential Integrate-and-Fire	28
2.3	Multiple Firing Patterns	31
2.4	Parameter Space	37
2.5	Boundaries in Parameter Space	40
2.6	Comparison with Cortical Neurons	42
2.7	Discussion	44
3	Can We Predict Every Spike?	51
3.1	What is a good stimulus to probe neurons?	52
3.2	How can we measure spike timing precision and reliability?	54
3.3	What are good neuron models?	57
3.4	Are all neurons predictable?	59
3.5	Conclusion	61
4	Improved Similarity Measures for Small Sets of Spike Trains	63
4.1	Introduction	63
4.2	The Space of Spike Trains	67
4.2.1	Spike Train Algebra	67
4.2.2	Scalar Products of Spike Trains	68
4.2.3	Ensemble of Stochastic Spike Trains	69
4.2.4	The PSTH as an Empirical Estimate	69
4.2.5	Intrinsic Variability and Reliability are Measures of the Variance	70
4.2.6	Similarity Measures Derived From an Inner Product	73
4.2.7	Similarity Measures loosely Related to an Inner Product	75
4.2.8	Comparing Sets of Spike Trains: The Distance between two PSTHs	77
4.3	Analytical Results	78
4.3.1	Train by Train Comparison has a Deterministic Bias	78
4.3.2	Sample Bias in the Norm of the Population Activity	79
4.3.3	Sample Bias in Distance and Angle Measures	81
4.3.4	Applications to Measures Loosely Related to an Inner Product	82
4.4	Examples	87
4.4.1	A Framework for Comparing The Different Measures	89
4.4.2	Numerical Methods	90
4.4.3	Case Study I – Match based on PSTH or Population Activity	90
4.4.4	Case Study II – Measures Loosely Related to an Inner Product	93
4.4.5	Case Study III - Phase code	94
4.4.6	Case Study IV - Spike Response Model	96
4.4.7	Case Study V - Classification	100
4.5	Fitting Cortical Neurons	102
4.5.1	Experimental details	102
4.5.2	Neural Model and Fitting	102

4.5.3	Deterministic Bias Penalizes the Maximum Likelihood Solution	103
4.5.4	New Ranking in INCF Competition	105
4.6	Conclusions	106
5	Parameter Extraction and Classification of Three Neuron Types Reveals two Different Adaptation Mechanisms	111
5.1	Introduction	111
5.2	Material and Methods	114
5.2.1	Model Dynamics	114
5.2.2	<i>In vitro</i> Two-photon Microscopy and Whole-cell Recordings	117
5.2.3	Stimulation Protocol: Synaptic-like current	118
5.2.4	Performance Measurements	119
5.2.5	Fitting Procedure	120
5.3	Results	124
5.3.1	Efficiency and Accuracy of the Fitting Method on Surrogate Data	124
5.3.2	Accuracy of Fitted Models on FS, NFS and Exc	125
5.3.3	Essential features for Subthreshold Voltage Prediction	127
5.3.4	Essential Features for Spike Time Prediction	131
5.3.5	Cell-type Classification	132
5.4	Discussion	133
5.4.1	Automatic Fitting Method	133
5.4.2	Effective Spike-triggered Adaptation	135
5.4.3	Interpretation of Model Parameters	136
5.4.4	Classification	137
5.4.5	How Good is Good?	138
6	Coding and Decoding with Adapting Neurons: a Population Approach	145
6.1	Introduction	145
6.2	Results	147
6.2.1	Encoding the Population Activity	147
6.2.2	Decoding the Population Activity	152
6.2.3	How the Effective Spike After-Potential Shapes the Transfer Function	152
6.3	Discussion	157
6.4	Materials and Methods	159
6.5	Appendix: Derivation of Eq. 6.2	159
7	Work in Progress	161
7.1	Power-law Adaptation	161
7.2	Active Dendrites	162
7.2.1	Description of the Model	164
7.2.2	Experimental Protocol	164
7.2.3	Fitting Methods	166
7.2.4	Results	166

Contents

Contribution	171
Bibliography	192
Curriculum Vitae	193

List of Figures

1.1	Positively and negatively charged ions are distributed inside and outside the cell.	2
1.2	The defining responses of the LIF model.	3
1.3	The filter κ can be measured in real neurons [Jolivet et al., 2006].	5
1.4	Comparing two simple mechanism of spike-triggered adaptation.	6
1.5	The responses of the LIF with a single linearized current.	9
1.6	The shape of the filter κ in the presence of resonance.	10
1.7	Where is the voltage threshold?	12
1.8	Experimental measurement of the nonlinearity for spike initiation.	14
1.9	Generalizations of the LIF include either refractoriness, adaptation, linearized currents, or smooth spike initiation.	15
1.10	Membrane potential recorded from four repetition of the same stimulus.	16
1.11	Multiple firing patterns are reproduced by merely tuning the parameters of a simple threshold model.	21
1.12	Overlayed traces of an AdEx model (red) and a fast-spiking interneuron (blue).	22
1.13	Initially bursting (A) and regular bursting (B) trajectories in both phase plane and as a time-series.	24
2.1	Phase plane representation of a step current injected in an AdEx model.	30
2.2	Types of excitability depend on the point of reset.	32
2.3	Phase diagram and associated traces illustrating two types of tonic spiking.	33
2.4	Phase plane representation of 8 firing patterns.	35
2.5	Irregular firing is chaos in the AdEx.	38
2.6	Parameter space exploration of the four bifurcation parameters.	47
2.7	Phase diagram and parameter space for a piecewise linear V -nullcline in the limit of separation of timescales.	48
2.8	Comparison of the AdEx with three types of cortical neurons on step current injections.	49
3.1	Schematic representation of the spike-timing prediction challenge.	53
3.2	Counting coincident spikes for the computation of the coincidence rate Γ_{nm} .	54
3.3	Predicting fast-spiking GABAergic neurons and pyramidal neurons from in 5 from the 2009 competition.	60
4.1	Similarity of Spike Trains Does Not Predict Similarity of Distributions.	66

List of Figures

4.2	Illustrating the Vector Space of Spike Trains.	71
4.3	Raster plots and firing rates illustrating the different case studies.	83
4.4	Discriminability profiles for three similarity measures of spike trains seen as vectors.	84
4.5	Discriminability profiles for three standard similarity measures.	88
4.6	Raster plots and firing rates illustrating the phase coding	94
4.7	Phase coding: Discriminability profiles.	95
4.8	The similarity measure M_D^* increases discriminability for all sample sizes greater than one.	96
4.9	Illustrating the effect of the model parameters on the PSTH.	98
4.10	The discriminability of the Victor and Purpura measure on surrogate data.	99
4.11	Classification of spike trains coming from different point processes is improved by using D_{spk}^* instead of the measure of Eq. 4.54 d	101
4.12	Fitting cortical neurons.	104
5.1	Stimulation protocol and dynamics of a IF model.	118
5.2	Fitting protocol.	122
5.3	Assessment of the fitting procedure by fitting a model to a model.	126
5.4	Examples of voltage traces and spike train predictions for the 3 neuron classes.	127
5.5	Instantaneous firing frequency of models as a function of the intensity of a step current, (f-I curve).	129
5.6	Essential Model Features for Subthreshold Voltage Prediction in GABAergic FS, GABAergic NFS and Exc.	139
5.7	Conductance vs current-based adaptation.	140
5.8	Essential Model Features for Spike Time Prediction in GABAergic FS, GABAergic NFS and Exc.	141
5.9	Principal components of the 3 studied cell types.	142
5.10	Effective spike-triggered adaptation for the three neuron classes.	143
6.1	Encoding the stimulus in adapting populations.	148
6.2	Decoding the stimulus from the population activity.	153
6.3	Validity of the analytical transfer function.	154
6.4	The relation between the shape of the spike-triggered adaptation and the transfer function.	156
7.1	Dynamic threshold and spike-triggered current on long time scales.	162
7.2	Schematic representation of the two-compartment model.	163
7.3	The two-compartment model fits qualitatively and quantitatively the electrophysiological recordings.	167

7.4	Fitted kernels of the two-compartment model. A The kernel for spike-triggered adaptation is negative and increases monotonically between 6 and 600 ms. B The back-propagating current reaching the dendrites is a short (2ms) and strong (900 pA) pulse. C The convolution kernel linking the current injected in the dendrite to the current reaching the soma. D The convolution kernel linking the current injected in the soma to the current reaching the dendrite.	168
7.5	The model reproduces the qualitative features of active dendrites reported in Larkum et al. [1999] and Larkum et al. [2004].	169

List of Tables

2.1	Parameters and cost for fits shown in Fig. 2.8 and for firing pattern example shown in Fig. 2.4.	43
4.1	Summary of spike train similarity measures.	74
4.2	Summary of similarity measures for comparing groups of spike trains.	85
4.3	Columns on the left: Evaluation of CF2 normalized by the intrinsic reliability of the data, as published on the INCF web page. Right: re-ordering with the bias-corrected measure CF2*.	106
5.1	Evaluation of the fitting procedure on different common types of stimulation.	125
7.1	List of Parameters and their fitted value for the two-compartment model.	165

1 Introduction: The Performance (and limits) of Simple Neuron Models

The study of neuronal populations with regards to coding, computation and learning relies on its primary building bloc: the single neuron. ¹ Describing the activity of single neurons can be done by mathematical models of various complexity. On one hand there are complex biophysical models and on the other hand there are the simpler integrate-and-fire models. In order to relate with higher functionalities such as computation or coding, it is not necessary to model all the spatio-temporal details of ionic flow and protein interactions. There is a level of description that is simple, that bridges the gap to higher functionalities, and that is sufficiently complete to match real neurons. In this chapter we start with the integrate-and-fire model and then consider a set of enhancements so as to model the behaviour of multiple types of neurons.

The formalism considered here takes a stimulating current $I(t)$ as an input, to which the neuron responds with a voltage trace $V(t)$, which contain multiple spikes. The input current can be injected experimentally *in vitro*. In a living brain, the input comes from synapses, gap junctions or voltage-dependent ion channels of the neuron's membrane.

1.1 Basic Threshold Models

Nerve cells communicate by action potentials – also called spikes. Each neuron gathers input from thousands of synapses to decide when to produce a spike. In the absence of input, the neurons would stay at a resting membrane potential around -70 mV. Most neurons will fire a spike when their membrane potential reaches a value around -55 to -50 mV. Action potential firing can be considered an all-or-nothing event. These action potentials are very stereotypical. This suggests that spikes can be replaced by unitary events that are generated by a threshold-crossing process.

¹Text copied from Naud, R. and Gerstner, W. to appear in textbook *Computational System Biology* edited by Nicolas Le Novère, Springer (full citation in the References).




Figure 1.1: Positively and negatively charged ions are distributed inside and outside the cell. Current $I(t)$ entering the cell will modify the difference in electric potential between the exterior and the interior (V). The dynamics of the LIF correspond to a RC-circuit composed of a conductance (g_L) in parallel with a capacitance (C). The electric supply corresponds to the resting potential (E_L).

1.1.1 The Perfect Integrate-and-Fire Model

Fundamentally, we can say that a single neuron accumulates electric charge supplied by an input current $I(t)$ on its membrane. This translates into an by increase of the membrane potential $V(t)$. When the voltage hits the threshold (V_T), a spike is said to be emitted and the voltage is reset to V_r . Mathematically we write:

$$\frac{dV}{dt} = \frac{1}{C} I(t) \tag{1.1}$$

$$\text{when } V(t) > V_T \text{ then } V(t) \rightarrow V_r. \tag{1.2}$$

Here, C is the total capacitance of the membrane. This equation is the first Kirchoff law for an impermeable membrane: the current injected can only load the capacitance. The greater the capacitance the greater the amount of current required to increase the potential by a given amount.

This system is called the (perfect or non-leaky) integrate-and-fire (IF) model. Solving the first-order differential equation shows the integration process explicitly; after a previous spike at t_0 the voltage at time t is given by:

$$V(t) = V_r + \frac{1}{C} \int_{t_0}^t I(s) ds. \tag{1.3}$$

Here a pulse of current will never be forgotten. In other words one could inject a small pulse




Figure 1.2: The defining responses of the LIF model. A short but strong pulse will make a marked increase in potential which will then decay exponentially. A subthreshold step of current leads to exponential relaxation to the steady-state voltage, and to an exponential relaxation back to resting potential after the end of the step. A supra-threshold step of current leads to tonic firing. If a sinusoidal wave of increasing frequency is injected in the model, only the lowest frequencies will respond largely. The higher frequencies will be attenuated.

of current every hour and their repercussions on the voltage will cumulate to eventually make the model neuron fire. This conflicts with the behaviour of real neurons, which are not perfect integrators but “leaky” ones. In fact, the IF model is used here for didactic purposes only because it summarizes the central idea: integrate and fire.

1.1.2 The Leaky Integrate-and-Fire Model

Real neuronal membranes are leaky. Ions can diffuse through the neuronal membrane. The membrane of neurons can be seen as providing a limited electrical conductance (g_L) for charges crossing the cellular membrane. The difference in electric potential at equilibrium depends on the local concentration of ions and is often called the equilibrium (or resting) potential (E_0). This additional feature leads to the more realistic Leaky-Integrate-and-Fire (LIF) model [Lapicque, 1907]:

$$C \frac{dV}{dt} = -g_L(V - E_0) + I(t) \tag{1.4}$$

$$\text{when } V(t) > V_T \text{ then } V(t) \rightarrow V_r. \tag{1.5}$$

Chapter 1. Introduction: The Performance (and limits) of Simple Neuron Models

Again, this is the Kirchoff law for conservation of charge. The current injected can either leak out or accumulate on the membrane. The effect of a short current pulse will cause a transient increase in voltage. This can be seen by looking at the solution of the linear differential equation given a previous spike at time \hat{t}_0 :

$$V(t) = E_0 + \eta_r(t - \hat{t}_0) + \int_0^{t - \hat{t}_0} \kappa(s) I(t - s) ds \quad (1.6)$$

$$\eta_r = (V_r - E_0) e^{-t/\tau} \Theta(t) \quad (1.7)$$

$$\kappa(t) = \frac{1}{C} e^{-t/\tau} \Theta(t) \quad (1.8)$$

where $\Theta(t)$ is the Heaviside function and $\tau = C/g_L$ is the membrane time constant. In Eq. 1.6, three terms the voltage. The first term is the equilibrium potential. The second term is the effect of voltage reset which acts as an initial condition for the integration of the differential equation and gives rise to the function η_r . Note that far away from the last spike (\hat{t}_0) this term vanishes. The last term – made of the convolution of the filter $\kappa(t)$ with the current – is the influence of input current on the voltage. We see that the voltage is integrating the current but the current at an earlier time has a smaller effect than current at later times. The membrane time constant of real neurons can vary between 10 and 50 ms. The theory of signal processing tells us that the membrane acts as a low-pass filter of the current (Fig. 1.2). In fact, input current fluctuating slowly is more efficient at driving the voltage than current fluctuating very rapidly.

There is another way to implement the reset. Instead resetting to fixed value, we assume that whenever the voltage equals V_T , there is a sudden decrease in voltage to V_r caused by a short negative pulse of current, which reflects the fact that the membrane loses charge when the neuron sends out a spike. We write the LIF equations differently, using the Dirac delta function ($\delta(t)$):

$$C \frac{dV}{dt} = -g_L(V - E_0) + I(t) - C(V_T - V_r) \sum_i \delta(t - \hat{t}_i) \quad (1.9)$$

where the sum runs on all spike times $\hat{t}_i \in \{\hat{t}\}$, defined as the times where $V(t) = V_T$. The integrated form is now:

$$V(t) = E_0 + \sum_i \eta_a(t - \hat{t}_i) + \int_0^\infty \kappa(s) I(t - s) ds \quad (1.10)$$

$$\eta_a(t) = -(V_T - V_r) e^{-t/\tau} \Theta(t). \quad (1.11)$$

The two different ways to implement the reset yield slightly different formalisms. Even though




Figure 1.3: The filter κ can be measured in real neurons [Jolivet et al., 2006]. Here the shape of the filter is shown as measured in the soma of a pyramidal neuron of the layer 5 of the cortex (gray circles). Data points before 5 ms are not shown because they bear a heavy artefact due to the electrode used for recording. In black is a fit of Eq. 1.8 with $C = 408$ pF and $\tau = 11.4$ ms.

including the leak made the mathematical neuron model more realistic, it is not sufficient to describe experiments with real neurons. We also need a process to account for adaptation [Benda and Herz, 2003, Rauch et al., 2003, Jolivet et al., 2006], and this is the topic of the next section.

1.2 Refractoriness and Adaptation

Refractoriness prevents a second spike immediately after a first one was emitted. One can distinguish between an absolute and a relative refractory period. During the absolute refractory period, no spike can be emitted, no matter the strength of the stimulus. The duration of the spike is often taken as the absolute refractory period since it is impossible to emit a spike while one is being generated. During the relative refractory period it is possible to fire a spike, but a stronger stimulus is required. In this case the current required depends on the time since the last spike. Manifestly, the absolute refractory period always precedes the relative refractory period, and the absolute refractory period can be seen as a very strong relative refractory period.

Spike-frequency adaptation, on the other hand, is the phenomenon whereby a constant stimulus gives rise to a firing frequency that slows down in time (Fig. 1.4). Here it is the cumulative effect of previous spikes that prevents further spiking. In other words: the more a neuron has fired in the recent past, the less it is likely to spike again. How long can this history dependence be? Multiple studies have pointed out that spikes emitted one or even ten seconds




Figure 1.4: Comparing two simple mechanisms of spike-triggered adaptation. A step current (top) is injected in a model with spike-triggered hyperpolarizing current (left) or in a model with moving threshold (right). Both models can produce the same firing patterns, but the voltage trace differs qualitatively.

earlier can still reduce the instantaneous firing rate [La Camera et al., 2004, Lundstrom et al., 2008].

Refractoriness and adaptation are two similar but distinct effects, and we need to define the difference in precise terms. Although refractoriness mostly affects the earliest times after a spike and adaptation the latest times, this distinction is not adequate: spike-triggered currents can cumulate even at small time scales. It is more convenient to distinguish the two processes based on the history-dependence whereby refractoriness prevents further spiking as a function of time since the last spike only, while adaptation implies a dependency on the time of all the previous spikes. In other words, adaptation is a refractoriness that cumulates over the spikes. Equivalently, refractoriness can be distinguished from adaptation by the type of reset: a fixed reset like in Eq. 1.5 leads to dependency on the previous spike only (see Eq. 1.7) and hence to no adaptation. A relative reset allows the effect of multiple spikes to cumulate and can lead to spike-frequency adaptation.

In terms that are compatible with threshold models, both refractoriness and adaptation can be mediated either by a transient increase of the threshold after each spike or by the transient increase of a hyperpolarizing current. These will be discussed next.

1.2.1 Spike-Triggered Current

Some ion channels seem to be essential to the action potential but influence very weakly the subthreshold dynamics. Some other ion channels play no role in defining the shape of the spike, are partially activated during a spike, and their level of activation decays exponentially after the spike. These are ion channels that can mediate adaptation or refractoriness. Such ion channels have voltage-dependent sensitivity to the membrane potential: at high voltage they rapidly open, at low voltage they slowly close. Since a spike is a short period of high voltage this creates a short jump in the level of activation which will slowly decay after the spike. Such a situation can be implemented in the LIF by adding an hyperpolarizing current w which is incremented by b each time there is a spike and otherwise decays towards zero with a time constant τ_w [Baldissera et al., 1976, Izhikevich and Desai, 2003, Benda and Herz, 2003] :

$$C \frac{dV}{dt} = -g_L(V - E_0) - w + I(t) - C(V_T - V_r) \sum_i \delta(t - \hat{t}_i) \quad (1.12)$$

$$\tau_w \frac{dw}{dt} = -w + b\tau_w \sum_i \delta(t - \hat{t}_i). \quad (1.13)$$

$$(1.14)$$

The above system of equations is a simple mathematical model for a neuron with spike-frequency adaptation. The current w is triggered by the spikes and will move the membrane potential away from the threshold when $b < 0$. This equation can be integrated to yield:

$$V(t) = E_0 + \sum_i \eta_a(t - \hat{t}_i) + \int_0^\infty \kappa(s) I(t - s) ds \quad (1.15)$$

$$\eta_a(t) = \frac{b\tau_w}{C(\tau_w - \tau)} [e^{-t/\tau_w} - e^{-t/\tau}] \Theta(t) - (V_T - V_r) e^{-t/\tau} \Theta(t) \quad (1.16)$$

where $\kappa(t)$ is the same as in Eq. 1.9. The spike-triggered current that cumulates over the spikes is reflected in a stereotypical change in voltage η_a that can also cumulate over the spikes. Such a spike-triggered current can also make refractoriness if we replace its cumulative reset by a fixed reset :

$$\tau_w \frac{dw}{dt} = -w + (b - w_-)\tau_w \sum_i \delta(t - \hat{t}_i) \quad (1.17)$$

where w_- is w just before time t so that at each time instead of incrementing by b , we increment to b . In this case the amount w of refractory current depends only on the time since

the last spike. Integration of this yields:

$$\eta_r(t) = \frac{b\tau\tau_w}{C(\tau_w - \tau)} [e^{-t/\tau_w} - e^{-t/\tau}] \Theta(t) \quad (1.18)$$

The shape of the spike after potential can be mediated by a handful of ion channels. Likely candidates for mediating a spike triggered current of the type described above must have a slow to medium activation at supra-threshold potentials and a very slow inactivation or de-activation at subthreshold potentials. An action potential will then induce a small increase in the number of open channels which could cumulate over the spikes. The time constant of the hyperpolarizing current τ_w relates to the time constant for the closure of the ion channels at subthreshold potentials. Typical examples are: slow potassium current I_K with de-activation time constant around 30-40 ms [Korngreen and Sakmann, 2000], muscarinic potassium current, I_M , with de-activation time constant around 30 to 100 ms [Passmore et al., 2003] or the calcium-dependent potassium current $I_{K[Ca]}$ which can have a time constant in the order of multiple seconds [Schwindt et al., 1989]. Finally, active dendritic processes can also induce current to flow back into the somatic compartment after each spike [Doiron et al., 2007]. In this case the current is depolarizing rather than hyperpolarizing, leading to facilitation rather than adaptation.

1.2.2 Moving Threshold

Spike-triggered currents are not the only way to implement refractoriness and adaptation in neuron models. Multiple experiments have shown that the effective threshold of neurons is dynamic [Hill, 1936, Fuortes and Mantegazzini, 1962, Azouz and Gray, 2000, 2003, Badel et al., 2008b]. If instead of adding a spike-triggered current we let the threshold V_T be dynamic: the threshold can increase by δV_T each time there is a spike and decay exponentially with time constant τ_T to the minimum threshold $V_T^{(0)}$. This is summarized by the supplementary equations:

$$\tau_T \frac{dV_T}{dt} = -(V_T - V_T^{(0)}) \quad (1.19)$$

$$\text{when } V(t) \geq V_T(t) \text{ then } V(t) \rightarrow V_T \quad (1.20)$$

$$\text{and } V_T(t) \rightarrow V_T(t) + \delta V_T. \quad (1.21)$$

Again, the moving threshold can implement adaptation (as with Eq.'s 1.19-1.21 above) or refractoriness if we replace the relative reset by a fixed reset: $V_T(t) \rightarrow V_T^{(0)} + \delta V_T$.

It is often possible to find parameters for which a model with a moving threshold will yield the same spikes times than a model with a spike-triggered current. Indeed moving the membrane potential away from the threshold with a spike-triggered current is equivalent to moving the




Figure 1.5: The responses of the LIF with a single linearized current. A short but strong pulse will generate a jump in potential which will then relaxes with an undershoot. A subthreshold step of current leads to a characteristic overshoot after the onset and an undershoot after the offset of the current. A supra-threshold step of current leads to firing with spike-frequency adaptation. If a sinusoidal wave of increasing frequency is injected in the model, the lowest frequencies and the highest frequencies will be attenuated. In-between frequencies will yield the greatest response.

threshold away from the membrane potential. In particular, when only the spike times are predicted, we can put Eq.'s 1.19-1.21 in the form of Eq. 1.15 by keeping a fixed threshold and adding to $\eta_a(t)$:

$$\delta V_T e^{-t/\tau_T} \Theta(t). \quad (1.22)$$

It is not yet clear which biophysical mechanisms are responsible for moving thresholds. One likely candidate is the sodium channel inactivation [Azouz and Gray, 2000]. An increase in sodium channels inactivation can increase the voltage threshold for spike initiation. Inactivated sodium channels de-inactivate with a time constant of 2-6 ms [Huguenard et al., 1988]. Furthermore, it is believed that sodium channels can de-inactivate on time scales as long as multiple seconds [Baker and Bostock, 1998].

1.3 Linearized Subthreshold Currents

There are ion channels influencing principally the shape of the spike, some the refractoriness, and others the adaptation of neurons. However, there are also channels whose dynamics depends and influences only the subthreshold potentials. An example is the hyperpolarization activated cation current I_h which start to activate around -70 mV. The first-order effect of such




Figure 1.6: The shape of the filter κ in the presence of resonance. Here the shape of the filter is shown as measured in the apical pyramidal neuron of the layer 5 of the cortex (gray points). Data points before 5 ms are not shown because they bear a heavy artefact from due to the measurement electrode. In black is shown a fit of Eq. 1.23-1.24 ($a = 13.6$ nS, $g_L = 35.0$ nS, $C = 168$ pF, $\tau_w = 15.5$ ms). Data a courtesy of Brice Bathellier.

currents can be added to the LIF equations (see [Mauro et al., 1970] for details on the Taylor expansion):

$$C \frac{dV}{dt} = -g_L(V - E_0) - w + I(t) \quad (1.23)$$

$$\tau_w \frac{dw}{dt} = a(V - E_w) - w \quad (1.24)$$

$$\text{if } V(t) > V_T \text{ then } V(t) \rightarrow V_r. \quad (1.25)$$

Here a regulates the magnitude of the subthreshold current and τ_w rules the time constant of the coupling with the voltage. E_w should correspond to the average voltage, we will assume that $E_w = E_0$ for the following treatments. When a is negative, w is said to cause subthreshold facilitation. The response properties will resemble the LIF as (in Fig. 1.2 and 1.3) but with a longer impulse-response function. When a is positive, w is said to generate subthreshold adaptation. For a sufficiently strong positive a we see the emergence of resonance as shown in Fig. 1.5. This model can be called the resonate and fire (RF; [Izhikevich, 2001, Richardson et al., 2003]).

This system of equations can be mapped to the equations of a damped oscillator with a driving force. It is a well-studied system that comes from the dynamics of a mass hanging on a spring in a viscous medium. We know that this system has three dynamical regimes:

- $4C\tau_w(g_L + a) < (g_L\tau_w + C)^2$ overdamped,
- $4C\tau_w(g_L + a) = (g_L\tau_w + C)^2$ critically damped,
- $4C\tau_w(g_L + a) > (g_L\tau_w + C)^2$ underdamped.

Overdamped and critically damped systems have no resonating frequency. It is only when the system is underdamped that resonance appears. Such resonance is seen in multiple types of neurons, typically in some cortical interneurons [Markram et al., 2004], in mesencephalic V neurons [Wu et al., 2005] and in the apical dendrites of pyramidal neurons [Cook et al., 2007].

What are the main characteristics of a resonating membrane? Contrasting with the standard LIF, the response to a current pulse in Fig. 1.5 and 1.6 is not a single exponential. Instead, the voltage makes a short undershoot before relaxing to the resting potential. Similarly, when the input is a step current, there is a substantial overshoot at the onset and undershoot after the offset of the current. The neuron model resonates around a characteristic frequency for which it will respond with maximal amplitude:

$$\omega = \sqrt{\frac{g_L + a}{C\tau_w} - \frac{(g_L\tau_w + C)^2}{4C^2\tau_w^2}}. \quad (1.26)$$

Resonating membranes are bandpass filters as we can see from the response to a sinusoidal wave of increasing frequency (Fig 1.5). The shape of the filter can be written as:

$$\kappa(t) = \exp(-t/\tau_\omega) \left[\frac{1}{C} \cos \omega t + \frac{\tau_\omega \tau_w + 1}{g_L + a} \sin \omega t \right] \Theta(t) \quad (1.27)$$

where the decay time constant is:

$$\tau_\omega = \frac{2\tau_w\tau}{\tau_w + \tau}. \quad (1.28)$$

Figure 1.6 shows the shape of the filter, as measured in a neuron with resonance.

1.4 Nonlinear Integrate-and-Fire Models

Do neurons really have a voltage threshold? Imagine that a neuron was to receive a stimulus that brings its membrane potential to a value which triggers the spike but then the stimulus is stopped. The membrane potential would continue to increase even in the absence of stimulus, and produce the action potential. Can we say that a spike is produced whenever the membrane potential reaches this threshold voltage? No. At the earliest times of the action potential, a negative current can veto the spike even though the membrane potential was above the threshold for spike initiation. Another example that makes the conceptualization of a threshold dubious is shown in Fig. 1.7. Here the voltage threshold measured from a current pulse is significantly different from the voltage threshold measured with a current step.




Figure 1.7: Where is the voltage threshold? Current pulses and current steps show different voltage thresholds that can not be accounted for by a moving threshold.

Real neurons have sodium ion channels that gradually open as a function of the membrane potential and time. If the channel is open, positive sodium ions flow into the cell, which increases the membrane potential even further. This positive feedback is responsible for the upswing of the action potential. Although this strong positive feedback is hard to stop, it can be stopped by a sufficiently strong hyperpolarizing current, thus allowing the membrane potential to increase above the activation threshold of the sodium current.

Sodium ion channels responsible for the upswing of the action potential react very fast. So fast that the time it takes to reach their voltage-dependent level of activation is negligible. Therefore these channels can be seen as currents with a magnitude depending nonlinearly on the membrane potential, which suggests that a simplified description of spike generation should be possible. This section explores the LIF augmented with a nonlinear term for smooth spike initiation.

1.4.1 The Exponential Integrate-and-Fire Model

Let us assume that the transmembrane current to is some function of V , so that the membrane dynamics is an equation of the type:

$$C \frac{dV}{dt} = F(V) + I(t), \quad (1.29)$$

where $F(V)$ is the current flowing through the membrane. For the perfect IF it is zero ($F(V) = 0$), for the LIF it is linear with a negative slope ($F(V) = -g_L(V - E_0)$). We can speculate on the shape of the non-linearity. The simplest non-linearity would arguably be the quadratic : $F(V) = -g_L(V - E_0)(V - V_T)$ (Quadratic Integrate and Fire, QIF; [Latham et al., 2000]). However

this implies that the dynamics at hyperpolarized potentials is non-linear, which conflicts with experimental observations. Other possible models could be made with cubic or even quartic functions of V [Touboul and Brette, 2008]. An equally simple non-linearity is the exponential function:

$$F(V) = -g_L(V - E_0) + g_L\Delta_T \exp\left(\frac{V - V_T}{\Delta_T}\right) \quad (1.30)$$

where Δ_T is called the slope factor that regulates the sharpness of the spike initiation. The first term on the right hand side of Eq. 1.30 is identical to that of the LIF. The second one accounts for nonlinearity in spike initiation.

The Exponential Integrate-and-Fire (EIF; [Fourcaud-Trocme et al., 2003]) model integrates the current according to Eq. 1.29 and 1.30 and resets the dynamics to V_r (*i.e.* produces a spike) once the simulated potential reaches a value θ . The exact value of θ does not matter, as long as $\theta \gg V_T + \Delta_T$. As in the LIF, we have to reset the dynamics once we have detected a spike. The value at which we stop the numerical integration should not be confused with the threshold for spike initiation. We reset the dynamics once we are sure the spike *has been* initiated. This can be at a membrane potential of 0, 10 mV or infinity. In fact, because of the exponential nonlinearity V goes from $V_T + \Delta_T$ to infinity in a negligible amount of time.

Which $F(V)$ is the best? This can be measured experimentally (provided that we can estimate membrane capacitance beforehand [Badel et al., 2008a,b]). Fig. 1.8 shows the function $F(V)$ as measured in pyramidal neurons of the cortex. Choosing $F(V)$ as a linear plus exponential allow a good fit to the experimental data. Similar curves are observed in neocortex interneurons [Badel et al., 2008a].

1.5 Unifying Perspectives

We have seen that the LIF can be augmented with mechanisms for refractoriness, adaptation, subthreshold resonances and smooth spike initiation. Models combining these features can be classified in two categories depending on the presence of a non-linearity. This is because the linear system of equations can be integrated analytically, whereas integration is generally not possible for the non-linear systems of equations. When integration can be carried out, the dynamics can be studied with signal processing theory. When this is not possible, the dynamical system is scrutinized with bifurcation theory.




Figure 1.8: Experimental measurement of the nonlinearity for spike initiation. Example of the measured $F(V)$ in black circles for a layer 5 pyramidal neuron of the barrel cortex. The errorbars indicate one standard error of the mean and the blue line is a fit of Eq. 1.30. Data is a courtesy of Laurent Badel and Sandrine Lefort [Badel et al., 2008b].

1.5.1 The Adaptive Exponential Integrate-and-Fire Model

One obvious way to combine all the features is to add to the EIF model several linearized currents with cumulative spike-triggered adaptation:

$$C \frac{dV}{dt} = -g_L(V - E_0) + g_L \Delta_T \exp\left(\frac{V - V_T}{\Delta_T}\right) + I(t) - \sum_{i=1}^N w_i \quad (1.31)$$

$$\tau_i \frac{dw_i}{dt} = a_i(V_i - E_0) - w_i \quad (1.32)$$

$$\text{if } V(t) > V_T \text{ then } V(t) \rightarrow V_T \quad (1.33)$$

$$\text{and } w_i(t) \rightarrow w_i(t) + b_i \quad (1.34)$$

where each additional current w_i can be tuned by adapting its subthreshold coupling constant a_i and its spike-triggered jump size b_i . The simplest version of this framework assumes $N = 1$ and it is known as the Adaptive Exponential Integrate-and-Fire (AdEx; [Brette and Gerstner, 2005, Gerstner and Brette, 2009]). This model compares very well with many types of real




Figure 1.9: Generalizations of the LIF include either refractoriness, adaptation, linearized currents, or smooth spike initiation. Various regroupments have various names. For instance, the refractory-Exponential-Integrate-and-Fire (rEIF) regroups refractoriness, smooth spike initiation and the features of a LIF [Badel et al., 2008b].

neurons, as we will see in Sect. 1.7.

1.5.2 Integrated Models

For some neurons the spike initiation is sharp enough and can be neglected. In fact, if the slope factor $\Delta \rightarrow 0$ in Eq. 1.31, then the AdEx turns into a linear model with a sharp threshold. As we have seen in Sect. 1.2, the solution to the linear dynamical system can be cast in the form:

$$V(t) = E_0 + \int_0^\infty \kappa(s)I(t-s)ds + \sum_i \eta_a(t - \hat{t}_i) \quad (1.35)$$

where $\kappa(t)$ is the input filter and $\eta_a(t)$ is the shape of the spike with its cumulative tail. The sum runs on all the spike times $\{\hat{t}\}$ defined as the times where the voltage crosses the threshold ($V > V_T$). To be consistent with the processes seen in the previous sections the functions $\kappa(t)$ and $\eta_a(t)$ must be sums of exponentials. However, for the sake of fitting a model to experimental data, any type of basis function can be used. For arbitrary shape of the kernels $\kappa(t)$ and $\eta_a(t)$, this model is known as the cumulative Spike Response Model (SRM, [Gerstner et al., 1996b] and Gerstner [2008]) or more recently as the Generalized Integrate-and-Fire (GIF, Paninski [2004]).

The simplified Spike Response Model (SRM₀; Gerstner [2008]) is another related model worth pointing out. In the SRM₀ the sum in Eq. 1.35 extends to the last spike only. This makes a purely refractory model without spike-frequency adaptation.

Figure 1.9 summarizes the nomenclature for the combinations of generalizations. As we will see in the next section, the formalism of the SRM and SRM₀ models bridges the gap to a more general class of spiking models where the influence of noise is taken into account.




Figure 1.10: Membrane potential recorded from four repetitions of the same stimulus. Spikes are missed others are shifted and this variability is intrinsic to the neuron.

1.6 Noise

Variability is ubiquitous in the nervous system. In many ways this variability is seen as a feature rather than a defect [Faisal et al., 2008]. It is important therefore not to ignore all the noise, but to listen carefully in order to understand what it is trying to communicate.

In view of the intrinsic probabilistic nature of the neurons, it is difficult to predict the exact spike times because a neuron will fire with some jitter around an average spike time (Fig. 1.10). Rather, the models of neuronal behaviour must predict the probability to emit a spike in a given time interval. The probability $p(t)$ of observing a spike in a given small interval of δt defines the instantaneous firing rate:

$$p(t) = r(t)\delta t. \quad (1.36)$$

Mathematically speaking, $r(t)$ is a stochastic intensity, a terminology borrowed from the theory of point-processes [Daley and Vere-Jones, 1988]. In the context of neuron models $r(t)$ is called the firing intensity or instantaneous firing rate. It is also related to the experimentalist concept of a Peri-Stimulus Time Histogram (PSTH).

Chapter 12 describes the implication of noise at the biophysical level. Here we describe how noise influences the level of description of the LIF model. Furthermore, we describe a simple framework through which models such as the SRM are related to models of the firing intensity $r(t)$.

1.6.1 Synaptic Noise

Synaptic noise deals with the input $I(t)$. In this chapter $I(t)$ refers to the current arriving in the region which is solely responsible for spike initiation. This signal can be seen as being noisy: because some synaptic events happened without a pre-synaptic spike, because the unreliability of axonal propagation prevented a pre-synaptic spike from getting to the synapse, or because the variability in vesicle release and receptor channel opening make the amplitude of the postsynaptic event variable. From yet another point of view, one may be interested in the current input coming from an identified subpopulation of pre-synaptic neurons. In this view the rest of the pre-synaptic neurons form a considerable background synaptic noise. In any case, the synaptic noise is added to a deterministic current:

$$I(t) \rightarrow I(t) + \xi(t, V) \quad (1.37)$$

where $\xi(t)$ is the synaptic noise. The dependance of the noise on V is added because synapses make changes in conductance which must be multiplied by V to yield the current. Despite that, the dependance on V can often be neglected and the synaptic noise is considered as a time-dependent current $\xi(t)$.

There are multiple stochastic models of synaptic noise. Maybe the simplest situation takes into account a random ocurrence of synaptic events, $\hat{t}^{(pre)}$ each bringing in an exponentially decaying pulse of current:

$$\xi(t) = c \sum_{\hat{t}_i \in \{\hat{t}^{(pre)}\}} \exp\left(-\frac{t - \hat{t}_i}{\tau_s}\right) \theta(t - \hat{t}_i) \quad (1.38)$$

where τ_s is the time constant and c is the amplitude of the post-synaptic current decay. If the synaptic events occur randomly with frequency ν_e , the mean, the variance and the autocorrelation of the noise are [Stein, 1965]:

$$\mu \equiv \langle \xi \rangle = c \tau_s \nu_e \quad (1.39)$$

$$\sigma^2 \equiv \text{Var}[\xi] = \frac{c^2 \tau_s \nu_e}{2} \quad (1.40)$$

$$\langle [\xi(t) - \mu][\xi(t+s) - \mu] \rangle = \sigma^2 e^{-s/\tau_s}. \quad (1.41)$$

In the limit of small synaptic jumps ($c \rightarrow 0$) and large frequency of synaptic events ($\nu_e \rightarrow \infty$), the synaptic noise can be seen as a gaussian noise with exponential auto-correlation function. The dynamics of such a noise current is often written with the equation:

$$\xi(t + dt) = \xi(t) + \frac{(\mu - \xi(t))}{\tau_s} dt + \sigma G_t \sqrt{\frac{2dt}{\tau_s}} \quad (1.42)$$

where G_t is a random number taken from a standard normal distribution and dt is the step

size. This is called an Ornstein-Uhlenbeck process, a similar equation rules the position of a particle undergoing Brownian motion in an attracting potential. Eq. 1.42 is the diffusion approximation for synaptic inputs.

1.6.2 Electrical Noise

Electrical noise groups the thermal noise and the channel noise. The thermal noise (also called Nyquist or Johnson noise) adds fluctuation to the current passing through the membrane due to the thermal agitation of the ions. In this case the variance of the voltage fluctuations at rest is [Manwani and Koch, 1999]: $k_B T B / g_L$, where k_B is the Boltzmann constant, T is the temperature and B is the bandwidth of the system. Thermal noise is not the main source of electrical noise as it is three orders of magnitude smaller than the channel noise [Faisal et al., 2008].

Channel noise is due to the stochastic opening and closing of the ion channels (which itself is of thermal origin). This noise creates fluctuations in current that depend on the membrane potential following the activation profile of ion channels. Noise due to the Na-channels responsible for spike initiation can explain how the probability of firing depends on the amplitude of the stimulation when the stimulation consists of a short pulse of current [White et al., 2000]. Noise due to Na ion channels is therefore seen as an important source of noise which adds variability to the threshold for spike initiation.

Next section will explore the idea of a stochastic threshold further. More details on stochastic models of ion channels can be found in Chapter 12.

1.6.3 Generalized Linear Models

Consider the noiseless dynamics for $V(t)$ (as given by the SRM₀ or SRM) and replace the fixed voltage threshold with a stochastic voltage threshold:

$$V_T \rightarrow \theta + \zeta(t) \tag{1.43}$$

where θ is the average – or deterministic – threshold, and $\zeta(t)$ is a zero-mean white noise. This type of noise is called ‘*escape noise*’ [Gerstner and Kistler, 2002] and relates to the escape rate in models of chemical reactions [van Kampen, 1992]. In this scenario, the probability of finding the membrane potential above the threshold depends solely on the instantaneous difference between the voltage and the average threshold. We can write in general terms that the firing intensity is a nonlinear function of the modelled voltage trace:

$$r(t) = f(V(t) - \theta). \tag{1.44}$$

The monotonically increasing nonlinear function f is the cumulative distribution function of $\zeta(t)$.

Models such as the SRM₀ or the SRM have an explicit formulation for $V(t)$ that we can substitute in Eq. 1.44. These formulations for $V(t)$ require the knowledge of the spiking history. In this case, the firing intensity is dependent on the knowledge of the previous spike times. We will label this intensity differently, $\lambda(t|\{\hat{t}\}_t)$, since it does not equal the PSTH observed experimentally anymore. Writing the convolution operation of Eq. 1.35 with an asterisk and substituting the voltage in Eq. 1.44 by an explicit formula, we have:

$$\lambda = f \left(\kappa * I + \sum_i \eta_a(t - \hat{t}_i) + E_0 - \theta \right). \quad (1.45)$$

When all kernel κ and η_a are expressed as a linear combination of basis functions, the firing intensity would be a linear model if f were linear. With the nonlinear link-function f , this is instead a Generalized Linear Models (GLM). GLMs have convenient properties in terms of finding the parameters of the model and estimating the validity of the estimates [McCullagh and Nelder, 1989]. One of these properties is that the likelihood of observing a given set of spike times is a convex function of the parameters (when the link function is strictly convex [Paninski, 2004]). This means that it is always possible to find the best set of parameters to explain the data. It is a remarkable fact that only the knowledge of the spike times observed in response to a given stimulus is sufficient to estimate the filter κ and the shape of the adaptation function η .

The GLM model above depends on all the previous spikes, and therefore shows spike-frequency adaptation through the kernel $\eta_a(t)$. By excluding the possible influence of cumulative adaptation, it is possible to make a purely refractory stochastic model by dropping the dependence on all the previous spikes but the last one. This framework allows the theorems of renewal theory [Cox, 1962] to be applied, and to study the behaviour of networks of neurons analytically [Gerstner and Kistler, 2002].

In the same vein, it is possible to ignore completely the refractoriness and consider only the filtering of the input:

$$v = f(\kappa * I). \quad (1.46)$$

Despite the fact that it appears as too crude an assumption, we can gain considerable knowledge on the functional relationship between external stimuli and neuronal response. This model referred to as the Linear-Nonlinear Poisson model (LNP) was extensively used to describe the response of single neuron in the retina, thalamus or cortex as a function of the visual stimulus. When multiple neurons pave the way between the stimulus and the spikes generated by the LNP model, the filter function no longer represents the membrane filter of the cell but rather the linear filter corresponding to successive stages of processing before the signal reaches the neurons.

The mere fact that it is possible to make a decent firing-rate prediction with such a simple model makes a strong claim about the role of the neurons and the neural code. A claim that could be challenged by experiments pointing towards missing features. In the next section we further elaborate on the models described in the previous sections. We want to make sure we find the *simplest* description possible for a given set of experiments, but not *simpler*.

1.7 Advantages of Simple Models

How good are the simple models discussed in this chapter? Before addressing this question one needs to define what is meant by a good model. Suppose we focus on an experiment which injects time-varying current in a neuron that otherwise receives no input. What can be reproduced by the model? the average firing rate? the PSTH? the timing of each spike? The subthreshold potential? A model that is able to reproduce the averaged firing-rate may not be sufficient to underly the fast computational capabilities of neuron networks [Thorpe et al., 1996]. On the other hand, modelling the exact time-course of the membrane potential may not be necessary given that a later neuron only receives the spikes, and no information about the membrane potential of the neuron is transmitted. Perhaps the most appropriate task of a model is to predict the occurrence of spikes of a neuron receiving *in-vivo*-like input. Before evaluating the performance at predicting the spike times, we assess the ability of simple models to reproduce qualitatively the firing patterns observed in various neuron types.

1.7.1 Variety of Firing Patterns

Neurons throughout the nervous system display various types of excitability. The diversity is best illustrated in experiments injecting a predefined stimulus in cells that otherwise receive no input. For example, if an identical step current is used to drive multiple cells any differences in the observed firing patterns between the cells must be attributed to intrinsic mechanisms rather than the stimulation pattern. It is common to classify neurons according to their initial response to the step current as well as according to different steady state responses [Markram et al., 2004]. Consequently, the onset of firing is categorized as being either delayed, bursting or tonic. On the other hand the steady state can be tonic, adapting, bursting or irregular. Simple threshold models can reproduce all the firing patterns observed experimentally (Figure 1.11). The study of excitability types in such simple models sheds light on the basic principles contributing to the neuronal diversity.

Delayed spiking. Delayed spiking can be due to smooth spike initiation as in the EIF or to linearized subthreshold currents. Indeed, the EIF can produce delayed spiking when the stimulating current is slightly greater than $g_L(V_T - E_0 - \Delta)$. Another possibility is that a subthreshold current slowly activates at higher voltage to depolarize the membrane further. For instance Eq.'s 1.23-1.24 may lead to delayed spiking onset when $a < 0$ and $\tau_w > \tau$.




Figure 1.11: Multiple firing patterns are reproduced by merely tuning the parameters of a simple threshold model. Here the AdEx fires with tonic bursts, initial burst, spike-frequency adaptation, or delay. For each set of parameters, the model is simulated on a step current with low (close to current threshold for firing) or high (well above the firing threshold).

Bursting. Bursting can arise from many different causes and it is possible to define multiple types of bursting [Izhikevich, 2007]. Perhaps the simplest bursting model is the LIF with adaptation (Eq.'s 1.12-1.13). The high firing-frequency during a burst increases the adaptation current to a point where the neuron can no longer spike until its level of adaptation has decreased substantially. In this case the inter-burst interval is related to the time constant of adaptation. A hallmark of this type of bursting is the slowing down of the firing during the burst.

Transient spiking. Upon the onset of the stimulus, the neuron fires one or multiple spikes and then remains quiescent, even if the stimulus is maintained for a very long time. Spike-triggered adaptation or a moving threshold cannot account for this pattern. Transient spiking is due to a resonance with a subthreshold current (Fig. 1.5).

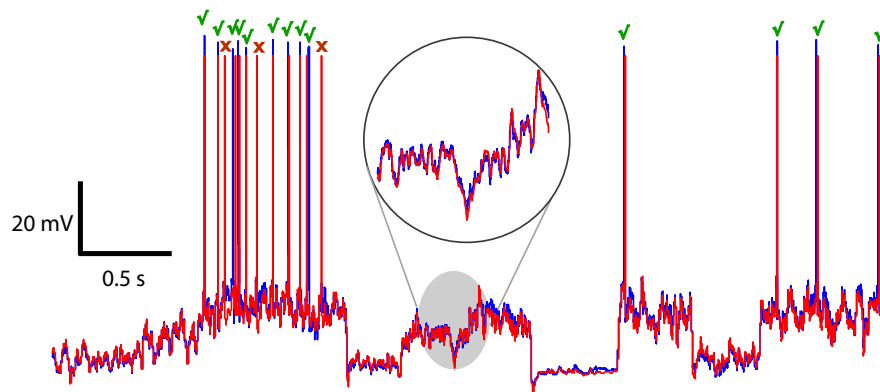


Figure 1.12: Overlaid traces of an AdEx model (red) and a fast-spiking interneuron (blue). Most of the spikes are predicted correctly by the model (green check marks) but some extra spike are predicted (red crosses). The subthreshold voltage traces match very well (inset). The data is a courtesy of Skander Mensi and Michael Avermann [Mensi et al., Under Review].

Adaptation. We have seen that spike-frequency adaptation is brought by either hyperpolarizing currents or a moving threshold which cumulates over multiple spikes.

This brief description emphasizes only the most important firing patterns, we will discuss the analysis of the bursting and initial bursting firing patterns further in Section 1.7.3. The main observation is that a two-dimensional model such as the AdEx is capable of describing a large variety of different firing patterns. The question now is whether these models can describe not only qualitative features of firing but also quantitative ones. Concretely, can we predict the timing of spikes in experiments using simple threshold models?

1.7.2 Spike-Time Prediction

Models of the threshold type can predict the spike times of real neurons *in vitro*. It is important to focus on the prediction performance and not simply on reproducing those spike times that were used to calibrate the model parameters. Otherwise, a very complex model could reproduce perfectly the data used for fitting while it would fail completely to reproduce the response to novel stimulus (a phenomenon called overfitting). The correct procedure is therefore to separate the data into a training set and a test set. The first is used for finding the best model parameters and the second to test the performance.

In vitro it is possible to simulate realistic conditions by injecting a fluctuating current in the soma of the neuron. For instance the injected $I(t)$ can be taken to be as in Eq. 1.38 or Eq. 1.42 as it would be expected from a high number of synaptic events affecting the soma. This current when injected in the soma drives the neuron to fire spikes. ‘Injecting’ this current in the mathematical neuron models will give a similar voltage trace. After determining the model parameters that yield the best performance on the training set, the neuron model can be used

to predict the voltage trace and the spike times on the test set. As we can see from Fig. 1.12, the AdEx is capable of predicting the subthreshold voltage and spike timings to an impressive degree of accuracy.

Deterministic models such as the SRM₀ with a dynamic threshold [Jolivet et al., 2006], the AdEx [Jolivet et al., 2008a], the rEIF [Badel et al., 2008b], the GIF [Mensi et al., Under Review]² or other similar models [Kobayashi et al., 2009] have been fitted to such *in vitro* experiments and their predictive performance evaluated. To evaluate the performance of deterministic models, the number of predicted spikes that fall within ± 4 ms of the observed spikes are counted. When discounting for the number of spikes that can be predicted by chance [Kistler et al., 1997], a coincidence factor is obtained ranging from zero (chance level) to one (perfect prediction). This coincidence factor ranged from 0.3 to 0.6 for pyramidal neurons of the layer 5, and from 0.5 to 0.8 for fast-spiking interneurons of the layer 5 of the cortex.

It turns out that these performances are very close to optimal, as we can see if we consider that a real neuron would not reproduce exactly the same spike times after multiple repetitions of the same stimulus (as shown in Fig. 1.10). The best coincidence factor achievable by the models is the intrinsic reliability R , which is the average of the coincidence factor across all pairs of experimental spike trains generated with the same stimulus. This value can be seen as an upper bound on the coincidence factor achievable by the mathematical models. Scaling the model-to-neuron coincidence factor by the intrinsic reliability and multiplying by 100 gives a measure of the percentage of the predictable spikes that were predicted. For models like the AdEx or the SRM, this number ranged from 60 to 82 % for pyramidal neurons, and from 60 to 100 % for fast-spiking interneurons. Simpler models do not share this predictive power: the LIF only accounts for 46 to 48 % of the predictable portion of spikes.

Models from the GLM family have been used to predict the spike times of neurons in the retina, thalamus or cortex from the knowledge of the light stimulus. Almost perfect prediction can be obtained in the retina and in the thalamus [Pillow et al., 2005]. Furthermore, it has been shown that refractoriness or adaptation is required for good prediction [Berry and Meister, 1998, Pillow et al., 2005, Truccolo et al., 2010]. Furthermore, the quality of the prediction can be improved by taking into account the coupling between adjacent cells [Pillow et al., 2008].

1.7.3 Ease of Mathematical Analysis

The greater simplicity of the neuron models discussed in this chapter compared to biophysical models of the Hodgkin-Huxley type have another advantage: the ease of mathematical analysis. This is particularly advantageous for studying neuron networks and for investigating synaptic plasticity. This is a vast field of study where many macroscopic properties of neuron networks can be scrutinized: learning, oscillations, synchrony, travelling waves, coding, and possibly others (see Hoppensteadt and Izhikevich [1997], Dayan and Abbott [2001], and Gerstner and Kistler [2002] for introductions). Paving the way to these exciting fields, mathematical analysis

²Chapter 5




Figure 1.13: Initially bursting (A) and regular bursting (B) trajectories in both phase plane and as a time-series. The initial state of the neuron model before the step current is marked with a cross and is associated with the V -nullcline shown as a dash line. After the step increase in current the V -nullcline shifts upwards (full line). The w -nullcline is shown in green and the trajectories are shown in blue with each subsequent reset marked with a square. The first and last resets are labeled with their ordinal number. Regular bursting is distinguished from initial bursting by the presence of two voltage resets between the trajectory of the fifth spike and the V -nullcline.

has yielded important insights that relate the function of single neurons with that of networks.

The characteristic of the response to various types of stimulations can often be described in mathematical terms. For a constant current I , the firing frequency of the LIF is given by:

$$\nu = \left[\tau \ln \left(\frac{g_L(V_r - E_0) - I}{g_L(V_T - E_0) - I} \right) \right]^{-1}. \quad (1.47)$$

The mean frequency when the current is an Ornstein-Uhlenbeck process (Eq. 1.42) can be written as an integral formula [Johannesma, 1968]:

$$\nu = \left[\tau \sqrt{\pi} \int_{\sqrt{\tau}(g_L(V_r - E_0) - \mu)/(\sigma\sqrt{2\tau_s})}^{\sqrt{\tau}(g_L(V_T - E_0) - \mu)/(\sigma\sqrt{2\tau_s})} e^{x^2} [1 + \operatorname{erf}(x)] dx \right]^{-1} \quad (1.48)$$

where $\operatorname{erf}(x)$ is the error function. When the non-linearity for spike initiation is taken into account, it is not possible to arrive at exact solutions anymore. Yet, for some parameter values there are approximations that can be worked in. This way, approximated solutions can be written for the EIF or the AdEx receiving gaussian white noise [Fourcaud-Trocme et al., 2003, Richardson, 2009]. When there is a strong effect of adaptation it is not possible to arrive at closed-form solutions for ν and one must rely on numerical integration of the appropriate Fokker-Planck equations can be used [Muller et al., 2007, Toyozumi et al., 2009].

Mathematical analysis is equally successful to relate the observed types of firing patterns with the parameter values. We will illustrate the application of bifurcation theory with an example related to Fig. 1.11: distinguishing initial bursting from a regular bursting using phase plane analysis.

Repetitive bursting is created in the AdEx model by constantly alternating between slow and

fast interspike intervals. When the injection is a step increase in current (as in Fig. 1.11) this corresponds to a sudden shift of the V -nullcline (*i.e.* the set of points where $\frac{dV}{dt} = 0$) in the phase plane. Before the step, the state of the neuron is at the stable fixed point which is sitting at the intersection between the V - and the w -nullcline. After the step, the stable fixed point has disappeared and this results in repetitive firing. The distinction between an initially bursting AdEx model and a regular bursting one is made by considering the location of the resets in the phase plane. Fig. 1.13 shows that both types of bursting have spike resets above and below the V -nullcline. Resetting above the V -nullcline brings a much larger interspike interval than resetting below. To achieve this, a spike reset above the V -nullcline must be able to make at least one spike reset below the V -nullcline before being mapped above again. This is seen in Fig. 1.13B where there is sufficient space below the V -nullcline and above the trajectory of the fifth spike for at least one reset. This leads the repetitive bursting of two spikes. On the other hand, in Fig. 1.13A the fourth spike – which was reset above the V -nullcline – is followed by another reset above the V -nullcline. This leads to the end of the initial burst which is followed only by long interspike intervals. By observing at the phase planes in Fig. 1.13, one can conclude that a higher voltage reset helps to bring about repetitive bursting.

Similar conclusions can be drawn for all the firing patterns [Naud et al., 2008]³. Furthermore, it is sometimes possible to have analytical expressions relating the parameter values and the firing patterns [Touboul and Brette, 2008]. Such conclusions are possible mainly because the complexity of the model was kept low. Indeed, the structure of firing patterns have been studied in other simple models [Izhikevich, 2007, Mihalas and Niebur, 2009].

1.8 Limits

The neuron models presented in this chapter present an idealized picture. The first idealization is that all these models consider the neuron as a point with no spatial structure. This point could represent the axon initial segment where the spike is initiated. However, real neurons receive input distributed not only in the soma but also in their dendritic arborizations. Dendrites bring multiple types of non-linear interactions between the inputs. Dendritic ion channels combine with basic properties of AMPA, GABA and NMDA synapses to make dendritic output to the soma dependent on the spatio-temporal pattern of excitation. Spatially extensive models (see Chapter 11) may be required to correctly translate input spikes into current arriving at the axon initial segment, but to what extent? This remains to be addressed experimentally.

Another central assumption prevalent throughout this chapter is that the spike does not change its length or shape with different stimuli. While this is seen as a good approximation for cortical neurons firing at low frequency [Bean, 2007], some neurons have action potential shapes that vary strongly as a function of either stimulus, firing frequency or neuromodulation. For instance, the interneurons taking part in the motor pattern generation of flight in locusts

³Chapter 2

Chapter 1. Introduction: The Performance (and limits) of Simple Neuron Models

display spikes that reduce by half their amplitude and width when part of a burst of activity [Wolf and Pearson, 1989]. Similarly, half-blown spikes are observed frequently in the lobster stomatogastric pattern generator [Clemens et al., 1998] or in complex spikes of cartwheel cells in the dorsal cochlear nucleus [Golding and Oertel, 1997].

The simple models described in this chapter can be seen as a high-level description of the complex biophysical mechanisms mediated by ion channels, ion pumps, and various chemical reactions involving neurotransmitters. At this level of description the molecular cascades for the effect of – for example – acetylcholine are not modelled. Though the effect of neuromodulators can be calibrated and cast in stereotypical modification of the simple model [Slee et al., 2005], it is not an intrinsic feature and the calibration has to be performed for each scenario. Complex biochemical as well as biophysical models are the norm if pharmacological procedures are studied (see Chapter 17-18).

To conclude, simple neuron models capture essential features that are required for transducing input into spikes. If we bear in mind the inherent limitations, the study of simple neuron models will continue bringing new insights about the role of single neurons. In particular, the roles of adaptation and variability are only starting to be considered.

2 Firing patterns in the adaptive exponential integrate-and-fire model

For simulations of large spiking neuron networks, an accurate, simple and versatile single-neuron modeling framework is required. ¹ Here we explore the versatility of a simple two-equation model: the adaptive exponential integrate-and-fire neuron. We show that this model generates multiple firing patterns depending on the choice of parameter values, and present a phase diagram describing the transition from one firing type to another. We give an analytical criterion to distinguish between continuous adaption, initial bursting, regular bursting and two types of tonic spiking. Also, we report that the deterministic model is capable of producing irregular spiking when stimulated with constant current, indicating low-dimensional chaos. Lastly, the simple model is fitted to real experiments of cortical neurons under step current stimulation. The results provide support for the suitability of simple models such as the adaptive exponential integrate-and-fire neuron for large network simulations.

2.1 Introduction

Large-scale simulations of cortical activity and theoretical investigations of neuronal dynamics require neuron models that are mathematically tractable, biologically relevant and computationally fast. Moreover, a modeling framework should be sufficiently versatile to span the whole diversity of neuron types by tuning a restricted number of parameters, avoiding the need of a new model for each class of neuron. Modeling the complete gating dynamics of ion channel densities in neuronal membranes satisfies only two of these five requirements: biological relevance and versatility. On the other hand, modeling a neuron either as a coincidence detector, as a resettable integrator or as a stochastic point process may fail to catch important aspects of single-neuron behavior. Small modifications to these simple models can bring them closer to reality. Lately multiple advances were made in that direction.

In the presence of high synaptic bombardment, modeling accurately the spike initiation is crucial and the Leaky Integrate-and-Fire (LIF) must be augmented by an exponential term to

¹Text copied from Naud, R., Marcille, N., Clopath, C., and Gerstner, W., *Biological Cybernetics*, 2008 (full citation in the References).

faithfully process fast inputs signals (Fourcaud-Trocme et al. [2003]). An additional recovery variable is important to capture adaptation and resonance properties (Richardson et al. [2003], Izhikevich and Desai [2003]). A simple quadratic model of spike initiation with a linearly dependent recovery variable and a reset in the state variables is sufficient to account for most types of firing patterns observed in the central nervous system (Izhikevich [2007]), but spike firing in that model occurs with an unrealistic delay. Unlike a quadratic dependence on voltage (Latham et al. [2000]), an exponential nonlinearity (Fourcaud-Trocme et al. [2003]) keeps the subthreshold dynamics linear and matches direct measurements in cortical neurons (Badel et al. [2007]). The exponential model combined with an adaptation variable - called the adaptive exponential integrate-and-fire (AdEx)² - is simple because it is described by only two equations and a reset condition. It is by construction more realistic than the LIF, and it was shown to predict with high accuracy the spike timing of a conductance-based Hodgkin and Huxley model (Brette and Gerstner [2005]) and the spike timing of real pyramidal neurons under noisy current injection (Jolivet et al. [2006]).

In this paper we explore the versatility and the biological relevance of the AdEx model. We show that the AdEx reproduces multiple firing patterns and study the correspondence between the parameters and the firing types. Finally, we fit the model to experimental traces and obtain sets of parameters describing cortical fast spiking interneurons and regular spiking pyramidal neurons.

2.2 Adaptive Exponential Integrate-and-Fire

The adaptive Exponential Integrate-and-Fire model (AdEx) describes the evolution of the membrane potential $V(t)$ when a current $I(t)$ is injected. It consists of a system of two differential equations:

$$C \frac{dV}{dt} = -g_L(V - E_L) + g_L \Delta_T \exp\left(\frac{V - V_T}{\Delta_T}\right) + I - w, \quad (2.1)$$

$$\tau_w \frac{dw}{dt} = a(V - E_L) - w. \quad (2.2)$$

When the current drives the potential beyond V_T the exponential term actuates a positive feedback which leads to the upswing of the AP. The exponential is related to the quasi-instantaneous reaction of the activation variable of the sodium channel in a Hodgkin-Huxley-type neuron model. The upswing is stopped at a reset threshold which we fix at 0 mV. The

²Also referred to as aEIF.

downswing of the action potential, is replaced by the reset condition:

$$\text{if } V > 0 \text{ mV then } \begin{cases} V \rightarrow V_r \\ w \rightarrow w_r = w + b. \end{cases} \quad (2.3)$$

We emphasize that voltage reset is *to* a fixed value V_r whereas w -reset is *by* a fixed amount b . Therefore the adaptation variable w can accumulate during a spike train whereas voltage does not.

There are nine parameters required to define the evolution of the membrane potential (V) and the adaptation current (w). The nine parameters can be separated into scaling parameters and bifurcation parameters. The scaling parameters are the parameters responsible for scaling the time axis, for the stretch and for the offset of state variables. The five scaling parameters are: total capacitance (C), total leak conductance (g_L), effective rest potential (E_L), threshold slope factor (Δ_T), effective threshold potential (V_T). Absorbing the parameters C and g_L into the time scale $\tau_m = C/g_L$, and using Δ_T and V_T to set the scale and offset of the membrane potential, then after appropriate rescaling of I and w , Eq.'s 2.1-2.2 can be reduced to a system of equations with dimensionless variables and only four parameters (Touboul and Brette [2008]). The resulting four parameters are bifurcation parameters and are directly proportional to the conductance a , the time constant τ_w the spike triggered adaptation b , and the reset potential V_r . Modifying these parameters brings qualitative changes in the behavior of the system, such that different firing patterns become possible. Note that a controls the sensitivity of the adaptation current to voltage even in the absence of spikes. This voltage coupling can arise from linearized ion channels (Sabah and Leibovic [1969], Mauro et al. [1970], Koch [1999], Richardson et al. [2003]) or from the interaction with a passive dendritic compartment which is not modeled explicitly. Voltage coupling via the parameter a with $a > 0$ acts like a negative feedback and leads to adaptation. This parameter can arise from the linearization of "resonant currents" (Hutcheon and Yarom [2000]) and can give rise to damped subthreshold oscillations. In the AdEx we assume that the voltage coupling (characterized by a) and the spike-triggered adaptation (characterized by b) have a similar time constant and can thus be lumped together in a single equation (Eq. 2.1). The biophysical mechanisms for spike-triggered adaptation are ion channels such as I_M , $I_{K(slow)}$, or $I_{K(Ca)}$ (Benda and Herz [2003]). Note that the spike triggered adaptation has contributions from both the a -term and the b -term. Having two parameters (a and b) allows to decouple the spike-triggered adaptation from the voltage coupling.

The role of the bifurcation parameters is best understood through phase plane analysis (see Izhikevich (2007)). Briefly, phase plane analysis involves plotting the state variables relative to each other. Nullclines represent the area in phase space in which a given variable remains constant. The V -nullcline (or w -nullcline) is defined as the set of points with $\frac{dV}{dt} = 0$ (or $\frac{dw}{dt} = 0$ respectively). The shape and position of the nullclines depends on the parameters of the model. For instance the minimum of the V -nullcline is given by the parameter V_T . The

Chapter 2. Firing patterns in the adaptive exponential integrate-and-fire model

slope of the left branch of the V -nullcline is proportional to the leak conductance g_L . Changing the current in Eq. 2.1 involves a vertical shift in the V -nullcline without changing its shape (Fig.'s 2.1-2.3). The intersection of the two nullclines defines fixed points which can be stable or unstable. In particular, these fixed points can change position, merge and disappear or simply loose or gain stability upon changes in one or several parameters. At a bifurcation point, a change in the stability of the pattern of fixed points occurs, and this qualitatively modifies the behavior of the system. In the system of Eq.'s 2.1-2.2, the choice of a and τ_w determines whether an increase in current induces a loss of stability via an Andronov-Hopf or via a saddle-node bifurcation (Fig. 2.1). When the current is increased such that the stable fixed point of Eq.'s 2.1-2.2 loses stability, repetitive spiking ensues and the current at which this occurs is the rheobase. A mathematical analysis (Touboul and Brette [2008]) has shown that




Figure 2.1: Phase plane representation of a step current injected in an AdEx model where (a) a saddle-node bifurcation is responsible for the loss of stability, and (b) the Andronov-Hopf bifurcation is responsible for the loss of stability. In the phase planes the trajectories of the first and second spikes are represented by blue squares and the state of rest is indicated by the blue cross. As the current increases, the V -nullcline shifts upwards. This makes the two fixed points move toward each other. (a) In the saddle-node bifurcation, the fixed points disappear after the stable fixed point merges with the unstable fixed point. The point where the two fixed points merge lies close to (but slightly to the right of) the voltage V_T , i.e. the minimum of the V -nullcline. (b) As the stable fixed point moves towards the right, the slope of the V -nullcline increases at the fixed point. If the slope of the w -nullcline is sufficiently high, this can lead to a loss of stability of the stable fixed point before the fixed points disappear. The w -nullcline is shown in green, the V -nullcline in the absence of current is the curved dash line, the V -nullcline in the presence of stimulating current is the curved solid line (black). Unstable fixed points are encircled. The scale bars corresponds to 20 mV vertically and 20 ms horizontally.

under the condition $a/g_L > \tau_m/\tau_w$ the transition occurs via the Andronov-Hopf bifurcation at

the value:

$$I_{AH} = (g_L + a) \left[V_T - E_L - \Delta_T + \Delta_T \ln \left(1 + \frac{\tau_m}{\tau_w} \right) \right] \quad (2.4)$$

$$+ \Delta_T g_L \left(\frac{a}{g_L} - \frac{\tau_m}{\tau_w} \right), \quad (2.5)$$

and otherwise ($a/g_L < \tau_m/\tau_w$) via the saddle-node bifurcation at the value

$$I_{SN} = (g_L + a) \left[V_T - E_L - \Delta_T + \Delta_T \ln \left(1 + \frac{a}{g_L} \right) \right]. \quad (2.6)$$

Together, I_{SN} or I_{AH} define the minimal value of a slow ramp current at which the AdEx starts spiking, *i.e.* the rheobase. In continuous two-dimensional neuron models such as the Morris-Lecar model, the type of bifurcation is directly related to the type of frequency-current (f-I) relation (Rinzel and Ermentrout [1998]). Model neurons with a saddle-node on invariant circle bifurcation have a continuous f-I curve (very low firing frequency for currents just above rheobase) and are classified as type I neurons, whereas neurons with an Andronov-Hopf-bifurcation or a saddle-node bifurcation off invariant circle have a discontinuous f-I curve with a jump to finite firing frequency just above rheobase and are classified as type II (Izhikevich [2007]). A one-dimensional quadratic (Latham et al. [2000]) or exponential (Fourcaud-Trocme et al. [2003]) integrate-and-fire model with a reset to the resting potential has a saddle-node bifurcation, and is always type I. Interestingly, such a simple relation between bifurcation class and type of f-I curve does not hold for the AdEx model, because of an intricate interplay between the continuous two-dimensional dynamics and the discontinuous reset. Different values of the reset parameters V_r and b can yield a continuous f-I curve as in a type I neuron model, or a discontinuous f-I curve as in a type II neuron model (Fig. 2.2c and 2.2d). Very low firing frequencies in a neuron with saddle-node bifurcation just above rheobase are possible if the limit cycle passes through the region where at rheobase current, the saddle and the node merged. In this region ('ghost' of the saddle-node bifurcation) the speed of the trajectory is very low leading to long interspike intervals. However, for some combinations of the reset values V_r and b , the trajectory in the AdEx model does *not* pass nearby the lost fixed point, but passes further away, leading to a finite period of the (Fig. 2.2c and 2.2d). This is the case where the system just *below* rheobase is bistable, *i.e.*, at the same constant subthreshold input, both periodic firing and constant membrane potential are possible unless the reset V_r is above V_T . We note that for a small adaptation parameter $a \rightarrow 0$, such a bistability is not possible. To summarize, the value of the reset parameters have a crucial role in determining the neuronal firing patterns - and this will be explored in the next section.

2.3 Multiple Firing Patterns

In order to study the range of firing patterns accessible with the AdEx, we simulate the injection of a step current. This is the most common experimental paradigm used by electrophysiolo-




Figure 2.2: Types of excitability depend on the point of reset. When $a = 0$ (a) only type I excitability is possible. The phase plane at a current just after rheobase shows that the trajectories are forced to pass through the ghost of the saddle-node bifurcation independently of the starting point (three steady-state trajectories are shown for three different values of the voltage reset). The frequency-current plot on the right shows the f - i curve corresponding to the three different reset conditions (consistent line types; full, dash or dot/dash). When $a > 0$ it is possible to have type I and type II depending on the reset (b). The region in pink shows the basin of attraction of the stable fixed point just before ($I_{SN} - 0.1$ pA) it is to lose stability via a saddle-node bifurcation. The resets at -50 and -60 mV result in trajectories that pass very near the V -nullcline, and are therefore very slow. The reset at -80 mV is outside the ghost of the attraction basin so that its trajectory passes further away from the ghost of the saddle-node bifurcation. The insets show the enlarged areas enclosed by the gray rectangles. The arrow on the I -axis of the f - I plot indicates the current that was used to draw the phase plane. The conventions of line colors is the same as for Fig. 2.1.




Figure 2.3: Phase diagram and associated traces illustrating two types of tonic spiking: tonic spiking with sharp reset (a) and tonic spiking with broad reset (b). The only modifications needed to change the neuron model (a; sharp resets) into the neuron model (b; broad resets) is an increase of the spike triggered adaptation b and increase of the voltage reset, V_r . In both cases, some degree of adaptation is seen, yet both traces do not belong to the continuously adapting class because there is no substantial adaptation beyond the first two inter-spike intervals. The convention for symbols, line colors and scale bars was the same as in Fig. 2.1.

gists to study and classify neurons (e.g. Markram et al. [2004]). Mathematically, this situation corresponds to the solution of Eq. 2.1-2.2 with constant current and initial values $V(0) = E_L$ and $w(0) = 0$. Similar to what is seen in real neurons, the response of the model to the step current is very diverse, and depends on the choice of parameters. In Fig. 2.4 we show an example of typical firing patterns that can be generated by varying the parameters of the AdEx. The parameters associated with each example are given in Table 2.1. In this section we will describe how the different firing patterns arise with our simple model.

Sharp vs Broad Spike After Potential (SAP) The AdEx can produce adapting and tonic traces of two qualitatively different types. In Fig. 2.3a we see an example of sharp SAPs where the potential increases monotonically after a rapid downswing of the action potential. This type of reset is commonly seen in fast spiking interneurons (Bean [2007]), and corresponds to a low value of the voltage reset V_r , combined with a weak spike-triggered adaptation b . In the phase plane we see that the reset is made at a point below the V -nullcline.

Broad SAPs, on the other hand, are observed in regular spiking pyramidal neurons (Connors and Gutnick [1990]) and in the continuously adapting interneurons. A broad SAP is recognized by its low curvature at all times after the spike (Fig. 2.3b).

The two types of SAPs correspond to two different spiking trajectories in the phase plane of the AdEx model. These are determined by the location of the reset point in the phase plane. If the reset point is above the V -nullcline (where $\frac{dV}{dt} < 0$ everywhere), the voltage will decrease before increasing in preparation of a spike, this trajectory is termed a broad SAP or broad reset. If the reset point is below the V -nullcline, the spiking trajectory starts to increase immediately after the reset ($\frac{dV}{dt} > 0$ everywhere below the V -nullcline), and the SAP appears as a sharp reset. We can write this distinction as an analytical relation that depends on the reset point

(V_r, w_r) . There is a broad reset if

$$w_r > -g_L(V_r - E_L) + g_L \Delta_T \exp\left(\frac{V_r - V_T}{\Delta_T}\right) + I, \quad (2.7)$$

and otherwise the reset is sharp.

Tonic vs Adapting The simplest type of spiking pattern is the regular discharge of action potentials (tonic firing, see Fig. 2.4a). This firing pattern is the only firing pattern that a standard leaky or non-leaky integrate-and-fire model can generate subject to constant current injection. In the framework of the AdEx, it corresponds to the absence of spike-triggered adaptation and adaptation sensitivity to subthreshold voltage ($a, b = 0$). Most neurons, however, show some level of spike-frequency adaptation. In this firing pattern, the inter-spike interval (ISI) grows during a sustained stimulus (Fig. 2.4b). The classification between adapting and non-adapting can be drawn from the adaptation index:

$$A = \frac{1}{N - k - 1} \sum_{i=k}^N \frac{ISI_i - ISI_{i-1}}{ISI_i + ISI_{i-1}} \quad (2.8)$$

where $k \geq 2$ is used to disregard any initial transient. Consistent with other studies (Druckmann et al. [2007]), we take $k = 4$; i.e. we disregard the first two inter-spike intervals. Typically fast spiking interneurons have an adaptation index, A , of 0.005 whereas regular spiking pyramidal neurons have $A = 0.015$ (with $N = 15$ to 40 spikes, Druckmann et al. [2007]). The adaptation index will depend on the number of spikes considered, in this article, we will compute the adaptation index with N fixed to 20 spikes, and take $A_c = 0.01$ as a critical value for classifying a spike train as adapting ($A \leq A_c$) or non-adapting ($A > A_c$).

Initial Bursting vs Adapting Initial bursting denotes a group of spikes that were emitted at a frequency considerably greater than the steady-state frequency. This definition is very ambiguous and in many experimental traces initial bursting is indistinguishable from pronounced adaptation. In the framework of the AdEx, a clear definition becomes apparent. Initial bursting arises when the spiking starts with one or several sharp reset followed by broad resets only (Fig. 2.4c). With this definition, there is a qualitative difference between the SAP forming the initial burst and the SAP forming the tonic spiking that follows, as observed in experiments (Markram et al. [2004]).

Regular Bursting Regular bursting appears in a scenario similar to initial bursting except that the trajectory starting after the first broad reset projects below at least one of the previous reset points in the phase plane, such that the next reset point is below the V -nullcline; i. e. the next reset is sharp (Fig. 2.4d). This situation leads to an alternation between sharp and broad resets. Regular bursting is made possible with a V_r higher than the effective threshold V_T (note that V_T corresponds to the minimum of the V -nullcline), so as to ‘shield’ some reset




Figure 2.4: Phase plane representation of 8 firing patterns. Firing patterns observed during a step current stimulation are: (a) tonic spiking, (b) adaptation, (c) initial burst, (d) regular bursting, (e) delayed accelerating, (f) delayed regular bursting, (g) transient spiking and (h) irregular spiking. The voltage traces are shown with a scale bar that corresponds to 100 ms and 20 mV. The graphs on the left show the traces in the phase planes as a trajectory (blue line) in the two state variables ($V(t)$, $w(t)$). The w -nullcline (green) is a straight line, the V -nullcline before current stimulation is the curved dashed black line, and in the presence of stimulation, the curved solid line (black). The stable fixed point in (g) is indicated with a black, filled circle, and all the other symbols refer to the same convention as in Fig.'s 2.1-2.2. Comparing (b) with (c) illustrates that reset points jumping above the V -nullcline lead to initial bursting. Comparing (c) with (d) clarifies that regular bursting is obtained when the first broad reset generate a trajectory that passes below at least one of the previous reset points.

points on the right hand side of the V -nullcline.

Delayed Spiking and Acceleration A negative a acts as a positive feedback which can be responsible for delayed initiation. Delayed spiking appears when injecting a current close to the rheobase. Because of $a < 0$, the (hyperpolarizing) adaptation current w is slowly decreasing at a depolarized V , allowing the neuron to eventually spike once the adaptation has decreased sufficiently (Fig. 2.4e). In the phase plane, the trajectory of the first spike stays below, but follows the contour of the V -nullcline. As it approaches the V -nullcline, the trajectory slows down since the magnitude of $|\frac{dV}{dt}|$ decreases. With $a < 0$, the system of equations of the AdEx can lead to spike frequency acceleration (Fig. 2.4e, similar to experiments by Beierlein et al. [2003]) if the spike-triggered adaptation b is weak, as shown in Fig. 2.4e. For greater values of spike-triggered adaptation, it is also possible to get a delayed adapting trace (not shown) and a delayed bursting trace (Fig. 2.4f). Note that in both Fig. 2.4e and 2.4f the trajectory starts with negative slope ($\frac{dw}{dt} < 0$ while $\frac{dV}{dt} > 0$) reflecting the decreasing adaptation current.

Rebound or Transient Spikes Post-inhibitory rebound is seen in several types of neurons. When a hyperpolarizing step current is released abruptly, some neurons will spike one or several time(s) before reaching their state of rest. This phenomenon is very similar to transient spiking observed during a step of depolarizing current. In both cases a stable fixed point remains, even after the sudden increase in current. A spike is nevertheless produced because the adaptation current is too slow to compensate for the sharp change in current. If the applied current were increased gradually to the same current, the neuron would not spike because it would have had time to adapt. In the phase plane, rebound corresponds to a situation where the initial condition is located outside the area of trajectories converging to the stable fixed point, but the reset point is located inside this basin of attraction (Fig. 2.4g).

Irregular Spiking Irregular spiking can occur in an AdEx model despite the fact that the equations are deterministic. Irregular spiking is manifest when the interspike interval keeps on changing without periodicity during the observation interval. There is an alternation of sharp and broad resets, but — unlike the case of initial bursting or regular bursting firing patterns — in the case of irregular spiking the sequence is not periodic, reminiscent of stuttering cells observed in some nest basket cells (Wang et al. [2002]). The irregular firing pattern appears for a restricted set of parameters, and the volume occupied by the irregular spiking pattern in the parameter space seems to be patchy. We verified that the behavior was not due to a numerical artifact: Simulating the model neuron with the same parameter set with different numerical integration methods (reducing the forward-Euler time step, using Runge-Kutta or Adams-Bashforth-Moulton) always produced an irregular spiking pattern. Even though the exact spike train was not identical, spiking was aperiodic so that the classification of the spike pattern (adaptive, bursting, irregular) did not change after a change in the numerical method. Though other regions of the parameter space may produce irregular spiking, the

pattern was found to be most noticeable at negative a , large spike-triggered adaptation (b) and high voltage reset (Fig. 2.4h).

For a given set of parameters, irregular spiking emerges for a restricted range of injection current Fig. 2.5a. We checked that for this set of parameter the AdEx model showed chaotic behavior (Strogatz [1994]) by testing the dependence of the numerical integration on the initial conditions. We perturbed the initial conditions by a very small value $\delta = 10^{-12}$ and evaluated the error, ξ_i , in the adaptation variable at each spike:

$$\xi_i = (w(\hat{t}_i) - w_\delta(\hat{t}_i))^2 \quad (2.9)$$

where $w(\hat{t}_i)$ and $w_\delta(\hat{t}_i)$ is the adaptation current at the time of the spike i for the unperturbed and perturbed initial condition, respectively. Fig. 2.5b shows that this error grows exponentially with the number of spikes. The slope in the semi-log plot is 2.56. Therefore, there is a great sensitivity of the solution on the initial conditions, this confirms that irregular spiking is due to chaos. A recurrence relation (map) can be found by plotting the magnitude of the interspike interval with respect to the preceding interspike interval. After a very large number of spikes ($n = 1240$), the interval map appears as a continuous function (Fig. 2.5c). The points in Fig. 2.5c all lie on a line of very small thickness, illustrating that the system is deterministic and neither artifacts nor noise is responsible for the observed irregularity in the interspike interval pattern.

2.4 Parameter Space

Given the definitions for each firing pattern, we can now investigate how these firing patterns depend on the specific set of parameters. Since this problem depends only on the bifurcation parameters, we fixed the scaling parameters to realistic values : $C = 100$ pF, $g_L = 10$ nS, $E_L = -70$ mV, $V_T = -50$ mV and $\Delta_T = 2$ mV. This corresponds to a membrane time constant $\tau = C/g_L$ of 10 ms. Resting potential (E_L), effective threshold (V_T) and slope factor (Δ_T) are in the range seen in experiments (Badel et al. [2007]). Changing the scaling parameters does not change the set of firing patterns accessible for a given set of bifurcation parameters, but may modify the amount of current necessary to go from one firing type to another. In addition, some firing patterns exist only close to the rheobase current (delayed spiking, transient spiking). In this section we will consider only the firing patterns that appear at a step current twice the amplitude of the rheobase current. The rheobase was determined according to analytical expressions given in Touboul and Brette [2008] and corresponds to the (quasi-stationary) ramp current at which the neuron begins to spike because the stable fixed point losses stability. The firing pattern is likely to change for different step current amplitudes, but we limit our study to just one amplitude of current step.

For each set of the bifurcation parameters, Eq.'s 2.1-2.3 were solved with the Adams-Bashforth-Moulton numerical integration method (ode113 in Matlab), we stopped the simulation after 50 spikes or 16 seconds in model time whatever occurred earlier. We characterized the firing




Figure 2.5: Irregular firing is chaos in the AdEx. Spikes times of an irregular spiking model are shown in (1) for three different amplitudes of the stimulating current step. At medium current amplitude ($I = 150$ pA) it spikes without periodicity, this current amplitude was used to make (b) and (c). The numerical integration of an irregular spiking model depends heavily on the initial conditions (b), such that ξ grows exponentially with the number of spikes simulated. The stars denote a modification of the initial condition in w only, the diamonds is a modification in V and the circles a modification in both w and V . (c) A linear fit shows a slope of 2.56 (full line). The interval map (c) between each interspike interval and the preceding one appears as a thin, continuous curve ($n = 1240$ spikes). The parameters for the irregular spiking model are given in Table 2.1.

pattern according to the following simple rules:

- tonic: strictly sharp resets or strictly broad resets and $-0.01 < A < 0.01$;
- adapting: strictly sharp resets or strictly broad resets and $A > 0.01$;
- accelerating: strictly sharp resets or strictly broad resets and $A < -0.01$;
- initial bursting: ordered sequence going from sharp to broad resets (e.g. sharp - sharp - broad - broad - broad);
- regular bursting: alternation between broad and sharp SAP such that the number of sharp resets between each broad reset is constant;
- irregular spiking: alternation between broad and sharp SAP such that the number of sharp resets between each broad reset is not constant after the third broad reset;

All programs were written in Matlab (The Mathworks, Natick, MA) and ran on a personal computer.

In Fig. 2.6, the distribution of firing patterns is shown as a function of the reset parameters V_r and b with fixed values of a and τ_w . The fixed τ_w was 5 ms or 100 ms such that the w -variable can be interpreted as a refractory current or an adaptation current, respectively. The a parameter was fixed to 0.001 nS or 30 nS, corresponding to a system losing stability via a saddle-node or Andronov-Hopf bifurcation, respectively. The range of reset parameters considered was limited to realistic values for spiking neurons ($V_r \in [-70, -40]$ mV, and $b \in [0, 400]$ pA).

The adaptive models ($\tau_w = 100$ ms, Fig. 2.6a, 2.6c, 2.6e) have bursting firing patterns extending to larger areas in the parameter space than the refractory models ($\tau_w = 5$ ms, Fig. 2.6b, 2.6d, 2.6f), which have their diversity almost uniquely constrained to resets above V_T . In all cases, the border between bursting (initial or regular) and tonic-adapting firing patterns can be approximated by a curve with linear plus exponential terms, similar to a shifted version of the V -nullcline. This can be related to the qualitative arguments in Sect. 2.3 (see subsection on initial bursting above and Fig. 2.4e as well as the section 2.5 below). In Fig. 2.6a and 2.6b, the tonic spiking (red areas) at low b and at high b corresponds to tonic with sharp resets and tonic with broad resets, respectively.

An adaptive current with high subthreshold adaptation yields predominantly bursting (initial and regular) and adapting firing patterns. However, a strong subthreshold adaptation is not sufficient on its own to model adapting and bursting patterns (Fig. 2.6d), but needs to be combined with high values of V_r . The refractory time constant in that case is often smaller than the interspike interval, this prevents cumulative increase of w since the adaptation current decays almost completely between spikes. Parameters combinations leading to irregular spiking (black pixels in Fig. 2.6) make a sparse and patchy structure that is entirely contained

in the larger region of regular bursting. This firing pattern appears predominantly close to the border between initial and regular bursting.

Delayed spiking and transient spiking are absent from Fig. 2.6 because these firing patterns do not exist when the current is well above the rheobase. Delayed spiking will appear at injection current slightly above the rheobase and are more salient at low or negative a . Transient spiking is produced with current not sufficient to make the stable fixed point loose stability, and this firing pattern depends heavily on a as it can occur only for sizable a . The accelerating firing pattern is also associated with current close to rheobase. Nevertheless Fig. 2.6a shows that accelerating firing pattern persists at higher step current amplitudes given that b is very small and V_r is within a small range below V_T .

2.5 Boundaries in Parameter Space

In Fig. 2.6c, the boundary between initial bursting and tonic spiking appears as a line that can be described as linear plus exponential terms. Hence the shape of the boundary reflects the shape of the V -nullcline. In this section we analyze why this is so.

To study the relation between the the boundary and the shape of the V -nullcline we consider a slightly simpler neuron model with a piecewise linear V -nullcline:

$$\begin{aligned} C \frac{dV}{dt} &= F(V) - w + I \\ \tau_w \frac{dw}{dt} &= a(V - E_L) - w. \end{aligned} \quad (2.10)$$

where the nonlinearity in the first equation is a piecewise linear function

$$F(V) = \begin{cases} -g_L(V - E_L) & \text{if } V \leq V_T \\ g_L \Delta_T (V - E) & \text{otherwise} \end{cases} \quad (2.11)$$

with $E = V_T + \frac{V_T - E_L}{\Delta_T}$

We assume that the evolution of the w -variable is much slower than the evolution of the voltage V . The voltage equation is characterized by a time scale $\tau = C/g_L$, the w variable by a time scale τ_w . We observe that, in the limit of separation of timescales ($\tau/\tau_w \ll 1$), the trajectories in the (V, w) phase space are nearly horizontal since $dV/dt \gg dw/dt$ - unless one of the trajectories gets close to the V -nullcline. Note that, by definition dV/dt vanishes on the V -nullcline, and because of continuity, dV/dt is small in the neighborhood of the nullcline. It can be shown that, in the limit of a separation of time scales, the trajectory follows the left branch V nullcline staying at a vertical distance $X(V)$

$$X(V) = \frac{\tau}{\tau_w g_L} [I - (a + g_L)(V - E_L)]. \quad (2.12)$$

below it (e.g., Gerstner and Kistler [2002]). In the case of piecewise linear nullclines, the trajectory can hence be obtained analytically.

More importantly for the present study, we can use the above result to get an approximation for the evolution of the adaptation variable w during spiking. After a spike, a trajectory restarts at the reset point (V_r, w_r) with a w -value w_r calculated from the reset condition 2.3. Along the trajectory, both V and w change. A trajectory ends if V hits the upper bound $V = 0$ at which point the dynamics is reset and the next trajectory starts. We focus on a function M that maps the starting point (V_r, w_r) of the trajectory to the value w_e of the adaptation variable at the end of the trajectory. We note that this function depends on V_r and write $w_e = M(w_r, V_r)$.

Let us now calculate w_e . With $F(V)$ defined in Eq. 2.11 it is useful to split the problem in two cases depending on whether the voltage reset V_r leads to a value above or below V_T . If $V_r < V_T$, separation of timescales means that V will increase too fast for w to change significantly unless the trajectory encounters the V -nullcline. This is illustrated in Fig. 2.7b where we see two trajectories: one with low w_r and another with high w_r . The former travels straight to high voltages with little change in w , the latter approaches rapidly the neighborhood of the V -nullcline which it follows at a distance $X(V)$ before going to high V . Therefore, with $V_r \leq V_T$:

$$M(w_r, V_r) = \begin{cases} w_r & \text{if } w_r < F(V_T) + I - X(V_T), \\ F(V_T) + I - X(V_T) & \text{otherwise.} \end{cases} \quad (2.13)$$

If $V_r > V_T$, then the trajectory evolves (nearly) horizontally to the right (increasing potential) unless the initial point is situated above the V -nullcline in which case the movement starts towards the left. This is illustrated in Fig. 2.7c where we see that trajectories starting above the V -nullcline again follow $F(V) + I$ with a distance $X(V)$. This gives, for $V_r > V_T$:

$$M(w_r, V_r) = \begin{cases} w_r & \text{if } w_r < F(V_r) + I, \\ F(V_T) + I - X(V_T) & \text{otherwise.} \end{cases} \quad (2.14)$$

Given M , it is now straightforward to classify tonic, initial bursting or regular bursting according to the definitions of mentioned above. We simply have to note that (i) w_e at the moment of the first spike is still close to its resting value $w_e \approx 0$; and (ii) in the case of periodic spiking the reset w_r is given by $w_r = w_e + b$. Hence, If $V_r \leq V_T$

- tonic spiking with broad resets if $b > F(V_r) + I$,
- tonic spiking with sharp resets if $b < F(V_r) - F(V_T) + X(V_T)$,
- initial bursting if $F(V_r) + I > b > F(V_r) - F(V_T) + X(V_T)$.

If $V_r > V_T$

- tonic spiking with broad resets if $b > F(V_r) + I$,
- regular bursting if $b < F(V_r) - F(V_T) + X(V_T)$,

- initial bursting if $F(V_r) + I > b > F(V_r) - F(V_T) + X(V_T)$.

These relations can be used to draw a picture of the parameter space for the case of a piecewise linear V -nullcline $F(V)$ (Fig. 2.7a) which holds in the limit of separation of timescales. Note that in the AdEx model of Fig. 2.6c we have $\tau/\tau_w = 0.1$ so that the assumption of a separation of timescales is justifiable. Indeed, in the AdEx model the boundaries between the firing patterns in Fig. 2.6c reflects the shape of the nullclines as should be expected based on the results of this section.

2.6 Comparison with Cortical Neurons

We can test that the AdEx model accurately reproduces the firing patterns of real neurons by comparing with experimental injections of step currents into neurons of the cortex (data, courtesy of Henry Markram, Maria Toledo-Rodriguez and Felix Schürmann, see Markram *et al.* (2004) and Toledo-Rodriguez *et al.* (2004) for the complete details on the experiments). Briefly, the experiments consist of 2-5 repetitions of 2 second step current injections with three different amplitudes. The amplitude of the steps ranged from 100 to 200 pA. The electrophysiological class was defined according to Markram *et al.* (2004) for the inhibitory neurons and according to Connors and Gutnick (1990) for the excitatory neurons. We will be considering only three different classes. In particular, we will compare the AdEx with two types of inhibitory neurons (continuous accommodating, cAD, and continuous non-adapting, cNA) and one type of excitatory neuron (regular spiking, RS).

In order to compare the experiments with the AdEx model, we fit the parameters of the model neuron per neuron. The fitting method was chosen for the ease of implementation and for the capability to handle an optimization problem with many local minima (see further below). Optimized model traces are compared with the experimental traces in Fig. 2.8. We can see that the AdEx model offers a good qualitative match akin to optimized Hodgkin and Huxley models (Druckmann *et al.* [2007]). The optimized parameters for each chosen cell are given in Table 2.1. From this table, we see that the inhibitory cells correspond to smaller membrane capacitance, consistent with the smaller size of these cells. The parameter a is low for all three cells and does not strongly influence the features used as optimization criteria. The timescale of the adaptation is the largest for the RS cells, which is expected because of the slow adaptation currents known to be present in these cells. Figures 2.8d-f show the steady-state $f-I$ curve for the models corresponding to three different cells. We see that the two interneuron models show type II $f-I$ curves of a high slope, while the RS neuron have an $f-I$ curve of I slow slope. Although the full $f-I$ curve was not available with the experiments we used, we can see a qualitative link with the stereotypical $f-I$ curves that are to be expected for these types of cells (Tateno *et al.* [2004]).

2.6. Comparison with Cortical Neurons

Table 2.1: Parameters and cost for fits shown in Fig. 2.8 and for firing pattern example shown in Fig. 2.4.

Type	C (pF)	g_L (nS)	E_L (mV)	V_T (mV)	Δ_T (mV)	a (nS)
Fig. 2.8, cNA	59	2.9	-62	-42	3.0	1.8
Fig. 2.8, cAD	83	1.7	-59	-56	5.5	2.0
Fig. 2.8, RS	104	4.3	-65	-52	0.8	-0.8
Fig. 2.4a	200	10	-70	-50	2	2
Fig. 2.4b	200	12	-70	-50	2	2
Fig. 2.4c	130	18	-58	-50	2	4
Fig. 2.4d	200	10	-58	-50	2	2
Fig. 2.4e	200	12	-70	-50	2	-10
Fig. 2.4f	200	12	-70	-50	2	-6
Fig. 2.4g	100	10	-65	-50	2	-10
Fig. 2.4h, Fig.2.5	100	12	-60	-50	2	-11

Type	τ_w (ms)	b (pA)	V_r (mV)	$C(\beta)$	I (pA)
Fig. 2.8, cNA	16	61	-54	8.4	184
Fig. 2.8, cAD	41	55	-54	10.4	116
Fig. 2.8, RS	88	65	-53	10.4	98
Fig. 2.4a	30	0	-58	-	500
Fig. 2.4b	300	60	-58	-	500
Fig. 2.4c	150	120	-50	-	400
Fig. 2.4d	120	100	-46	-	210
Fig. 2.4e	300	0	-58	-	300
Fig. 2.4f	300	0	-58	-	110
Fig. 2.4g	90	30	-47	-	350
Fig. 2.4h, Fig.2.5	130	30	-48	-	160

Optimization Methods Brette and Gerstner (2005) have proposed an experimental protocol to determine the parameters of the model. The experiments used here were done before the work of Brette and Gerstner was published, and no fitting method based on step current stimulation only has been reported for the AdEx model. Here we will describe the optimization methods used to find the best set of parameters (β). Inspired from Druckmann et al. [2007] and Vanier and Bower [1999], we used a MATLAB implementation of a genetic algorithm (The Mathworks, Natick, MA) to solve the optimization problem. The cost associated with a parameter set, $C(\beta)$, was defined with eight features of the observed responses. The six features are:

- f_1 : Number of spikes, n ,
- f_2 : First spike latency,
- f_3 : First inter-spike interval,
- f_4 : Last inter-spike interval,
- f_5 : First inter-spike minimum potential,

- f_6 : Waveform before the first spike.

The spike times were defined as zero-crossings of the voltage and the inter-spike minimum potential was taken to be the lowest voltage observed between the first and the second spike. For each feature we define a $\bar{\chi}^2$ which averages the χ^2 value across all three step current amplitudes. Illustrating this with the first feature we have:

$$\bar{\chi}_1^2 = \frac{1}{3} \sum_{j=1}^3 \frac{\langle n_j^{(obs)} \rangle - n_j^{(\beta)}}{Var[n_j^{(obs)}]} \quad (2.15)$$

where n_j is the number of spikes for stimulus amplitude j observed in experimental traces (*obs*) or in a model with parameters β . The angular brackets denote the average and Var is the variance of the observed features across the repetitions. A similar equation can be written for each feature, except for feature 6:

$$\bar{\chi}_6^2 = \frac{1}{3} \sum_{j=1}^3 \frac{\int (V^{(obs)}(t) - V^{(\beta)}(t))^2 dt}{(0.2mV)^2} \quad (2.16)$$

where the integral runs from the onset of the step to two standard deviation before the first spike mean latency. Finally, the cost associated with a parameter set is:

$$Cost = \sum_{i=1}^6 \bar{\chi}_i^2. \quad (2.17)$$

This is minimized with a genetic algorithm with a population of 100 individuals for 200 generations and a crossover fraction of 0.6.

2.7 Discussion

The AdEx model can produce multiple firing patterns depending on parameters. The model neuron can exhibit initial bursting, regularly bursting, tonic spiking, adapting, accelerating, irregular spiking, or show delayed initiation. We have drawn clear definitions of these firing patterns in terms of two types of spiking trajectories. The two types of spiking trajectories depends on whether the adaptation current immediately after spiking is strong enough to make the membrane potential decrease slowly before starting to increase in preparation for the next spike. The explicit distinction between continuous adaptation, initial bursting, and the description of irregular spiking extends previous work on firing patterns in model neurons (Izhikevich and Desai [2003], Izhikevich [2007], Touboul and Brette [2008]). Whereas Touboul and Brette [2009] focuses on types of excitability, stationary limit cycles and their relation to the underlying bifurcations, we studied the transient response to step currents since this is the prevalent experimental paradigm for classifying firing patterns.

Chaotic responses have been observed in neuron models as well as real neurons under periodic

forcing (Holden [1986], Brumberg and Gutkin [2007]). The type of chaos that we report does not depend on a periodic forcing. Rather it is closely related to the chaos observed in Hodgkin-Huxley equations (Chay and Rinzel [1985]). The presence of chaos in the AdEx model is restricted to small regions of the parameter space (Fig. 2.6). It is nevertheless possible that neurons access a similar chaotic regime. Evidences for the existence of chaos at all levels of the nervous system accumulate, despite technical limitations (Korn and Faure [2003]). Fig. 2.5 provides a method with which experiments could be made to verify the existence of the irregular firing pattern described. The irregular firing pattern appears in a narrow range of step current amplitudes (Fig. 2.5a) surrounded by a bursting at smaller and stronger current amplitude. Irregularity can be distinguished from noise by looking at the interval map, which is a thin continuous line (Fig. 2.5b) for the irregular spiking AdEx model. The AdEx model can be fit with good agreement to three types of cortical neurons, as can be seen from the comparison between model and experimental traces on the time scale of seconds (two first columns in Fig. 2.8). A closer look on the overlay unveils some discrepancies: the spike initiation of the cAD interneurons is not fully captured by the model. The first interspike interval is too long in the case for the RS cell, while the last interspike interval is too short in the cNA cell. This makes the trade-off between different features evident, and our simple model cannot fit with high precision both the initial burst of RS cells and the broad SAP observed in these cells. Important improvement is to be expected by adding another adaptation current but here we restricted the investigation to the capability of only one adaptation variable.

Our study has shown that the value of the voltage reset is an important parameter of the model. The voltage reset allows to include some, but not all refractory properties in a neuron model (Badel et al. [2007]). Refractoriness arises from the combined contributions of increased firing threshold, reduced membrane time constant and hyperpolarized voltage. If the point-neuron model is replaced by a multi-compartment model, the effective voltage reset depends also on the amount of current flowing from the dendritic compartments into the soma. This can lead to an After Spike Depolarization (ADP). The most convenient choice for a voltage reset would be the maximum of the ADP. Mainen and Sejnowski (1996) have shown previously that variable electrical coupling with an active dendritic compartment tunes the ADP and can be made responsible for multiple firing patterns. The AdEx provides a simpler framework to switch between firing patterns by changing reset parameters - and these changes could indeed reflect, amongst other influences, changes in electrical coupling.

We can conclude that the diversity of firing patterns is explained in the AdEX model by simple dynamical processes that can be completely analyzed in the phase plane. The AdEx model represents an attractive candidate neuron model for use in large-scale network simulations. Earlier studies have shown that it is sufficiently accurate for the prediction of spike timing when a RS pyramidal neuron receives noisy current injection at the soma (Jolivet et al. [2006]). Network simulations of large systems such as a column of the barrel cortex (Markram [2006]), or thalamocortical systems (Hill and Tononi [2005], Izhikevich and Edelman [2008]) require the description of different neuron types. In this paper we addressed this issue by providing parameter sets that describe three types of cortical neurons, within the framework of a single

Chapter 2. Firing patterns in the adaptive exponential integrate-and-fire model

and attractively simple neuron model. Further work will be needed to extend this to a larger number of neuron types and stimulation paradigms.




Figure 2.6: Parameter space exploration of the four bifurcation parameters. Tonic spiking in red, adapting in yellow, initial bursting in green, regular bursting in cyan, irregular spiking in black and accelerating in blue. The four-dimensional parameter space was reduced to six relevant planes: (a) adaptive time constant ($\tau_w = 100$ ms) and negative a ($a = -5$ nS), (b) refractory time constant ($\tau_w = 5$ ms) and negative a ($a = -5$ nS), (c) adaptive time constant ($\tau_w = 100$ ms) and low a ($a = 0.001$ nS), (d) refractory time constant ($\tau_w = 5$ ms) and low a ($a = 0.001$ nS), (e) adaptive time constant ($\tau_w = 100$ ms) and high a ($a = 30$ nS), (f) refractory time constant ($\tau_w = 5$ ms) and high a ($a = 30$ nS). The firing pattern was classified for injection of current twice the rheobase, according to the criteria exposed in Sect. 2.3.




Figure 2.7: Phase diagram and parameter space for a piecewise linear V -nullcline in the limit of separation of timescales. (a) The firing pattern predicted from the analytical expression given in the text for the limit of separation of timescales is shown with the same color scheme as in Fig. 2.6. (b) Trajectories for reset points where $V_r \leq V_T$ go to high potential with little change in w unless they must contour the V -nullcline. (c) When $V_r > V_T$ the trajectories go along the V -nullcline only when the reset point is situated above the $\frac{dV}{dt} = 0$. In both cases, the function M plateaus at a value $w_c = F(V_T) - X(V)$. The trajectories follow the V nullcline at a distance $X(V)$, as illustrated in the inset of (b). ($C = 100$ pF, $g_L = 10$ nS, $E_L = -70$ mV, $V_T = -50$ mV, $a = 0$ nS, $\Delta = 3$, $\tau_w = 2000$ ms).




Figure 2.8: Comparison of the AdEx with three types of cortical neurons on step current injections. From left to right: experimental traces (red), AdEx model (blue), and overlay of the traces during onset and offset of the current step. From top to bottom: cNA (a), cAD (b), and RS (c). The left scale bar shows 20 mV and 300 ms, the scale bar for the overlays shows 20 mV and 20 ms. The current injections corresponds to 150 pA for cNA, 105 pA for cAD, and 130 pA for RS. Only one of the 5 repetitions are shown for clarity. Across multiple repetitions of the same stimulus, the time of the first spike or the first interspike intervals may jitter around what is seen on this figure. The f - I curves (d, e, f) of the fitted models show a steeper slope for the interneurons (d and e, cNA and cAD respectively) and a slow, type-I slope for the RS cells (f).

3 Can We Predict Every Spike?

Is it possible to predict the spike times of a neuron with millisecond precision? ¹ In the classical picture of rate coding [Adrian, 1928], single spikes do not play a role, and the question would have to be answered negatively. For rate coding in a single-neuron, the relevant quantity to encode a stimulus such as pressure onto a touch sensor in the skin [Adrian, 1928] or presence of a light bar in the receptive field of a visual neuron [Hubel and Wiesel, 1959] is the number of spikes a neuron emits in a short time window of, e.g., 100ms. The timing of the spikes is considered as irrelevant. However, over the last 20 years many researchers have shown that it is not only the temporally averaged firing rate that carries information about the stimulus, but also the exact timing of spikes. For example, spike timing has shown to be relevant to encode force amplitude and direction in touch sensors of the skin [Johansson and Birznieks, 2004] as well as the whole-field visual movements [Bialek et al., 1991] or object movement [Gollisch and Meister, 2008] in visual neurons.

If spike timing is important, a whole series of questions arises: What is the precision of spike timing if the same stimulus is repeated several times? Do spikes always appear at the same time? What would be a sensible measure of spike timing precision and reliability? Can a neuron model match the spike timing precision of a real neuron? Does it matter which neuron or what stimulus we take? If so, what would be a useful stimulus?

To answer these related questions, let us think of the following experimental protocol. An experimentalist injects a time-dependent input of, say, 20 second duration into a single neuron. The neuron responds with spikes. The experimentalist now repeats the same stimulus sequence several times. At each repetition, the neuron responds with a spike train that may or may not look similar to the previous one: some spikes appear at exactly the same time during the stimulus sequence, some are missing, some are shifted by a few millisecond or appear at a completely different time. The information derived from this type of experiment which dates back to Bryant and Segundo (1976) and has been popularized by Mainen and Sejnowski (1995) should be sufficient to answer questions regarding precision and reliability of spike timing.

¹Text copied from Naud, R., and Gerstner, W., chapter to appear in *Spike Timing: Mechanisms and Function*, CRC Press (full citation in the References).

But here comes the challenge (Fig. 3.1). Let us suppose that I give you the time course of the input as well as the neuron's response in each trial, but only for the first 10 seconds of the data. For the second half of the stimulus sequence, I give you only the time course of the input. Your task is to predict the timing of the spikes of neuron.

Will you be able to predict the timing of the spikes using an appropriate neuron model? Is your model as reliable and as precise as the real neuron? What would be the best model to choose so as to solve the task?

The above challenge has been turned into a single-neuron modeling competition that was first run by Brain Mind Institute at the Ecole Polytechnique Fédérale de Lausanne (EPFL) in Switzerland [Jolivet et al., 2008b,a] and was officially handed over to the International Neuroinformatic Coordinating Facility (INCF) in Sweden in 2009. In this Chapter, we recapitulate the main questions and findings related to predicting spike times, with a special focus on the spike-time prediction competition of 2009 [Gerstner and Naud, 2009]. If not specified otherwise, the term spike timing competition refers in the following to part A of the 2009 competition, if we specify a different year we imply part A of the competitions in 2008 or 2007.

3.1 What is a good stimulus to probe neurons?

In classical electrophysiological experiment, an artificially generated input is injected in a neuron *in vitro* (Fig. 3.1). In principle, the time course of the input can be chosen arbitrarily and could consist of short or long steps of different amplitudes, sequences of steps, ramps, white noise, filtered noise or whatever comes to mind. But what is a 'good' stimulus?

Since the work of Hodgkin and Huxley [Hodgkin and Huxley, 1952] electrophysiologists have been using steps and ramps to characterize single-neuron responses. These types of stimulations are helpful to systematically probe the gating dynamics of ion channels under pharmacological manipulation. They can also give a qualitative classification of neuronal responses in terms of intrinsic firing patterns such as regular, fast-spiking, bursting [Connors and Gutnick, 1990, Markram et al., 2004] but they have very little resemblance with the type of stimulus a neuron would receive in its natural environment.

Inspired by signal processing theory, the pioneering studies of Bryant and Segundo (1976) and of Marmarelis and Marmarelis (1978) used white-noise stimulation instead of step currents. However, if the aim is to drive a neuron with a stimulus that resembles as much as possible the input it would receive an *in vivo* situation, a white-noise stimulus is not sufficient. Rather, a stimulus at the soma should replace the total current flowing from the synapses to the soma while the neuron receives presynaptic input. Following a line of earlier research [Stein, 1967a, Poliakov et al., 1996, Destexhe et al., 2003, Jolivet et al., 2006], the first spike timing competition in 2007 used an Ornstein-Uhlenbeck current injection with various means and variance to mimic the combined effect of a large number of synapses [Stein, 1967a, Jolivet et al., 2008a]. In 2008, the competition was modified [Jolivet et al., 2008a] to replace the dynamic current by




Figure 3.1: Schematic representation of the spike-timing prediction challenge. The same time-dependent input stimulus (left) is given to a mathematical neuron model and to a real neuron in an electrophysiological experiment. Part of the response of the real neuron is used to optimize the model parameters. The remaining part of the stimulus is injected into the model so as to predict the spike-times of the real neuron. The mathematical neuron model illustrated here is made of a linear filter of the input current (bottom full trace) and a dynamic threshold (dashed black line). Fig. adapted from [Gerstner and Naud, 2009]




Figure 3.2: Counting coincident spikes for the computation of the coincidence rate Γ_{nm} . The predicted spike train (bottom) is compared to the recorded spike train (top). A predicted spike is said to be coincident (black) if it falls between $\pm\Delta$ of a recorded spike and if that recorded spike was not counted as coincident with any other predicted spike. The prediction can miss recorded spikes (gray, top) or contain generate extra spikes (bottom, gray). Here the total number of coincident spikes is $N_{\text{coinc}} = 3$ while there were $N_m = 6$ spikes in the predicted and $N_n = 6$ spikes in the experimental spike train. (Adapted from Figure 1.1 in Jolivet [2005]).

dynamic inhibitory and excitatory conductances using dynamic clamp Destexhe et al. [2003]. Then in 2009 the injected current was changed to a current produced by the simulation of six populations of presynaptic neurons changing their firing rate every 200-500 ms.

3.2 How can we measure spike timing precision and reliability?

Suppose that a single neuron is driven with multiple repetitions of the same time-dependent stimulus. The response of the neuron is recorded in each trial, so that the stimulation protocol builds up a database containing one spike train for each repetition.

If we compare the spike trains across several repetitions, different types of variability are seen depending on the system and the variance of the input [Bryant and Segundo, 1976, Mainen et al., 1995, Jolivet et al., 2006]. Some spikes are seen at the same time for all repetitions, others appear at a specific time on half the repetitions and yet others do not seem to be related to a specific time. Several questions arise: First, how can we *quantitatively* compare one spike

3.2. How can we measure spike timing precision and reliability?

train of the neuron with another one recorded during a later repetition of the same stimulus sequence? Second, how can we quantify the reliability across the set of all spike trains recorded with the same stimulus? Finally, how can we compare the set of spike trains generated from neuronal recordings with a similar set of spike trains generated by a mathematical neuron model?

These are crucial questions which can be answered in different ways. One can compare one spike train with another one based on global features such as the intrinsic firing patterns in response to step stimuli [Connors and Gutnick, 1990, Markram et al., 2004] and ask whether a neuron model is able to reproduce the same intrinsic firing patterns [Izhikevich, 2007, Naud et al., 2008]². One can focus the quantitative comparison on the shape of spikes and adaptation patterns [Druckmann et al., 2007], on the inter-spike interval distribution [Chacron et al., 2005a] or the spike-count variability [Softky and Koch, 1993, Schaette et al., 2005].

If one focuses on spike timing, one may want to apply methods that compare spike trains in terms of a spike-train metrics [Victor and Purpura, 1996] or the coincidence rate [Kistler et al., 1997]. Both measures can be used to compare a spike train from a recorded neuron in repetition n with another spike train recorded in repetition m . Both measures can also be used to compare a spike train derived from a neuron model with a spike train recorded in one of the sessions with the real neuron. Obviously, a model which achieves an optimal match in terms of spike-train metrics will automatically account for global features of the spike trains, such as inter-spike interval distributions.

In the INCF competition the average coincident rate was used to quantify spike-time prediction performance. The average coincidence rate can be seen as a similarity measure between pairs of spike trains that is finally averaged across all available pairs. To compute the pairwise coincidence rate, one first finds the number of spikes from the model that fall within an interval of plus or minus 4 ms around a spike from the real neuron. This is called the number of coincident events N_{nm} (at resolution $\Delta = 4ms$). The coincidence rate is the ratio of the number of coincident events over the averaged number of events $0.5(N_n + N_m)$, where N_n is the number of spikes in the neuron spike train and N_m is the number of spike in the model spike train (Fig. 3.2). This ratio is then scaled by the number of coincident events, $N_{\text{poisson}} = 2\Delta \cdot N_m N_n / T$, that are expected from a Poisson model that fires stochastically at a fixed rate N_m / T . The scaled coincidence rate is

$$\Gamma_{nm} = \frac{N_{nm} - N_{\text{poisson}}}{\frac{1}{2}(1 - N_{\text{poisson}}/N_n)(N_n + N_m)}. \quad (3.1)$$

Finally, the pairwise coincidence rate Γ_{nm} is then averaged across all the possible pairings of spike trains from the model with those of the neuron and this gives the averaged coincidence rate $\overline{\Gamma_{nm}}$.

The coincidence rate Γ_{nm} enables us to compare the spike timings of the model (subscript m)

²Chapter 2

with that of the neuron (subscript n). If we want to know how reliable the neuron itself is, we need to measure how similar the spike trains are between two trials. To do the comparison between one neuronal spike train (index n) to another neuronal spike train (index n), we can use again the same coincidence measure that we now label Γ_{nn} so as to indicate the coincidence rate between two neuronal spike trains. The neuron-to-neuron coincidence rate, averaged over all pairs of available spike trains for the same stimulus, is a measure of the intrinsic reliability of the neuron and denoted as $\overline{\Gamma_{nn}}$. It provides an upper bound for modelling: on average, a neuron model cannot predict spikes better than the neuron itself. The averaged coincidence rate $\overline{\Gamma_{nm}}$ needs therefore to be compared to the upper bound provided by the intrinsic reliability $\overline{\Gamma_{nn}}$ of the neuron. Scaling $\overline{\Gamma_{nm}}$ by $\overline{\Gamma_{nn}}$ gives a number that can be interpreted as the fraction of the predictable spikes that are predicted by the model.

The coincidence rate, like the majority of spike-train metrics, has a time-scale parameter. In the above equation, the time-scale parameter Δ regulates the size of the coincidence window, and thus the level of precision of the prediction. For a very small coincidence window, the number of coincidences N_{coinc} goes to zero due to a jitter in spike timing and the finite number of spike trains. For a very large coincidence window, the coincidence rate goes to zero because it becomes insensitive to specific times of the spikes so that the difference between the prediction of a precisely tuned neuron model and that of a Poisson model with constant rate vanishes. In a large range between these two extrema, however, the model-to-neuron coincidence rate is significantly positive. Moreover, over the range roughly from 2 to 15 ms depending on the neuron and on experimental conditions [Jolivet et al., 2006, 2008b], the results measured in terms of $\overline{\Gamma_{nm}}$ do not depend on the choice of the time window Δ . In the INCF competition a window of $\Delta = 4\text{ms}$ was chosen.

It is useful to distinguish measures such as the coincidence rate [Kistler et al., 1997] or the spike train metrics [Victor and Purpura, 1996] which are both based on a comparison of a single spike train A with a second spike train B from measures that first average across all repetitions of an experiment to calculate the Peri-Stimulus Time-Histogram (PSTH) [Eggermont et al., 1983] before a comparison of the PSTH of neuron A with that of a neuron B (or of a model neuron) is performed. In the INCF challenge, rankings were based on the pairwise comparison of a model spike train with a real spike train, averaged *a posteriori* across all repetitions of the experiment so as to determine the average coincidence rate.

Does the average coincidence rate correspond to a comparison of the PSTH? The PSTH is calculated by averaging all the independent responses to the same input. A smoothing filter is then applied to the averaged spike trains. In many neuronal systems, the PSTH is made of a series of peaks and plateaus. The peaks correspond to spikes always coming at a precise time and the plateaus corresponds to times where spikes are emitted with no specific timing. A model reproducing such a PSTH can be said to predict the spike times because such a model will emit a spike precisely at times where the neuron emits precisely timed spikes. Indeed, normalized spike-train similarity measures such as $\overline{\Gamma_{nm}}/\overline{\Gamma_{nn}}$ calculate a quantity very similar to the variance of the experimental PSTH that is explained by the model PSTH [Naud et al.,

To appear; 2012]³. The time-scale parameter of the similarity measure is equivalent to the filter time-scale applied for smoothing the PSTH. The main difference between $\overline{\Gamma_{nm}}/\overline{\Gamma_{nn}}$ and the comparison of PSTHs is that an optimization of neuron models based on the comparison of PSTH attempts to match the spike-timing variability of the model to the variability of the data. In contrast, an optimization of neuron models based on the normalized coincidence rate gives a slight advantage to deterministic neuron models, i.e., those that do not correctly reproduce the intrinsic variability of the data [Naud et al., To appear; 2012]⁴.

Being aware of the similarity between PSTH comparison and coincidence rate scaled with intrinsic reliability is important to relate different studies to each other even though they use different evaluations criteria. It should be kept in mind, that there is an over-estimation of the prediction performance for deterministic models using the scaled coincidence rate $\overline{\Gamma_{nm}}/\overline{\Gamma_{nn}}$ with respect to variance-explained of the PSTH [Naud et al., To appear; 2012]⁵. However, there is also a small-sample bias that can over-estimate the variance explained when comparing PSTHs [David and Gallant, 2005, Petersen et al., 2008, Naud et al., To appear; 2012]⁶. Values of above 100 % are possible because of the bias, but also because a model can be fitted independently on each repetition and thus allowing to take into account the experimental drifts.

3.3 What are good neuron models?

Across several years and editions of the single-neuron modelling competition, various models participated in the challenge. The models ranged from the very simple integrate-and-fire models to complete biophysical models using the Hodgkin-and-Huxley formalism and a 3D reconstruction of the morphology from another neuron of the same class. The number of state-variables goes from one for the simplest models to a few hundred depending on the number of ion-channel types modelled and the number of compartments used in the discretization of space. Similarly, the number of parameters scaled with the number of state-variables with five parameters for the simplest models and close to a hundred parameters for the biophysical models. Between the two extremes, various other models were used. For instance, adding non-linear threshold, spike-triggered adaptation and sub-threshold adaptation, the Izhikevich model [Izhikevich, 2004] and the AdEx [Brette and Gerstner, 2005] models were used by some of the participants. Simpler versions of the complete biophysical models were also devised by reducing the number of ion-channels and the number of compartments (in the same line as Pospischil et al. [2008]).

A good prediction stems from the union of a good model and an efficient fitting method. The method for finding the optimal parameters should be efficient in the sense of providing a single set of optimal parameters with small computing resources. Many different methods

³Chapter 4

⁴*Idem*

⁵*Idem*

⁶*Idem*

Chapter 3. Can We Predict Every Spike?

were used by the various participants, some using the action potential shape, the subthreshold voltage dynamics and the spike times as the observables to fit, others using only the set of spike times as observables to fit.

The participants that used biophysical models approached the problem by first constraining a significant number of parameters with published measurements of ion channel dynamics. Most of the biophysical-model participations thus reduced the parameter space to the somatic ion-channel densities only. This leaves a number of free parameters equal to the number of ion-channel species [Druckmann et al., 2007]. These remaining parameters are then fit either by hand-tuning, stochastic optimization algorithms such as the genetic algorithm or by exhaustive search when the number of ion-channels species is low.

For many of the simpler models, the optimization methods available are more efficient. Some of the most effective participations performed an exhaustive search on a small number of crucial parameters [Kobayashi et al., 2009]. Convex optimization algorithms can be used to maximize the likelihood of observing the spike times [Paninski et al., 2004] this method lead to some of the top-ranking participations. Another noteworthy method for fitting involves a convex, two-step procedure where in the first step the optimal passive parameters and the spike shape are determined from the voltage trace, and on the second step the parameters regulating a dynamic threshold are determined by maximizing the likelihood of the observed spike trains [Mensi et al., Under Review]⁷.

The model that achieved the highest performance in Challenge A 2009 was an integrate-and-fire model provided with a dynamic threshold that jumps every times there is a spike and decays back to a baseline with three different time-constants [Kobayashi et al., 2009]. The ratio $\bar{\Gamma}_{nm}/\bar{\Gamma}_{nn}$ for this participation was of 76.2%. The participant winning this spike-time prediction competition had extracted the membrane time constant from the voltage trace, fixed the decay time-constants of the dynamic threshold to 10, 50 and 200 ms leaving three free parameters; one parameter regulating the amount by which the threshold jumps for each of the three time scales. The optimal set of the three parameters was found by conducting an exhaustive search for the set of parameters maximizing the average coincidence. This winning model happened to be the model with the smallest number of free parameters used in the competition. The small number of free parameters is not sufficient to explain the high performance, because for instance a simple leaky and integrate and fire cannot predict more than 38% on the same task. The winning participant used judicious insights in deciding which model to use, which parameters could be fixed *a priori* and which parameters required to be fitted to the specific neuron recorded.

A very small number of biophysical models participated in the competition, perhaps due to the difficulty of finding the optimal parameters for such complex models. There was a noteworthy participation using state-of-the-art optimization methods which was outside of the official competition because submitted after the deadline for the money prize. This

⁷Chapter 5

submission would have ranked third if it had been submitted before the deadline. Thus, within the framework of the competition, the prediction performance of simple models is as good if not slightly better than that of the biophysical models.

The relatively high prediction performance of models is not explained by the fact that the challenge is too easy. Off-the-shelf models such as a hand-tuned models of pyramidal neurons or a leaky-integrate-and-fire achieved a performance around $\bar{\Gamma}_{nm}/\bar{\Gamma}_{nn} = 40\%$. Perhaps the most important aspect for providing a good prediction in the competition was to take into account spike-frequency adaptation. The dynamic threshold of the winning submission is just one example, there are many other ways to implemented adaptation in single neurons and models that implement adaptation were systematically better than models that did not [Jolivet et al., 2008a].

So what is the best model? Accurate modelling of the refractory period is essential to predict the spike times [Kistler et al., 1997, Keat et al., 2001]. More generally, if the same neuron model has to predict spike timings for stimuli with different mean firing rates, the importance of spike-frequency adaptation was recognized explicitly [Pillow et al., 2005, Jolivet et al., 2006, 2008a]. State-of-the-art models now consist of models akin to the stochastic integrate-and-fire model but upgraded with an adaptation process. The adaptation makes the firing probability dependent on the timing of all recently emitted spikes. Such models are capable of predicting 75-100% of the predictable spikes (see next Section). This leaves little room for improving the accuracy of encoding models. Indeed, increasing the level of detail with conductance-based adaptation or Hodgkin-Huxley ion channels does not yield substantial increases in prediction performance [Druckmann et al., 2007, Mensi et al., Under Review]⁸.

3.4 Are all neurons predictable?

Can we predict the spike times in other systems than a Layer 5 pyramidal neuron in a cortical slice? The different editions of the competition showed that good spike-time prediction can be achieved for current injections of at least two types of dynamics ($\bar{\Gamma}_{nm}/\bar{\Gamma}_{nn} = 82.0\%$ in 2007 for Ornstein-Uhlenbeck dynamics, $\bar{\Gamma}_{nm}/\bar{\Gamma}_{nn} = 71.6\%$ in 2009 for multiple time-scale dynamics). The predictions are slightly better when the stimulation is given as fluctuating inhibitory and excitatory conductance ($\bar{\Gamma}_{nm}/\bar{\Gamma}_{nn} = 91.4\%$ for challenge A in 2008). The activity of L2/3 pyramidal neurons and non-Fast spiking GABA-ergic neurons can be predicted with similar performances [Mensi et al., Under Review]⁹. Prediction of the spike times of fast-spiking GABA-ergic neurons is systematically higher with $\bar{\Gamma}_{nm}/\bar{\Gamma}_{nn} =$ in the range of 100 percent in challenge B of 2009 (Fig. 3.3).

Real neurons receive their inputs from synapses distributed throughout their dendritic tree. The single-neuron model should model the dendritic integration of inputs. This dendritic integration is known to be highly non-linear especially in the thick tufted L5 pyramidal cells

⁸Chapter 5

⁹*Idem*




Figure 3.3: Predicting fast-spiking GABAergic neurons and pyramidal neurons from in 5 from the 2009 competition. A: The time-dependent current input models six populations of presynaptic neurons changing their firing rate every 200-500 ms. B-D Prediction of a GABAergic fast-spiking neuron. B: modelled (gray) and recorded (black) voltage traces. A zoom of 50 ms is shown in inset. C spike trains of each 7 repetitions of the recorded (black) compared modelled (gray) spike train. D PSTH of the model (gray) overlaid on the PSTH of the data (black) calculated from the average of the spike trains that are then filtered with a gaussian of 5 ms standard deviation. For the model we computed the PSTH from one thousand independent realizations while for the data we were restricted to the number of repetitions that could be recorded in the experiment. E-G same as for B-D but predicting the activity of a pyramidal neuron from the layer 5. (Figure a courtesy of Skander Mensi).

[Larkum et al., 1999]. To explore the dendritic dimension, dual electrode recordings were made with two independent injection sites: one in the soma and a second high in the dendritic tree. Spike-time prediction of $\bar{\Gamma}_{nm}/\bar{\Gamma}_{nn} = 83.8\%$ was achieved in Part C of the 2009 competition with a model similar to the one used by Larkum et al. [2004].

In the retina, spiking models of the Retinal Ganglion Cells (RGC) can predict 91 % of the variance of the PSTH [Pillow et al., 2005]. Similar performances have been observed in vivo where 41-92% of the variance of the PSTH Lateral Geniculate Nucleus (LGN) cells can be predicted from the activity of a single impinging RGC for spatially restricted visual stimulation [Carandini et al., 2007] as confirmed by submissions to Challenge D in the 2009 edition of the competition.

3.5 Conclusion

In summary, the prediction of precise spike timing on the millisecond time scale is similar to predicting the time-dependent firing rate on the milli-second time scale. High prediction performance is possible in many neuronal systems and depends strongly on the choice of neuron model and fitting method. One important model feature for high prediction performance is the presence of spike-frequency adaptation. The choice of the model formalism can also influence the fitting method that can be used. High quality prediction is most of the time associated with an efficient and convex fitting method.

Can we use these results to determine the best single-neuron model? The single-neuron model of choice should be able to generalize across all the different experimental protocols and also across the possible systems and neuron types with a mere change of the model's parameters. The original competition was rewarding only the participations that could generalize across more than one of the experimental protocols. The data of the challenge will remain available in the future for bench-marking purposes, leaving the possibility for such a deed to be accomplished¹⁰.

¹⁰<http://www.incf.org/>

4 Improved Similarity Measures for Small Sets of Spike Trains

Multiple measures have been developed to quantify the similarity between two spike trains.¹ These measures have been used for the quantification of the mismatch between neuron models and experiments as well as for the classification of neuronal responses in neuroprosthetic devices or electrophysiological experiments. Frequently, only a few spike trains are available in each class. We derive analytical expressions for the small-sample bias present when comparing estimators of the time-dependent firing intensity. We then exploit analogies between the comparison of firing intensities and previously used spike train metrics, and show that improved spike train measures can be successfully used for fitting neuron models to experimental data, for comparisons of spike trains, and classification of spike train data. In classification tasks the improved similarity measures can increase the recovered information. We demonstrate that when similarity measures are used for fitting mathematical models, all previous methods systematically underestimate the noise. Finally, we show a striking implication of this deterministic bias by re-evaluating the results of the Single-Neuron Prediction Challenge.

4.1 Introduction

In order to compare one spike train with another, many methods have been proposed; each relating to an underlying philosophy as to what feature of the spike train matters most. Comparing spike trains in terms of the time-series of interspike intervals [Victor and Purpura, 1996, 1997, Quiroga et al., 2002a, Kreuz et al., 2007] focuses on a local estimation of the firing rate. Comparing spike trains in terms of the presence of specific spike patterns [de Ruyter van Stevenick and Bialek, 1988, Lestienne, 1995, Victor and Purpura, 1996, 1997, Chi et al., 2007] is important if spike patterns such as sequences of action potentials have a greater importance than single spikes [Abeles, 1991, Lisman, 1997, Izhikevich et al., 2003, Eyherabide et al., 2009]. There are also many methods to compare spike trains in terms of spike timings: one can

¹Text copied from Naud, R., Gerhard, F., Mensi, S., and Gerstner, W., *Neural Computation*, to appear (full citation in the References).

Chapter 4. Improved Similarity Measures for Small Sets of Spike Trains

consider the minimum cost to transform one spike train into another by shifting, deleting or adding spikes [Victor and Purpura, 1996, 1997], one can also consider the similarities of spike trains in terms of how they would appear to a post-synaptic neuron [van Rossum, 2001, Houghton, 2009], or one could simply count the number of coincident spikes and consider the coincidence rate [Kistler et al., 1997, Quiroga et al., 2002a, Hunter and Milton, 2003].

Spike-train similarity measures were used for the classification of responses coming from different stimuli and thus to predict the stimulus that was presented to an animal (visual as in [Optican and Richmond, 1987] and in [Geisler et al., 1991] or auditory as in [Wang et al., 2007]). Decoding motor cortex activity with similarity measures can be used to predict the response of the animal such as the motion of its arm [Shpigelman et al., 2005, Eichhorn et al., 2004], which demonstrates the potential role of spike-train similarity measures for brain-machine interfaces and the construction of brain-controlled prosthesis. Quantifying the efficiency of the decoding can give an estimate of the amount of information recovered depending on the precision or the type of spike-train metric [de Ruyter van Stevenick and Bialek, 1988, Rieke et al., 1996, Victor and Purpura, 1996, 1997, Reich et al., 1997, 2000, 2001, DiLorenzo and Victor, 2003]. Spike-train similarity measures have also been used to quantify how well the spike-time predictions of a mathematical model match real data [Pillow et al., 2005, Jolivet et al., 2006, Badel et al., 2008b, Jolivet et al., 2008a, Gerstner and Naud, 2009, Kobayashi et al., 2009]. The study of neuronal noise and variability is yet another topic for which the problem of comparing spike trains arises (see [Tiesinga et al., 2008] for a review).

Neurons are said to have variable responses because different trials with the same stimulus give different spike trains. A small trial-to-trial variability is observed when the stimulus is directly injected into a neuron which receives no other input [Mainen and Sejnowski, 1995] and a much larger variability when one records a single neuron of the visual cortex during visual stimulation [Carandini, 2004]. The source of variability *in vivo* is probably channel noise [Faisal et al., 2008] in combination with cortical patterns impinging via a large number of synapses onto a single neuron [Calvin and Stevens, 1968, Arieli et al., 1996, Tsodyks and Markram, 1997]. Channel noise implies that spike generation is an intrinsically probabilistic process. One way to approach this problem is to average the spike-train similarity measures across several trials as in [Jolivet et al., 2006, 2008a], the other is to look at the mismatch between the probability distributions of the different spiking processes.

Let us start with an example of the first approach, a pairwise comparison of spike trains using two standard measures (Fig. 4.1). We have a set of 50 spike trains which could come from experiments but have been generated here for the sake of simplicity by a rate-normalized inhomogeneous Poisson process defined by a firing rate that is a superposition of Gaussians of different widths (Fig. 4.1a). After each spike, further spiking was blocked during an absolute refractory period of 3 ms. We now generate a second set of 50 spike trains with the same point process (Fig. 4.1b) and generate a third set with the widths of the Gaussians reduced a factor of four (Fig. 4.1c). Since the processes in Fig. 4.1a and Fig. 4.1b are identical, we might naïvely expect that the spike trains in Fig. 4.1a and Fig. 4.1b are more similar than

the spike trains generated with a different process in Fig. 4.1c. However, when we compare the spike trains using the well known Victor-Purpura metric [Victor and Purpura, 1996] and the coincidence rate [Kistler et al., 1997] as a dissimilarity or similarity measure, respectively, we find that both the Victor-Purpura metric and the coincidence factor imply that the spike trains from dissimilar processes are more similar than spike trains from the same process. There is an easy explanation for this apparent contradiction: the average distance between spike trains depends not only on the difference between the distributions, but also, on the variance of each distribution. For example, suppose that ten one-dimensional data points $\{x_1^{(1)}, \dots, x_{10}^{(1)}\}$ are drawn from a gaussian distribution of width $\sigma_1 = 1$ and another 10 data points $\{x_1^{(2)}, \dots, x_{10}^{(2)}\}$ from a Gaussian distribution of arbitrary width σ_2 . The average distance $\sum_{i,j}^{10} \|x_i^{(1)} - x_j^{(2)}\|^2$ is smaller if $\sigma_2 < \sigma_1$ than for equal width $\sigma_1 = \sigma_2$ (Fig. 4.1d). Similarly, asking for the average coincidence count across spike trains to be maximum is not the same as asking for the probability distributions to match.

The question arises whether it could be possible to directly compare distributions of spike trains, rather than pairwise distances. However, existing methods for estimating spike train probability distributions (parametric: [Paiva et al., 2009a] or entropy-based: [Optican and Richmond, 1987, Strong et al., 1998, Panzeri et al., 2007]) require the observation of a sufficiently high number of independent responses evoked by a given stimulus. To do this, one usually records the response of a given neuron to repeated presentations (trials) of the same stimulus. It is often impossible experimentally to accumulate enough trials to estimate these statistics adequately: When the neuron responses are recorded in patch-clamp, experimental drifts limit the total time of stable conditions to less than an hour, while on the other hand the presence of long-range adaptation in single neurons imposes a lower limit on the trial-to-trial wait period of at least 10 seconds [La Camera et al., 2006, Lundstrom et al., 2008]. Ignoring this lower limit will affect the independence of trials. This means that the similarity between the probabilistic spike processes often has to be calculated on a small number of samples where small-sample biases are known to occur.

Hence we arrive at the following dilemma: on one hand, we do not have enough data to estimate probability distributions of spike trains, on the other hand, optimisation based on naïve pairwise comparisons yields undesirable results (Fig. 4.1). One could conceive several ways to address this dilemma. One possibility would be to work with pairwise measures but use a Kilmogorov-Smirnov test on the within-class and across class distributions of pairwise distances. We briefly mention this method in Sect. 4.6. A second option is to base comparisons neither on the pairwise distances nor on the mismatch of the full probability distribution but only on the mean of the distributions, similar to some earlier approaches [Eggermont et al., 1983, Gawne et al., 1991, Brown et al., 2002, Paninski et al., 2005, Pillow et al., 2005, 2008]. This is the option explored in the main body of the present paper. Since in the case of spike trains the mean of the distribution is the Peri-Stimulus Time Histograms (PSTHs), we study how to measure differences between the PSTHs. We derive bias-corrected measures of vectorial distance and correlation coefficient between PSTHs and we estimate the variance of the distribution around the PSTH. A third option, which we study in depth is to apply

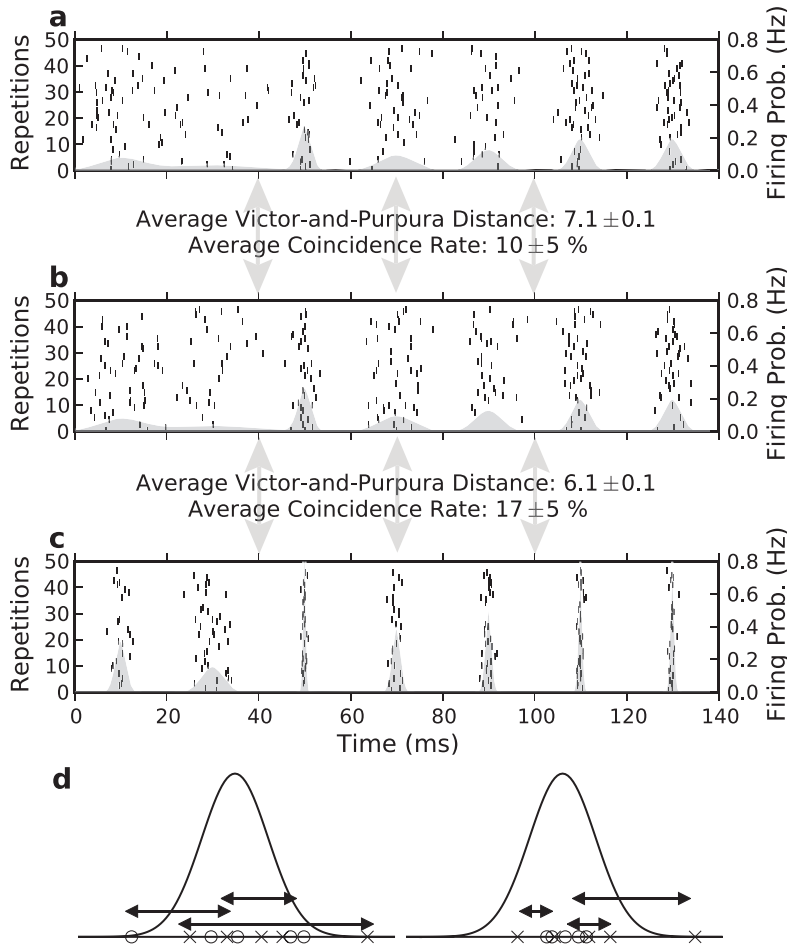


Figure 4.1: Similarity of Spike Trains Does Not Predict Similarity of Distributions **a** 50 spike trains (repetition number on vertical axis) containing an average of seven spikes (small vertical bars) per train were generated by a rate-normalized inhomogeneous Poisson processes with a refractory period of 3 ms. The instantaneous rate of the point process is indicated by the shaded area and consists of seven Gaussian bumps. **b** The same point process as in **a** was used to generate another set of 50 spike trains. **c** The jitter (*i. e.* the width of the Gaussian bumps in gray) was reduced by a factor of four to generate a different set of spike trains. The average coincidence rate [Kistler et al., 1997] and the average Victor-Purpura metric [Victor and Purpura, 1996] between each spike train in **a** and **b** and between **b** and **c** are indicated in the gap between the subplots. Notice that the coincidence rate is greater between the dissimilar processes in **b** and **c** than between different instantiations of the same point process in **b** and **a**. Similarly, the larger averaged Victor-Purpura distance between **a** and **b** indicates that the process in **b** is more dissimilar from the process in **a** than from process in **c**. In **d**, the one-dimensional example is illustrated. The average point-wise distance between the crosses and the circles is larger if crosses and circles are drawn from the same distribution than if the circles are all clustered at the center.

bias-correction to previously used pair-wise distance measures applicable to small sets of spike trains.

The paper is organized as follows. In order to establish the theoretical links between the different similarity measures, we review the binless vector-space framework ([Carnell and Richardson, 2005]; Sect. 4.2.1 and 4.2.2). A large set of seemingly distinct measures are encompassed by the vector-space framework; from coincidence rate [Kistler et al., 1997] to correlation coefficient of peri-stimulus time histograms [Eggermont et al., 1983, Gawne et al., 1991, Brown et al., 2002, Paninski et al., 2005], and these relationships are reviewed in Sect. 4.2.6 and 4.2.7. In Section 4.2.8 we discuss distance and angle measures between PSTHs. The analytical results on the deterministic bias (Section 4.3.1) and small-sample bias (Section 4.3.2) suggests new methods to compare sets of spike trains (Sections 4.3.3 and 4.3.4). In Section 4.4, we use a set of test cases to show the validity of our analytical results, and to show that these results generalize to other similarity measures that can not be cast in the vector-space framework [Eggermont et al., 1983, Victor and Purpura, 1996, van Rossum, 2001, Schreiber et al., 2003]. We then show that in classification tasks the improved similarity measures can increase the recovered information (Sect. 4.4.7). Finally we demonstrate that the improvements discussed in this article are crucial if the similarity measures are used for fitting or validating a stochastic neuron model (Section 4.4.6 and 4.5).

4.2 The Space of Spike Trains

4.2.1 Spike Train Algebra

Similar to [Carnell and Richardson, 2005] we define the spike train vector as a weighted sum of Dirac-delta-pulse:

$$\mathbf{S}_i = \sum_{m=1}^{n_i} w_i^{(m)} \delta(t - \hat{t}_i^{(m)}) \quad (4.1)$$

with $0 < \hat{t} < T$ where \hat{t} corresponds to a spike time, T to the total length of the spike train and n_i to the number of spikes in spike train i . This mathematical object makes sense only when it has been integrated. For instance the integral of such a spike train would be a mark-accumulator process Snyder et al. [1991]. A spike train will have all weights $w_i^{(m)}$ equal to one. The spike train vector is a mathematical generalisation since it permits spikes to have fractional weights $w_i^{(m)}$. The mathematical framework also enables operations such as addition and multiplication by a scalar on the spike-train vector. In general we can write addition as:

$$\mathbf{S}_i + \mathbf{S}_j = \sum_{m=1}^{n_i} w_i^{(m)} \delta(t - \hat{t}_i^{(m)}) + \sum_{p=1}^{n_j} w_j^{(p)} \delta(t - \hat{t}_j^{(p)}). \quad (4.2)$$

In the case where the spike $\hat{t}_m^{(i)}$ in spike-train i is at the same time as a spike $\hat{t}_p^{(j)}$ in spike-train j then the sum of the spike train vectors has a spike with a weight $w_m + w_p$ at this time. We

can also consider a spike train vector multiplied by a scalar:

$$a\mathbf{S}_i = \sum_{m=1}^{n_i} aw_i^{(m)} \delta(t - \hat{t}_i^{(m)}), \quad (4.3)$$

where $a \in \mathbb{R}$. Once addition and multiplication by a scalar are defined, it is easy to see that the closure axioms, the axioms for addition (commutativity, associativity, existence of zero element, existence of the negative), and the axioms for multiplication by a scalar (associativity, distributivity and existence of identity) of a linear vector space are satisfied. Such a linear vector space made of functions of time is also called a function space. Fig. 4.2a and b illustrate some of the properties of spike trains seen as vectors.

4.2.2 Scalar Products of Spike Trains

In order to define distances and angles between vectors, the function space must have an inner product. A linear space is said to have an inner (or scalar) product if for each vector pair \mathbf{S}_i and \mathbf{S}_j there exists a unique real number $\langle \mathbf{S}_i, \mathbf{S}_j \rangle$ satisfying the following axioms: commutativity, distributivity, associativity and positivity. Typically there are multiple candidates of inner products satisfying the above axioms. The inner products we consider here all have the general form:

$$\langle \mathbf{S}_i, \mathbf{S}_j \rangle = \int_0^T \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} K_{\Delta}(s, s') S_i(t-s) S_j(t-s') ds ds' dt, \quad (4.4)$$

where K_{Δ} is a two-dimensional coincidence kernel with a scaling parameter Δ . K_{Δ} is required to be a non-negative function with a global maximum at the origin. Moreover, $K_{\Delta}(s, s')$ should fall off rapidly so that $K_{\Delta}(s, s') \approx 0$ for all $s, s' > \Delta$. Typical examples of kernels include $K_{\Delta}(s, s') = h_1(s)h_2(s')$. For instance, $h_1(s) = h_2(s) = \frac{1}{\Delta} e^{-s/\Delta}$ is the kernel used in [van Rossum, 2001]. The scaling parameter should be small, $\Delta \ll T$. The limit $\Delta \rightarrow 0$ under the constraint of a fixed integral (e.g. $\int h(s) = 6$ ms), the kernel approaches a Dirac-delta function. With $K_{\Delta}(s, s') = \delta(s)\delta(s')$, we observe that $\langle \mathbf{S}_i, \mathbf{S}_i \rangle = \int_0^T S_i(t) dt = n_i$ where n_i is the number of spikes in \mathbf{S}_i . This inner product between spike trains is part of the reproducing kernel Hilbert space (RKHS) framework [Paiva et al., 2009a] and can be seen as a specific type of spike-pattern matching (SPM) kernel. We will see in Sect. 4.2.6 how several standard similarity measures relate to this class of inner product.

For a comparison of spike-trains seen as vectors the notions of angular separation, distance and norm of spike trains are particularly important. The squared norm of a spike train will be written $\|\mathbf{S}_i\|^2 = \langle \mathbf{S}_i, \mathbf{S}_i \rangle$. Then the distance, D , between two spike-train vectors is:

$$D^2 = \|\mathbf{S}_i - \mathbf{S}_j\|^2 = \langle \mathbf{S}_i - \mathbf{S}_j, \mathbf{S}_i - \mathbf{S}_j \rangle^2 = \|\mathbf{S}_i\|^2 + \|\mathbf{S}_j\|^2 - 2\langle \mathbf{S}_i, \mathbf{S}_j \rangle. \quad (4.5)$$

The cosine of the angle between \mathbf{S}_i and \mathbf{S}_j is:

$$\cos\theta_{ij} = \frac{\langle \mathbf{S}_j, \mathbf{S}_i \rangle}{\|\mathbf{S}_i\| \|\mathbf{S}_j\|}. \quad (4.6)$$

Fig. 4.2e and f illustrate the concepts of angle and distance for spike trains seen as vectors.

4.2.3 Ensemble of Stochastic Spike Trains

We consider spike trains arising from a stochastic process. Examples are theoretical processes (Poisson or other point processes) as well as data from neuronal recordings. From a mathematical point of view, we are interested in expected values of spike trains and functions thereof. From a practical point of view, we need to ask how well these mathematical properties can be estimated from a finite set of spike trains (*e.g.* derived from experimental recordings). We define the firing intensity as the expectation over multiple measurements of a (potentially time-dependent) spiking process generating spikes with unit weight ($w = 1$):

$$v(t) = E[S(t)]. \quad (4.7)$$

Note that $\mathbf{v} \equiv v(t)$ belongs to the same function space as the spike trains \mathbf{S} (Fig. 4.2c and d). Distance, norm and angle between different firing intensities have, therefore, a straightforward definition. It is also possible to take expectations of inner products in the form of Eq. 4.4. Using Eq. 4.4 we can calculate the expected correlation of independent spike trains drawn from the same process (across several independent trials). Because of independence:

$$E[\langle \mathbf{S}_i, \mathbf{S}_j \rangle_{i \neq j}] = \langle E[\mathbf{S}_i], E[\mathbf{S}_j] \rangle = \langle \mathbf{v}, \mathbf{v} \rangle = \|\mathbf{v}\|^2. \quad (4.8)$$

Therefore the expectation of the inner product of two independent spike trains belonging to the same probabilistic process is the squared norm of the firing intensity.

4.2.4 The PSTH as an Empirical Estimate

Let us now switch from the mathematical to the practical point of view. We consider a probabilistic time-dependent point process x from which we draw N_X independent and identically distributed spike trains which we label $\mathbf{S}_i^{(x)}$ (where i runs from 1 to N_X). This process has a firing intensity $v_x(t)$ defined as a mathematical object in Eq. 4.7. In case of an inhomogeneous Poisson process, the rate $v_x(t)$ would completely define the point process, but we keep the treatment general and allow the point process to have arbitrary history dependence. Thus our treatment includes advanced spiking neuron models such as GLMs [Truccolo et al., 2005, Pillow et al., 2008] as well as experimental recordings.

The finite set of N_X spike trains is denoted by the capital letter X . We suppose that the sample of N_X spike trains is the result of a measurement and that this sample contains all the

information available about process x . It is natural to look at the averaged vector:

$$\hat{\mathbf{v}}_X = \frac{1}{N_X} \sum_{i=1}^{N_X} \mathbf{S}_i^{(x)}. \quad (4.9)$$

This resembles the equation for the center-of-mass in classical mechanics, but we will refer to $\hat{\mathbf{v}}$ as the population activity [Gerstner, 2000, Gerstner and Kistler, 2002] since multiple repetitions of the same stimulus can be seen as multiple identical, unconnected neurons receiving the same stimulus. Other names include pooled train [Lindner, 2006] and superposition of point processes [Daley and Vere-Jones, 1988] although this last denomination often assumes sparsity of the individual spike trains which we do not require here. The population activity becomes identical to a Peri-Stimulus Time Histogram (PSTH; [Gerstein and Kiang, 1960]) if we were to replace \mathbf{S} by the vector made of the time-series of binned spike counts. This procedure gives a discrete histogram. Alternatively, one could convolve the spike trains with a Gaussian before or after averaging in order to get a smooth PSTH, also called spike density. In a stochastic framework, we can see $\hat{\mathbf{v}}_X$ as an estimator of the firing intensity since its expected value is \mathbf{v}_x .

In our approach for comparing ensembles of spike trains, the quantity $\hat{\mathbf{v}}_X$ has a central importance. We will derive similarity measures that quantify the match between $\hat{\mathbf{v}}$ of different ensembles. Moreover, we will discuss a quantity that relates to the variance around \mathbf{v}_x seen as a central vector. Such ideas may sound deceptively similar to the ‘central spike-train’ discussed in the literature of spike-train metrics. A central spike train, or prototype spike train, is a spike train that minimizes the sum of the squared distances to all the spike trains of the ensemble [Schoenberg and Tranbarger, 2008]. The quantity $\hat{\mathbf{v}}_X$ is not a prototype spike train, it is actually not a spike train at all but rather the rescaled sum of multiple spike trains. Our approach is built on the theory of point processes where the firing intensity is of central importance, regardless of the process being Poisson. Our approach assumes that the spiking process is such that a time-dependent firing intensity is informative.

4.2.5 Intrinsic Variability and Reliability are Measures of the Variance

In order to link the notion of intrinsic reliability with the vector space picture, we first define \hat{L}_X as the average of the norms of each spike train within the sample X :

$$\hat{L}_X = \frac{1}{N_X} \sum_{i=1}^{N_X} \|\mathbf{S}_i^{(x)}\|^2. \quad (4.10)$$

We note that \hat{L}_X is related to the averaged spike count. L_X is exactly the averaged spike count if the inner product satisfies i) $\int \int K_\Delta(s, s') ds ds' = 1$ and ii) $K_\Delta(s, s') = 0$ whenever either s or s' is greater than the minimum interspike interval of any of the spike trains considered. The interpretation $L_X \sim$ spike count is helpful for the discussion in the remainder of this section.

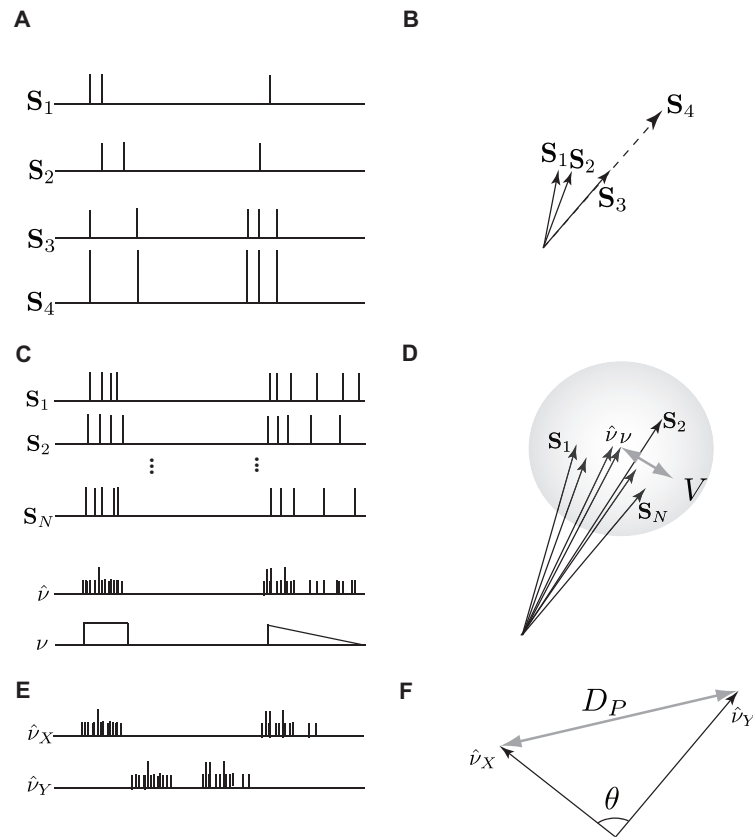


Figure 4.2: **A:** Four spike trains with the spike train S_4 having the same spike times as S_3 but with twice the weight for each spike. **B:** The abstract representation of the four spike trains in **A** as vectors. Vector S_3 is slightly longer than S_1 and S_2 since the length is related to the spike count. Vector S_4 is twice as long because each spike in the spike train (timing-wise equivalent to S_3) has a weight of two. **C:** The population activity \hat{v} , or PSTH, is the sum of all spikes divided by the total number N of spike trains. The population activity \hat{v} is an estimator of the instantaneous firing rate $\nu(t)$. **D:** The abstract representation shows the vector ν as slightly different from its estimator \hat{v} . The population activity vector lies in the center of the spike trains S_i seen as vectors. The mean squared distance from the population activity is the intrinsic variability V . **E:** Example of two population activity vectors \hat{v}_X and \hat{v}_Y that show no overlap in spike timing. **F:** The angle between two vectors of population activity is $\pi/2$ when no spikes are coincident at resolution Δ , such as in the examples of **E**.

Chapter 4. Improved Similarity Measures for Small Sets of Spike Trains

Next, we consider the averaged squared deviation from the population activity, or PSTH:

$$\hat{V}_X = \frac{1}{N_X - 1} \sum_{i=1}^{N_X} \|\mathbf{s}_i^{(x)} - \hat{\mathbf{v}}_X\|^2. \quad (4.11)$$

A similar measure was used by [Croner et al., 1993] to quantify the intrinsic variability in the responses to repeated representations of the same stimulus. We use the $N_X - 1$ normalization such that the sample variance is an unbiased estimator of the variance: $E[\hat{V}_X] = V_x$ for all N_X . Using Eq.'s 4.9 - 4.11, it is easy to show that the variance is proportional to the difference between the squared norm of the population activity and the averaged squared norms within the sample:

$$\hat{V}_X = \frac{N_X}{N_X - 1} (\hat{L}_X - \|\hat{\mathbf{v}}_X\|^2). \quad (4.12)$$

where $\hat{\mathbf{v}}_X$ is the empirical population activity defined in Eq. 4.9. The sample variance \hat{V}_X measures the amount of variability across the set of spike trains. To show this, we insert the definition of $\hat{\mathbf{v}}$ and separate from the double sum the elements containing the norm of the spike trains, which yields:

$$\hat{V}_X = \hat{L}_X - \frac{2}{N_X(N_X - 1)} \sum_i \sum_{j < i} \langle \mathbf{s}_i^{(x)}, \mathbf{s}_j^{(x)} \rangle. \quad (4.13)$$

Under the assumptions indicated after Eq. 4.10, we can state that the sample variance \hat{V}_X is the averaged spike count minus the average number of spikes that are coincident from one repetition to the next. One can thus think of \hat{V}_X as the average number of spikes that are not coincident from one repetition to the next. We refer to \hat{V}_X as the intrinsic variability.

If we take the expected value of the formula above, we can get an expression of the variance, using Eq.'s 4.10 and 4.13:

$$V_x = E \left[\langle \mathbf{s}_i^{(x)}, \mathbf{s}_i^{(x)} \rangle \right]_i - E \left[\langle \mathbf{s}_i^{(x)}, \mathbf{s}_j^{(x)} \rangle \right]_{i,j|i \neq j}. \quad (4.14)$$

We have used indices to the expected value in order to specify that the expectation is done in the first case on the ensemble of spike trains i and in the second term the expectation is on the same ensemble but without the terms corresponding to the inner product of a spike train to itself. Eq. 4.14 will be important for the results of Sect. 4.3.2.

Since the variance quantifies the amount of variability within the set of spike trains, it may be convenient to have a number ranging between zero and one (as for instance in [Schreiber et al., 2003]) to quantify the reliability. We define \hat{R}_X , the estimator of the intrinsic reliability, as:

$$\hat{R}_X = 1 - \frac{\hat{V}_X}{\hat{L}_X} = \frac{\frac{2}{(N_X - 1)N_X} \sum_i \sum_{j < i} \langle \mathbf{s}_i^{(x)}, \mathbf{s}_j^{(x)} \rangle}{\frac{1}{N_X} \sum_{i=1}^{N_X} \|\mathbf{s}_i^{(x)}\|^2}. \quad (4.15)$$

We thus conclude that the intrinsic variability (Eq. 4.13), the intrinsic reliability (Eq. 4.15) and the variance of the spike trains around the population activity (or PSTH; Eq. 4.11) all measure the same quantity but with different normalizations.

4.2.6 Similarity Measures Derived From an Inner Product

[Schrauwen and Campenhout, 2007] and [Paiva et al., 2009a] have discussed the link between spike train vectors and standard measures such as those by [Victor and Purpura, 1996, van Rossum, 2001, Hunter and Milton, 2003, Schreiber et al., 2003]. We briefly review these measures in order to link the formalism introduced in the previous section with known similarity measures. In the list of standard measures we add the measure by [Kistler et al., 1997] and the cross-correlation of spike densities [MacPherson and Aldridge, 1979, David and Gallant, 2005, Petersen et al., 2008]. For brevity, we omit the measure of [Houghton, 2009] for which the link with vector spaces has already been discussed in [Paiva et al., 2010]. We also omit the event synchronization of [Quiroga et al., 2002b] which is closely related to the coincidence factor of [Kistler et al., 1997]. Finally, we only consider spike-time measures, since inter-spike interval, gamma-phase-based, template-based or feature-based similarity measures pertain to different statistical quantities than the instantaneous firing intensity [Victor and Purpura, 1996, Quiroga et al., 2002b, Tiesinga, 2004, Christen et al., 2006, Chi et al., 2007, Kreuz et al., 2007, 2009, Druckmann et al., 2007, Naud et al., 2008]².

Table 4.1 shows the relationship between the space of spike trains and the coefficient of Correlation between PSTHs [Eggermont et al., 1983, Gawne et al., 1991, Paninski et al., 2005], van Rossum's Distance [van Rossum, 2001] and the coincidence factor [Kistler et al., 1997]. The different measures correspond to different kernels K_Δ that define the inner product Eq. 4.4. h_g , h_e and h_r are Gaussian, exponential and rectangular coincidence kernel, respectively;

$$h_g(s; \Delta) = \frac{1}{\sqrt{2\pi\Delta^2}} \exp\left(-\frac{s^2}{2\Delta^2}\right), \quad (4.16)$$

$$h_e(s; \Delta) = \Theta(s) e^{-s/\Delta}, \quad (4.17)$$

$$h_r(s; \Delta) = \Theta(s + \Delta) \Theta(\Delta - s). \quad (4.18)$$

where $\Theta(s)$ is the Heaviside function, with $\Theta(s) = 1$ for $s > 0$ and zero otherwise.

We have included in Table 4.1 the generic angle and distance measures that we introduce in Sect. 4.3. We also have included in Table 4.1 the measures loosely related to an inner product for completeness: Victor and Purpura metric [Victor and Purpura, 1996], the coincidence factor without replacement [Gerstner and Naud, 2009] and the Hunter and Milton measure [Hunter and Milton, 2003]. The quantities $C(\mathbf{S}_i, \mathbf{S}_j)$, $N_{coinc}(\mathbf{S}_i, \mathbf{S}_j)$ and $\mathbf{u}(\mathbf{S}_i \rightarrow \mathbf{S}_j)$ will be defined in the next section.

²Chapter 2

Chapter 4. Improved Similarity Measures for Small Sets of Spike Trains

Name	Original Measure	$K_{\Delta}(s, s')$	References
Coefficient of Correlation	$\rho_{XY} = \frac{\langle \hat{\mathbf{v}}_X, \hat{\mathbf{v}}_Y \rangle}{\ \hat{\mathbf{v}}_X\ \ \hat{\mathbf{v}}_Y\ }$	$h_g(s)h_g(s')$	[Eggermont et al., 1983] [Gawne et al., 1991] [Paninski et al., 2005] [David and Gallant, 2005] [Petersen et al., 2008]
Van Rossum's Distance	$D_{VR}^2 = \ \mathbf{S}_i - \mathbf{S}_j\ ^2$	$h_e(s)h_e(s')$	[van Rossum, 2001]
Coincidence Factor	$CF = \frac{\langle \mathbf{S}_i, \mathbf{S}_j \rangle - 2n_i n_j \Delta / T}{1/2(n_i + n_j)(1 - 2n_i \Delta / T)}$	$\delta(s)h_r(s')$	[Kistler et al., 1997]
Victor-Purpura Spike Metric	$D_{spk} = n_i + n_j - 2C(\mathbf{S}_i, \mathbf{S}_j)$	-	[Victor and Purpura, 1996]
Scaled Victor-Purpura	$VP = \frac{2C(\mathbf{S}_i, \mathbf{S}_j)}{n_i + n_j}$	-	[Kreiman et al., 2000]
Coincidence Factor without replacement	$CF2 = \frac{N_{coinc}(\mathbf{S}_i, \mathbf{S}_j) - 2n_i n_j \Delta / T}{1/2(n_i + n_j)(1 - 2n_i \Delta / T)}$	-	[Gerstner and Naud, 2009]
Hunter-Milton Measure	$HM = \frac{1}{n_i} \sum_k h_e(u_k(\mathbf{S}_i \rightarrow \mathbf{S}_j))$	-	[Hunter and Milton, 2003]
Angular Separation	$M_a = \frac{\langle \hat{\mathbf{v}}_X, \hat{\mathbf{v}}_Y \rangle}{\ \hat{\mathbf{v}}_X\ \ \hat{\mathbf{v}}_Y\ }$	unspecified	Section 4.2.8 [Eggermont et al., 1983] [Gawne et al., 1991] [Paninski et al., 2005] [David and Gallant, 2005] [Petersen et al., 2008]
Distance between population activities	$D_P^2 = \ \hat{\mathbf{v}}_X - \hat{\mathbf{v}}_Y\ ^2$	unspecified	Section 4.2.8
Scaled Distance	$M_D = \frac{2\langle \hat{\mathbf{v}}_X, \hat{\mathbf{v}}_Y \rangle}{\ \hat{\mathbf{v}}_X\ ^2 + \ \hat{\mathbf{v}}_Y\ ^2}$	unspecified	Section 4.2.8

Table 4.1: Summary of spike train similarity measures. Here $h_g(s; \Delta) = \frac{1}{\sqrt{2\pi\Delta^2}} \exp\left(-\frac{s^2}{2\Delta^2}\right)$ defines a Gaussian filter while $h_e(s; \Delta) = \Theta(s)\Theta(s/\Delta)$ is an exponential and $h_r(s; \Delta) = \Theta(s + \Delta)\Theta(\Delta - s)$ is a rectangular one. The kernel K_{Δ} defines the scalar product \langle, \rangle as in Eq. 4.4.

4.2.7 Similarity Measures loosely Related to an Inner Product

Coincidence Factor Without Replacement

The coincidence factor CF may count two coincidences if there are two spikes in \mathbf{S}_j that fall within $\pm\Delta$ of a spike in \mathbf{S}_i . An alternative method to count the number of coincident spikes is to find the maximum number of spikes in \mathbf{S}_i coincident with a spike in \mathbf{S}_j under the constraint that a spike in \mathbf{S}_i can be coincident with at maximum one spike in \mathbf{S}_j (this method was used in the INCF quantitative spike-time prediction in 2009 [Gerstner and Naud, 2009]). In analogy to common probability models we call this procedure a coincidence measure *without replacement*. This counting method can lead to a slightly different number of coincident spikes which we call N_{coinc} . The *coincidence factor without replacement* (CF2)³ is given by replacing $\langle \mathbf{S}_i, \mathbf{S}_j \rangle$ by $N_{coinc}(\mathbf{S}_i, \mathbf{S}_j)$ in the equation of CF (Table 4.1). This similarity measure cannot be seen as an inner product on spike-train vectors because it does not satisfy all the axioms for an inner product enumerated in Sect. 4.2.2 (distributivity and associativity are not fulfilled). However, if the minimum interspike interval is larger than 2Δ , then the difference between $\langle \mathbf{S}_i, \mathbf{S}_j \rangle$ and $N_{coinc}(\mathbf{S}_i, \mathbf{S}_j)$ vanishes since a single spike of \mathbf{S}_i can never coincide with two spikes of \mathbf{S}_j .

Victor and Purpura Metric

[Victor and Purpura, 1996] developed three similarity metrics which measure the minimum cost to transform a spike train $S_i(t)$ into another spike train $S_j(t)$. One of the three metrics, $D_{\text{spk}}(S_i, S_j)$, has been used extensively to study the importance of spike timing in real neurons (see [Victor, 2005] for a review). To calculate D_{spk} one uses a set of three transformations, each having an associated cost: i) shift a spike with a cost q per unit time, ii) remove a spike with a cost of one, iii) add a spike with a cost of one. Note that if a spike is shifted by an interval larger than $\Delta = 2/q$ in order to achieve a match, it is always better to remove the spike and recreate it at the new location. Hence the shifting cost increases linearly until a value of 2 and stays constant thereafter. Defined this way, $D_{\text{spk}}(\mathbf{S}_i, \mathbf{S}_j)$ is greater or equal to zero and cannot exceed the total number of spikes in \mathbf{S}_i and \mathbf{S}_j together [Victor and Purpura, 1996]. Similar to the situation discussed for CF2, D_{spk} cannot be mapped to an inner product on spike train vectors, but the distinction arises only if the minimum interspike interval is shorter than 2Δ .

If the minimum interspike interval is larger than 2Δ , there can be at most one spike in i within $\pm\Delta$ of a spike in j . Then we don't have to worry about the matching between the spike trains. In this case we can count the number of coincidences with an inner product $\langle \mathbf{S}_i, \mathbf{S}_j \rangle$ so that

³Also described in Chapter 3.

the Victor-Purpura metric is:

$$D_{\text{spk}}^2(\mathbf{S}_i, \mathbf{S}_j) = \|\mathbf{S}_i\|^2 + \|\mathbf{S}_j\|^2 - 2\langle \mathbf{S}_i, \mathbf{S}_j \rangle, \quad (4.19)$$

$$\mathcal{K}_\Delta(s, s') = h_t(s; \Delta)\delta(s'), \quad (4.20)$$

$$h_t(s; \Delta) = (1 - |t|/\Delta)\Theta(1 - |t|/\Delta), \quad (4.21)$$

where h_t is a triangular function. This function quantifies the advantage of shifting a spike by s instead of deleting and adding it at a different time. Here $\Delta = 2/q$ corresponds to the amount of shifting for which it is as costly to remove and add a spike. Eq. 4.19 tells us that instead of cumulating the cost to make spike train i identical to spike train j , we may start from the maximum cost $n_i + n_j = \|\mathbf{S}_i\|^2 + \|\mathbf{S}_j\|^2$ and reduce this according to the inner product defined by Eq.'s 4.20-4.21.

In the general case, the minimal interspike interval may be smaller than 2Δ . Instead of the scalar product $\langle \mathbf{S}_i, \mathbf{S}_j \rangle$ we work directly with the quantity $C(\mathbf{S}_i, \mathbf{S}_j)$ that we derive from the Victor-Purpura metric by the definition

$$C(\mathbf{S}_i, \mathbf{S}_j) = \frac{1}{2} [n_i + n_j - D_{\text{spk}}(\mathbf{S}_i, \mathbf{S}_j)] \quad (4.22)$$

where D_{spk} is evaluated as in [Victor and Purpura, 1996]. Finally, in order to compare D_{spk} with the same criteria and on the same scale as the previous similarity measures, we create a measure equivalent to D_{spk} but ranging from zero (worst match) to one (best match) by normalizing with the total number of spikes [Kreiman et al., 2000]:

$$\text{VP}(\mathbf{S}_i, \mathbf{S}_j) = \frac{n_i + n_j - D_{\text{spk}}(\mathbf{S}_i, \mathbf{S}_j)}{n_i + n_j} = \frac{C(\mathbf{S}_i, \mathbf{S}_j)}{\frac{1}{2}(n_i + n_j)}. \quad (4.23)$$

Hunter and Milton

Let $\mathbf{u}(\mathbf{S}_i \rightarrow \mathbf{S}_j)$ be a vector with elements, $u_k(\mathbf{S}_i \rightarrow \mathbf{S}_j) > 0$, made of the times between spike k in the spike train i and the nearest spike in spike train j (where $1 < k < n_i$). The average of each n_i components of \mathbf{u} could be taken as a measure of how spike train i is different from spike train j . Instead of simply averaging the components of $\mathbf{u}(\mathbf{S}_i \rightarrow \mathbf{S}_j)$, one could also consider averages of nonlinear transformations such as h_e . This generalization leads to what is known as the Hunter and Milton measure [Hunter and Milton, 2003]:

$$\text{HM}(\mathbf{S}_i, \mathbf{S}_j) = \frac{1}{n_i} \sum_{k=0}^{n_i} h_e(u_k(\mathbf{S}_i \rightarrow \mathbf{S}_j); \Delta). \quad (4.24)$$

[Chi and Margoliash, 2001] have used the same measure but without the nonlinear function h_e . Here again associativity and distributivity are not fulfilled and, formally, we cannot see HM as an inner product on the function space as defined in Sect. 4.2. HM can be loosely related to

an inner product the same way as VP was loosely related to an inner product.⁴ Moreover, HM can be defined as an inner product in the RKHS [Paiva et al., 2009a].

4.2.8 Comparing Sets of Spike Trains: The Distance between two PSTHs

When comparing sets of spike trains, we would like to know whether set X is distinct from set Y . More precisely, whether the spike trains in the set X were generated by a probabilistic process distinct from the one which generates the spike trains in set Y . For each set, we calculate the population activity vector Eq. 4.9. From the vector space perspective, we define two criteria to distinguish set X from set Y based on their population activity only: the angular separation between the two population activity vectors of each set, and the distance between the two population activity vectors. The angular separation between population activities is equivalent to the coefficient of correlation between PSTHs. Surprisingly, we will see that these criteria can be related to criteria matching the spike timings (angular separation) only or the spike times and the total number of spikes (distance).

Angular Separation Measures Temporal Coding Characteristics

Let the set X consist of N_X i.i.d. spike trains : $\{\mathbf{S}_1^{(x)}, \mathbf{S}_2^{(x)}, \dots, \mathbf{S}_{N_X}^{(x)}\}$ with population activity $\hat{\mathbf{v}}_X$, and the set Y of N_Y i.i.d. spike trains labeled $\mathbf{S}_j^{(y)}$ with population activity $\hat{\mathbf{v}}_Y$. The first similarity measure is the angular separation between the population activities:

$$M_a(X, Y) = \cos \hat{\theta}(X, Y) = \frac{\langle \hat{\mathbf{v}}_X, \hat{\mathbf{v}}_Y \rangle}{\|\hat{\mathbf{v}}_X\| \|\hat{\mathbf{v}}_Y\|}. \quad (4.25)$$

This quantity was extensively used as a similarity measure between populations of spike trains and it is called the correlation coefficient, or Pearson coefficient between two PSTHs [Eggermont et al., 1983, Gawne et al., 1991, Petersen et al., 2008, David and Gallant, 2005]. The cosine goes from zero (perpendicular population activities) to one (parallel population activities and hence perfect match). The spike-train vectors – and consequently the population activities – are always positive. The M_a measure (M stands for *match* and the index a for *angular*) is dependent on the time-course of the firing intensity only, but not its amplitude. In particular, it is independent of the total firing rate. To see this, consider $v_x(t) = A_1 f(t)$, and $v_y(t) = A_2 f(t)$, with scalar parameter $A_1 > 0$ and $A_2 > 0$, for which we have $\frac{\langle \mathbf{v}_x, \mathbf{v}_y \rangle}{\|\mathbf{v}_x\| \|\mathbf{v}_y\|}$ independent of the respective amplitudes A_1 and A_2 .

⁴To illustrate this one as to replace h_e by h_t in Eq. 4.24 and consider spike trains having no interspike intervals smaller than 2Δ . The the measure reduces to $\text{HM} = \frac{1}{n_i} \iiint h_t(s; \Delta) \delta(s') S_i(t-s) S_j(t-s') ds ds' dt$ which is the same inner product as for the VP distance (Eq. 4.20).

Distance Measures both Temporal Coding and Spike Count Characteristics

The second similarity measure is the squared distance between the population activities:

$$D_P^2(X, Y) = \|\hat{\mathbf{v}}_X - \hat{\mathbf{v}}_Y\|^2 = \|\hat{\mathbf{v}}_X\|^2 + \|\hat{\mathbf{v}}_Y\|^2 - 2\langle \hat{\mathbf{v}}_X, \hat{\mathbf{v}}_Y \rangle. \quad (4.26)$$

Here, clearly, both the shape and the amplitude of the firing intensity have to match for the distance to be zero. For the sake of comparison, let us scale and offset this measure so that (similar to the angular match M_a in Eq. 4.25) zero indicates the worst match and one the best match. The distance-based match M_D is defined as:

$$M_D(X, Y) = 1 - \frac{D_P(X, Y)}{D_{\max}} = \frac{2\langle \hat{\mathbf{v}}_X, \hat{\mathbf{v}}_Y \rangle}{\|\hat{\mathbf{v}}_X\|^2 + \|\hat{\mathbf{v}}_Y\|^2}, \quad (4.27)$$

where $D_{\max} = \|\hat{\mathbf{v}}_X\|^2 + \|\hat{\mathbf{v}}_Y\|^2$ is the maximum distance which is obtained for population activities with zero overlap.

4.3 Analytical Results

We start by considering pairwise comparisons of spike trains and show that there is a systematic bias towards noise-free models (Sect. 4.3.1). We then turn to a comparison of the means, *i.e.* the PSTHs. We show in Section 4.3.2 that for small sample sizes, the distance between two PSTHs is biased, but that this bias can be removed.

4.3.1 Train by Train Comparison has a Deterministic Bias

We start with the example from Fig. 4.1, *i.e.* a direct comparison of spike trains. To do so, we average the pairwise distance or similarity by comparing a spike train from set X to one of set Y (as was done with the coincidence factor in [Jolivet et al., 2008a]). We will demonstrate that such a pairwise comparison implies a big bias towards the deterministic process. To illustrate this aspect, let us consider the distance D_{VR} (see Table 4.2) between spike train vectors. In view of the one dimensional example (Fig. 4.1D), it is not surprising that this method wrongly measures the distance between the distributions of spike trains, since we are trying to estimate distance between distributions from element by element distances. First we note that the expected value of the element-by-element average of the distance is equal to the expected value of the distance between two arbitrary spike trains, one chosen from the ensemble X and the other from the ensemble Y :

$$E \left[\frac{1}{N_X N_Y} \sum_{i=1}^{N_X} \sum_{j=1}^{N_Y} \|\mathbf{s}_i^{(x)} - \mathbf{s}_j^{(y)}\|^2 \right] = E [\|\mathbf{S}^{(x)} - \mathbf{S}^{(y)}\|^2]. \quad (4.28)$$

Using the results of the previous section (in particular Eq. 4.7, 4.8 and 4.14) we can derive an expression of the expected value as a function of the underlying firing intensities. This gives:

$$E \left[\frac{1}{N_X N_Y} \sum_{i=1}^{N_X} \sum_{j=1}^{N_Y} \|\mathbf{s}_i^{(x)} - \mathbf{s}_j^{(y)}\|^2 \right] = \|\mathbf{v}_x - \mathbf{v}_y\|^2 + V_x + V_y. \quad (4.29)$$

In other words, the expected value of the element-by-element average of distances between spike trains is not the distance between the two firing intensities. The bias is proportional to the sum of the variability intrinsic to x and y . Moreover we note that the bias does not decrease as we increase the number of spike trains in the sample since it is independent of N_X and N_Y .

Now we come to an important deduction from Eq. 4.29. When the similarity measure is used for fitting a neuron model to experimental data, the variance terms $V_X + V_Y$ in Eq. 4.29 leads to what can be called a ‘deterministic bias’. Suppose that the distance in Eq. 4.29 is to be minimized by tuning the model, say y , to best fit the data, say x . Eq. 4.29 tells us that the best model may not have its firing characteristics equal to those of the neuron ($\|\mathbf{v}_x - \mathbf{v}_y\|^2 = 0$ and $V_x = V_y$) because it can be more advantageous to set the variance V_y of the model to zero at the expense of a minor mismatch $\|\mathbf{v}_x - \mathbf{v}_y\|^2 > 0$. Therefore the best match will be achieved with a process of less variability than the data, that is, a more deterministic model. Again, we stress the fact that this deterministic bias does not arise from the spike train but only from the fact that we are trying to assess the distance between probability distributions by looking at the averaged distance between its elements. A deterministic bias can arise when comparing a model with data or two data sets, and such a deterministic bias is not N -dependent. We explore this in detail with simulations and fits to real data in Sections 4.4 and 4.5.

4.3.2 Sample Bias in the Norm of the Population Activity

The results from Sect. 4.3.1 in particular Eq. 4.29 suggests that we should avoid averaged pairwise comparisons. It might be much better to compare directly the squared distances $\|\hat{\mathbf{v}}_X - \hat{\mathbf{v}}_Y\|^2$ of the mean rate of the set of spike trains. Let us again recall the one-dimensional example of the introduction. A set of $N = 10$ one dimensional data points $\{x_1, \dots, x_{10}\}$ will have a sample mean $\bar{x} = \frac{1}{N} \sum_{i=1}^{10} x_i$. It is well known that even if the sample mean is an unbiased estimator of the mean, the sample mean squared would be a biased estimator of the squared mean:

$$E[\bar{x}^2] - E[x]^2 = E[(\bar{x} - E[x])^2] = \frac{\sigma_x^2}{N}, \quad (4.30)$$

where σ_x^2 is the variance of the sample: $\sigma_x^2 = E[(x - E[x])^2]$. Only when the sample is large ($N \rightarrow \infty$) the bias between $E[\bar{x}^2]$ and $E[\bar{x}]^2$ vanishes. We refer to this effect in the following as a small-sample bias or short "sample bias".

The above one-dimensional example shows that the small-sample bias will appear in com-

Chapter 4. Improved Similarity Measures for Small Sets of Spike Trains

comparisons of spike trains whenever the similarity measure contains the norm of the (empirical) population activity. In the following, we expand on this simple instance of the small-sample bias and apply it to the space of spike trains.

We have seen (Eq. 4.8) that the expected value of the inner product $\langle \mathbf{S}_i^{(x)}, \mathbf{S}_j^{(x)} \rangle$ between two independent spike trains generated by the same point process is $\|\mathbf{v}_x\|^2$. The norm of the population activity $\|\hat{\mathbf{v}}_X\|^2$ is made principally of inner products between independent spike trains (off-diagonal terms), but contains also inner products of a spike train with itself (diagonal terms). To see this, we square the definition of the population activity in Eq. 4.9:

$$\|\hat{\mathbf{v}}_X\|^2 = \left\langle \frac{1}{N_X} \sum_i \mathbf{S}_i, \frac{1}{N_X} \sum_j \mathbf{S}_j \right\rangle = \frac{2}{N_X^2} \sum_i \sum_{j>i} \langle \mathbf{S}_i, \mathbf{S}_j \rangle + \frac{1}{N_X^2} \sum_i \langle \mathbf{S}_i, \mathbf{S}_i \rangle. \quad (4.31)$$

One expects that these diagonal terms (the last sum in Eq. 4.31) will make a vanishingly small contribution when $N \rightarrow \infty$, but when N is small they are not negligible. Therefore, the norm of the population activity is biased if it is used as an estimator of the true $\|\mathbf{v}_X\|^2$, and this bias propagates to all the similarity measures discussed so far (M_a , D_P , M_D). Taking the expectation of Eq. 4.31, we can now see from a comparison with Eq. 4.14 that using $\|\hat{\mathbf{v}}_X\|^2$ as the estimator for $\|\mathbf{v}_X\|^2$ is biased by:

$$E[\|\hat{\mathbf{v}}_X\|^2] - \|\mathbf{v}_X\|^2 = \frac{V_x}{N_X}, \quad (4.32)$$

in complete analogy to Eq. 4.30. But V_x can be estimated to order $(1/N_X)^2$ from \hat{V}_X as we can see from Eq.'s 4.13 and 4.14. Using \hat{V}_X and N_X , we can therefore build an unbiased estimator simply by subtracting the bias from the biased estimator. Let us define the off-diagonal terms:

$$\hat{C}_{XX}^* = \frac{2}{(N_X - 1)N_X} \sum_i \sum_{j<i} \langle \mathbf{S}_i^{(x)}, \mathbf{S}_j^{(x)} \rangle. \quad (4.33)$$

It is clear from Eq. 4.8 that $E[\hat{C}_{XX}^*] = \|\mathbf{v}_X\|^2$. Hence \hat{C}_{XX}^* is an unbiased estimator of the norm of the population activity. The C here stands for 'coincidences' as \hat{C}_{XX}^* counts the coincidences across repetitions, that is across several spike trains from the ensemble X . Across multiple tests with sample size N this new estimator has a variance $\text{Var}[\hat{C}_{XX}^*] = \frac{2}{N(N-1)} \text{Var}[\langle \mathbf{S}_i, \mathbf{S}_j \rangle_{i \neq j}]$ and is a member of a general class of unbiased estimators with minimum variance [Voinov and Nikulin, 1996]⁵. Eq. 4.33 is reminiscent of the equation of intrinsic reliability R_X , indeed we can write:

$$\hat{C}_{XX}^* = \hat{R}_X \hat{L}_X. \quad (4.34)$$

Informally, this means that number of coincidences across repetitions \hat{C}_{XX}^* equals the intrinsic

⁵We can see why the coefficient $\frac{2}{(N_X-1)N_X}$ replaces the coefficient $\frac{2}{N_X^2}$ in Eq. 4.31. The factor $\frac{2}{N_X^2}$ overestimates the number of terms in the corresponding sum by an amount which will become vanishingly small at large N_X : $\frac{2}{N_X^2} = \frac{2}{N_X(N_X-1)} \left(1 - \frac{1}{N_X}\right)$.

reliability times the number of spikes. In terms of the intrinsic variability estimator \hat{V}_X , the number of coincidences across repetitions is the total number of spikes minus the intrinsic variability: $\hat{C}_{XX}^* = \hat{L}_X - \hat{V}_X$.

4.3.3 Sample Bias in Distance and Angle Measures

The sample bias causes an over estimation of the distance measure D_P (defined in Eq. 4.26) proportional to the intrinsic variability V_x in each process. Again the analogy to the set of 1-dimensional data points discussed in the introduction is useful. The squared distance between the sample means $\bar{x} = \frac{1}{N} \sum_{i=1}^N x_i$ and $\bar{y} = \frac{1}{N} \sum_{i=1}^N y_i$ is given by: $(\bar{x} - \bar{y})^2 = \bar{x}^2 + \bar{y}^2 - 2\bar{x}\bar{y}$, which contains the biased quantities \bar{x}^2 and \bar{y}^2 (Sect. 4.3.2). In the case of spike trains, the bias-corrected estimator C_{XX}^* can be used to form an unbiased estimator of the distance between two population activities:

$$D_P^{*2}(X, Y) = \|\hat{\mathbf{v}}_X - \hat{\mathbf{v}}_Y\|^2 - \frac{\hat{V}_X}{N_X} - \frac{\hat{V}_Y}{N_Y} = \hat{C}_{XX}^* + \hat{C}_{YY}^* - 2\langle \hat{\mathbf{v}}_X, \hat{\mathbf{v}}_Y \rangle, \quad (4.35)$$

such that $E[D_P^{*2}(X, Y)] = \|\mathbf{v}_x - \mathbf{v}_y\|^2$. The reader may want to compare the right hand side of Eq. 4.26 with the right hand side of Eq. 4.35. Note that the expectation $E[D_P^*(X, Y)]$ can never fall below zero if X and Y are statistically independent point processes. Nevertheless, for small sample size $D_P^*(X, Y) < 0$ is possible so that D_P^* is not a metric⁶.

Because the sample bias causes an overestimate of $\|\hat{\mathbf{v}}_X\|^2$, this leads to an underestimation of the angular separation M_a and M_D . Replacing $\|\hat{\mathbf{v}}_X\|^2$ by \hat{C}_{XX}^* should remove the sample bias present in M_a and M_D . We therefore consider instead of Eq. 4.25 a bias-corrected angular separation:

$$M_a^*(X, Y) = \frac{\langle \hat{\mathbf{v}}_X, \hat{\mathbf{v}}_Y \rangle}{\sqrt{\hat{C}_{XX}^* \hat{C}_{YY}^*}}. \quad (4.36)$$

Similarly the distance measure of Eq. 4.27 is replaced by:

$$M_D^*(X, Y) = \frac{2\langle \hat{\mathbf{v}}_X, \hat{\mathbf{v}}_Y \rangle}{\hat{C}_{XX}^* + \hat{C}_{YY}^*}. \quad (4.37)$$

Note that when comparing a set of model data X with experimental data Y , the instantaneous firing intensity of the model is often known exactly. In this case the normalization should be done with $C_{XX}^* = \|\mathbf{v}_{model}\|^2$.

However, even if \hat{C}_{XX}^* is an unbiased estimator of $\|\mathbf{v}_x\|^2$ it does not imply that $(\hat{C}_{XX}^*)^{-1}$ is an unbiased estimator of $\|\mathbf{v}_x\|^{-2}$ [Voinov, 1985]. We expect that this additional bias will be

⁶ $D_P^*(X, Y) < 0$ will arise when the number of coincidences across repetitions \hat{C}_{XX}^* is smaller than the coincidence count between X and Y , $\hat{C}_{XX}^* + \hat{C}_{YY}^* < 2\langle \hat{\mathbf{v}}_X, \hat{\mathbf{v}}_Y \rangle$. For instance, two sets of spike trains each made of only two repetition of a homogeneous Poisson process could happen to have no coincidences across distinct repetitions ($\hat{C}_{XX}^* = \hat{C}_{YY}^* = 0$) but still coincidences between sets X and Y ($\langle \hat{\mathbf{v}}_X, \hat{\mathbf{v}}_Y \rangle \neq 0$).

sizable only when \hat{C}_{XX}^* is close to zero on the scale of its standard deviation. Before using the bias-corrected measures M_a^* and M_D^* described above, one should check if the magnitude of \hat{C}^* is very close to zero. We have observed that this can happen when there are very few spikes or when the point process has low intrinsic reliability, for instance a homogenous gamma process with random initiation and small firing rate. In this case D_p^* should be preferred to M_a^* or M_D^* . If an angular separation is required then one can use the method employed by [David and Gallant, 2005] to estimate the first order bias of $\|\hat{\mathbf{v}}_X\|^{-2}$.

The need to correct for the sample bias has been highlighted in the appendix of David and Gallant [2005]⁷ and of Petersen et al. [2008]⁸. In these studies they arrive at empirical estimates of the bias which accurately approximate the bias-corrected angular separation of Eq. 4.36. In contrast to these earlier approaches, our analytical de-biasing method presented here has minimal variance.

4.3.4 Applications to Measures Loosely Related to an Inner Product

In Sect. 4.3.1 we have seen that pairwise comparisons contain deterministic bias. Moreover, in Sect. 4.3.2 we have seen that working with squared quantities can give rise to a bias for small sample sizes. With these insights in mind, we now return to classical spike train metrics. The derivations we presented in Sect. 4.3.2-4.3.1 rely heavily on the inner product and the quantity $\hat{\mathbf{v}}$, *i. e.* the population activity defined as the empirical “mean” of the set of spike trains. We could go one step further and try to conciliate our approach with spike-train metrics. This is motivated by our point-process approach and strays from the formalisms of spike train metrics (see discussion in Sect. 4.6) The similarity measures obtained above encompass all the spike-train metrics that are related to an inner product, but we use the spike-train metrics to subserve a comparison of the population activity. In spike-train metrics loosely related to an inner product, there is no such thing as a population activity. In this section we exploit the analogy with the inner product to compute a distance that should relate to the distance between the distributions rather than the distance between the individual spike trains.

To do this, we have to identify a quantity that is related to $\|\hat{\mathbf{v}}\|^2$ and replace this quantity by a term analogous to \hat{C}_{XX}^* in Eq. 4.33. In this section we explain how this can be done for the Victor and Purpura metric (D_{spk}), for the coincidence rate without replacement (CF2) and for the Hunter and Milton measure (HM). The resulting measures are summarized in Table 4.2.

⁷David and Gallant [2005] remove the bias by first noting that $\|\hat{\mathbf{v}}\|^{-2} = \|\mathbf{v}\|^{-2} + \epsilon/N$, where ϵ scales the magnitude of the noise. By successively calculating $\|\hat{\mathbf{v}}\|^{-2}$ for subsets of the total sample of different sizes and plotting as a function of $1/N$, $\|\mathbf{v}\|^{-2}$ can be estimated by linear extrapolation to find the offset at $1/N = 0$. This method yields a good first-order bias correction that also removes bias due to the inverse.

⁸Using our present notation, Petersen et al. [2008] estimate $\|\mathbf{v}\|^2$ by taking the inner product of separate halves of the total available samples. This is a clever and simple way to remove the sample bias, though the additional bias due to the inverse may still be important. We note that because of the separation of the total sample in two subsamples this estimator has a variance that is roughly twice greater than $\text{Var}[\hat{C}^*]$, its variance would correspond to $\frac{4}{N^2} \text{Var}[\langle \mathbf{S}_i, \mathbf{S}_j \rangle_{i \neq j}]$.

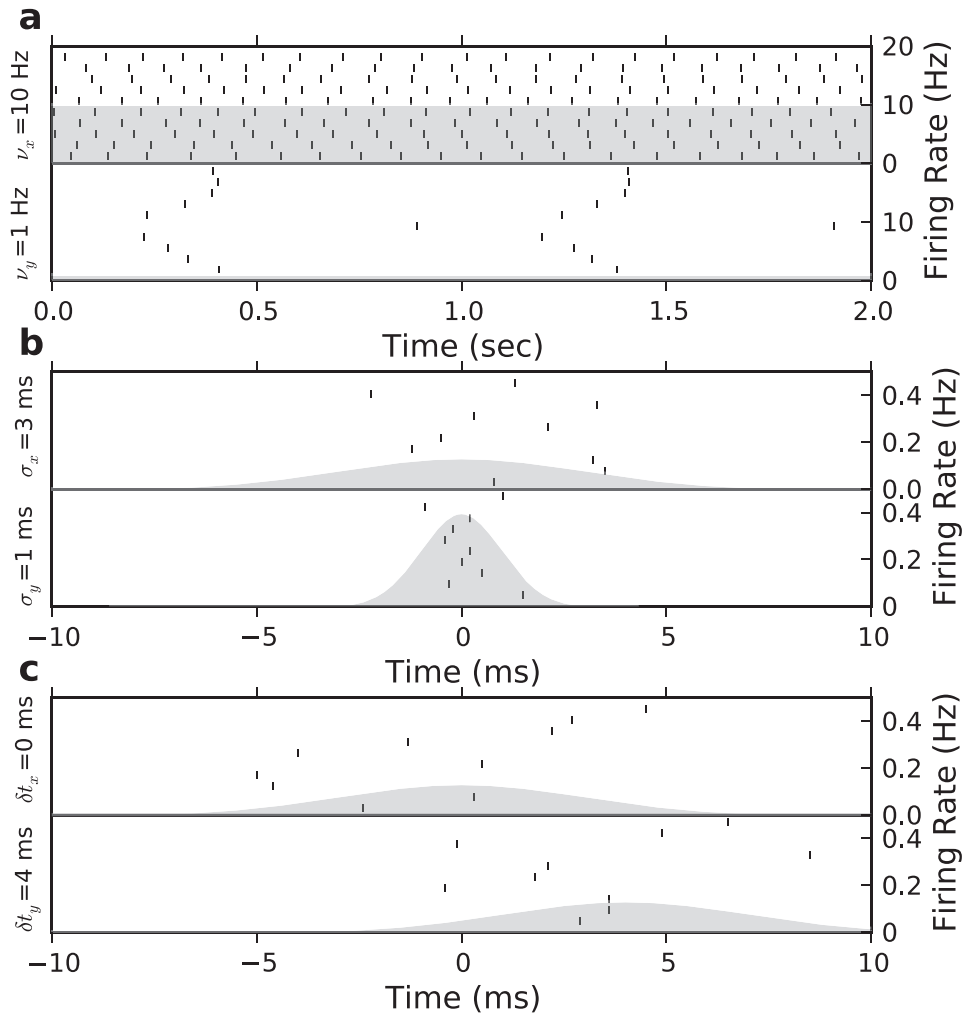


Figure 4.3: Raster plots and firing rates illustrating the different case studies. Two sets of $N=10$ spike trains (examples of X and Y) are shown representing: **a** comparing two gamma-processes, one with $v_x = 10$ Hz (top) the other with $v_y = 1$ Hz (bottom); **b** spike time jitter with $\sigma_x = 3$ ms (top) and $\sigma_y = 1$ ms (bottom); **c** spike time latency with $\delta t_x = 0$ ms (top) and $\delta t_y = 4$ ms (bottom). The shaded area shows the instantaneous firing rate of the process with its axis on the right.

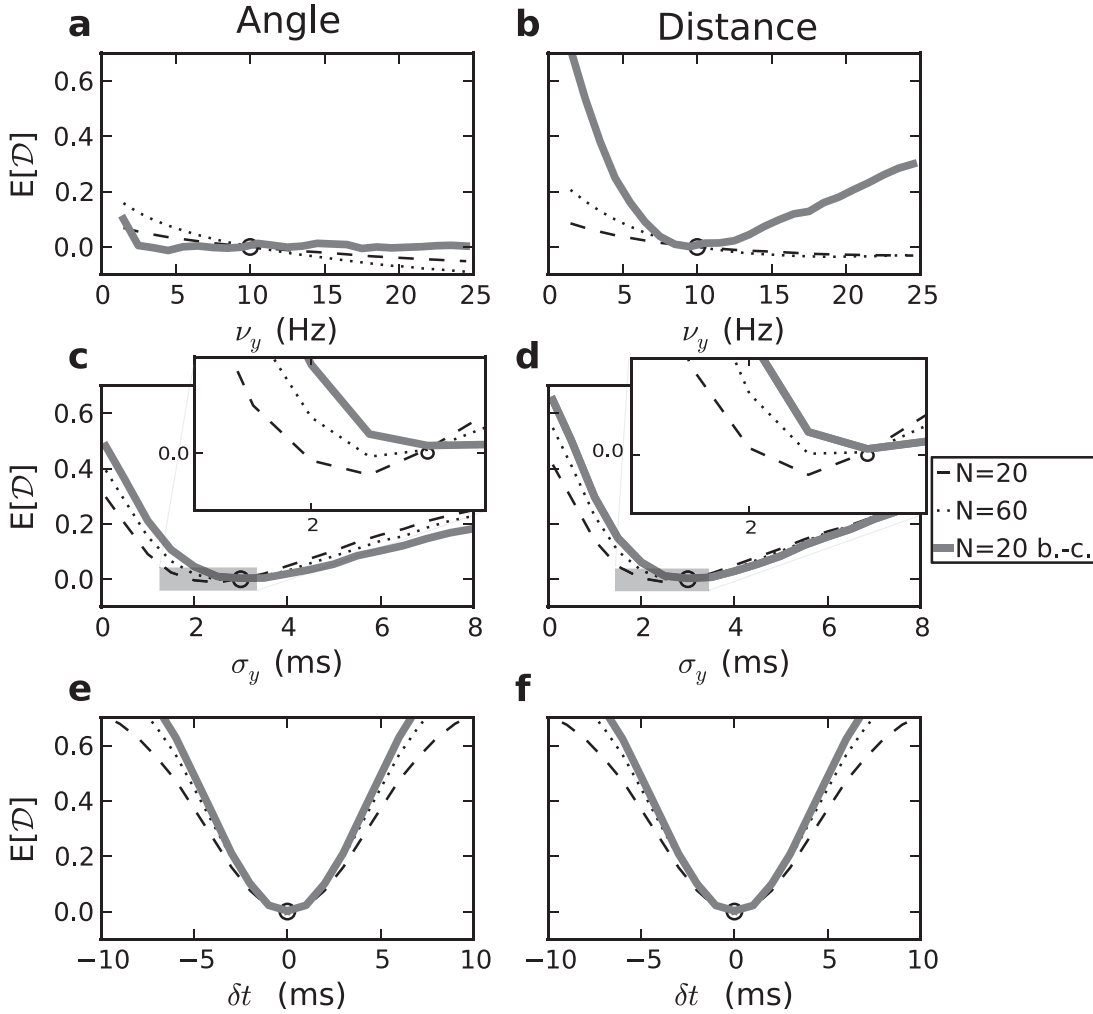


Figure 4.4: Discriminability profiles for three similarity measures of spike trains seen as vectors: $M_a - M_a^*$ (a, c, e) and $M_D - M_D^*$ (b, d, f). The discriminability is plotted for different values of either the firing intensity of a Gamma process (first row; a, b), the spike time jitter (second row; c, d), or the average latency (third row; e, f). The discriminability profiles are shown with N spike trains with dashes ($N=20$), dots ($N=60$) and a thick line for the bias-corrected measures ($N=20$). The reference value of the firing intensity, the jitter, or the latency in ensemble X and X' is indicated with a circle at $\mathcal{D} = 0$: $\nu_x = 10$ Hz in a and b, $\sigma_x = 3$ ms in c and d, $\delta t = 0$ ms in e and f. The zoomed in plot of the shaded area is shown as an inset for c and d.

Name	Group Measure	C_{XX}^*	Section
Angular Separation	$M_a^* = \frac{\langle \hat{\mathbf{v}}_X, \hat{\mathbf{v}}_Y \rangle}{\sqrt{\hat{C}_{XX} \hat{C}_{YY}}}$	$\frac{2}{N(N-1)} \sum_{i,j>i} \langle \mathbf{S}_i, \mathbf{S}_j \rangle$	4.3.3
Distance	$D_P^{*2} = \hat{C}_{XX}^* + \hat{C}_{YY}^* - 2\langle \hat{\mathbf{v}}_X, \hat{\mathbf{v}}_Y \rangle$	$\frac{2}{N(N-1)} \sum_{i,j>i} \langle \mathbf{S}_i, \mathbf{S}_j \rangle$	4.3.3
Scaled Distance	$M_D^* = \frac{2\langle \hat{\mathbf{v}}_X, \hat{\mathbf{v}}_Y \rangle}{\hat{C}_{XX}^* + \hat{C}_{YY}^*}$	$\frac{2}{N(N-1)} \sum_{i,j>i} \langle \mathbf{S}_i, \mathbf{S}_j \rangle$	4.3.3
Victor-Purpura Spike Metric	$D_{spk}^* = \hat{C}_{XX}^* + \hat{C}_{YY}^* - 2\hat{C}_{XY}$	$\frac{2}{N(N-1)} \sum_{i,j>i} C(\mathbf{S}_i, \mathbf{S}_j)$	4.3.4
Scaled Victor-Purpura	$VP = \frac{2\hat{C}_{XY}}{\hat{C}_{XX}^* + \hat{C}_{YY}^*}$	$\frac{2}{N(N-1)} \sum_{i,j>i} C(\mathbf{S}_i, \mathbf{S}_j)$	4.3.4
Coincidence Rate no replacement	$CF2^* = \frac{\hat{C}_{XY}}{\frac{1}{2}(\hat{C}_{XX}^* + \hat{C}_{YY}^*)}$	$\sum_{i,j>i} \frac{N_{\text{coinc}}(\mathbf{S}_i^{(x)}, \mathbf{S}_j^{(y)}) - 2n_i^{(x)}n_j^{(y)}\Delta/T}{N(N-1)/2}$	4.3.4
Hunter-Milton Measure	$HM^* = \frac{\hat{C}_{XY}}{\frac{1}{2}(\hat{C}_{XX}^* + \hat{C}_{YY}^*)}$	$\sum_{i,j>i} \frac{HM(\mathbf{S}_i^{(x)} \rightarrow \mathbf{S}_j^{(y)}) + HM(\mathbf{S}_j^{(y)} \rightarrow \mathbf{S}_i^{(x)})}{N(N-1)}$	4.3.4

Table 4.2: Summary of similarity measures for comparing groups of spike trains.

Victor and Purpura D_{spk}

Let us consider the element-by-element average of the Victor-Purpura metric D_{spk} defined in [Victor and Purpura, 1997]. We now use the definition of $C(\mathbf{S}_i, \mathbf{S}_j)$ in Eq. 4.22 so as to evaluate the average distance between N_X spike trains in set X and N_Y spike trains in set Y we have:

$$\frac{1}{N_X N_Y} \sum_{i=1}^{N_X} \sum_{j=1}^{N_Y} D_{spk}(\mathbf{S}_i^{(x)}, \mathbf{S}_j^{(y)}) = \bar{n}_X + \bar{n}_Y - \frac{2}{N_X N_Y} \sum_{i=1}^{N_X} \sum_{j=1}^{N_Y} C(\mathbf{S}_i^{(x)}, \mathbf{S}_j^{(y)}). \quad (4.38)$$

Since there are N_X spike trains in ensemble X and $n_i^{(x)}$ is the number of spikes in trial i , $\bar{n}_X = \sum_{i=1}^{N_X} n_i^{(x)} / N_X$ is just the average number of spikes in a train from ensemble X . Similarly, \bar{n}_Y is the average number of spikes per train across the spike trains in set Y . $C(\mathbf{S}_i^{(x)}, \mathbf{S}_j^{(y)})$ is a number discounting any coincidence from the maximum cost $n_i + n_j$ of transforming the spike train i into spike train j (see Sect. 4.2.7). In Eq. 4.19 we replaced C by the inner product defined with the triangular function (which holds only when the interspike interval is greater than 2Δ). The analogy is instructive; if C corresponds to the inner product then the average number of spikes in Eq. 4.38 should play a role equivalent to $\|\hat{\mathbf{v}}\|^2$ in Eq. 4.26.

Why and how should we replace the average number of spike per repetition in Eq. 4.38? We have argued that the similarity measure we are looking for should account for the similarity between statistical properties of the point processes. D_{spk} above indicates a distance of zero only when the spike trains (or the set of spike trains) are identical, spike per spike, through all the repetitions. Formulated this way, D_{spk} has no reasons to be zero when we compare spike trains coming from the same stochastic point process. For instance, two sets of spike trains coming from the same inhomogeneous Poisson process will not have all their spikes coinciding. We can create a measure that will have expected distance zero when comparing spike trains from the same point process by replacing the average number of spikes \bar{n}_X or \bar{n}_Y

by the average coincidence count across distinct pairs of repetition within each set, \hat{C}_{XX}^* for set X and \hat{C}_{YY}^* for set Y :

$$\hat{C}_{XX}^* = \frac{2}{N_X(N_X - 1)} \sum_{i=1}^{N_X} \sum_{j>i}^{N_X} C(\mathbf{S}_i^{(x)}, \mathbf{S}_j^{(x)}). \quad (4.39)$$

This leads to the new distance across sets of spike trains:

$$D_{spk}^*(X, Y) = \hat{C}_{XX}^* + \hat{C}_{YY}^* - 2\hat{C}_{XY}, \quad (4.40)$$

where $\hat{C}_{XY} \equiv \frac{1}{N_X} \frac{1}{N_Y} \sum_{i=1}^{N_X} \sum_{j=1}^{N_Y} C(\mathbf{S}_i^{(x)}, \mathbf{S}_j^{(y)})$ is the average coincidence count between individual spike trains of set X and individual spike trains of set Y (compare Eq. 4.40 with Eq. 4.35 to see that $\hat{C}_{XY} \sim \langle \hat{\mathbf{v}}_X, \hat{\mathbf{v}}_Y \rangle$). Then D_{spk}^* is zero when the average coincidence count across X and Y is the same as the average coincidence count across distinct repetitions within set X and within set Y . This distance is maximum when the processes x and y are so dissimilar that there are never any coincidences between spike trains from X and spike trains from Y . Similar to D_P^* treated in Sect. 4.3.2, D_{spk}^* can not be seen as a metric since it is greater or equal to zero only in expectation.

In Eq. 4.23 we performed a normalization of $D_{spk}(\mathbf{S}_i, \mathbf{S}_j) \in [0, n_i + n_j]$ into a measure $VP(\mathbf{S}_i, \mathbf{S}_j) \in [0, 1]$. For a scaled distance between sets of spike trains analogous to the VP distance, we introduce the expression:

$$VP^*(X, Y) = \frac{\hat{C}_{XY}}{\frac{1}{2}(\hat{C}_{XX}^* + \hat{C}_{YY}^*)}. \quad (4.41)$$

This is very similar to M_D^* but in Eq. 4.41 the quantity that we called average coincidence count (\hat{C}_{XX}^* , \hat{C}_{YY}^* , \hat{C}_{XY}) is based on element-by-element comparison rather than on the inner product on the population activity. Through numerical simulations, we will show in Sect. 4.3.4 that VP^* should be used instead of VP to properly distinguish between small sets of spike trains derived from different point processes.

Coincidence Factor without Replacement

In Section 4.2.7, we distinguished between two methods for counting the number of coincidence. In the first one, the number of coincidences can be counted using the inner product for CF in Table 4.1. In the second, we count the number of coincidence by allowing at most one spike from train i to be coincident with a spike in train j . We called this second quantity $N_{\text{coinc}}(\mathbf{S}_i^{(x)}, \mathbf{S}_j^{(y)})$. The analogy with the inner product is clear with $\langle \mathbf{S}_i, \mathbf{S}_j \rangle \sim N_{\text{coinc}}(\mathbf{S}_i, \mathbf{S}_j)$. Here, we keep the discounting term for the chance level (as in the equation for CF in Table 4.1) and identify:

$$\hat{C}_{XY} = \frac{1}{N_X} \frac{1}{N_Y} \sum_{i=1}^{N_X} \sum_{j=1}^{N_Y} \left(N_{\text{coinc}}(\mathbf{S}_i^{(x)}, \mathbf{S}_j^{(y)}) - 2n_i^{(x)} n_j^{(y)} \Delta / T \right) \quad (4.42)$$

as the averaged number of coincidences that is loosely related to an inner product. \hat{C}_{XY} counts the number of coincidences *without replacement* above the expected number of coincidences that one would obtain if both spike trains were produced by a Poisson process ($2n_i n_j \Delta / T$). Accordingly, we can define the average number of coincidences across distinct repetitions:

$$\hat{C}_{XX}^* = \frac{1}{N_X(N_X - 1)} \sum_{i=1}^{N_X} \sum_{j \neq i}^{N_X} \left(N_{\text{coinc}}(\mathbf{S}_i^{(x)}, \mathbf{S}_j^{(x)}) - 2n_i n_j \Delta / T \right). \quad (4.43)$$

Here again, the expectation of \hat{C}_{XY} will be always smaller than the expectation of \hat{C}_{XX}^* or \hat{C}_{YY}^* . We then formulate a new version of CF2 (distance-based) as:

$$\text{CF2}^* = \frac{\hat{C}_{XY}}{\frac{1}{2}(\hat{C}_{XX}^* + \hat{C}_{YY}^*)}. \quad (4.44)$$

We will show in the examples of Sect. 4.3.4 that Eq. 4.44 removes the deterministic bias from CF2.

Hunter and Milton

It is also possible to create a bias-corrected measure equivalent to HM by considering the average coincidence count:

$$\hat{C}_{XY} = \frac{1}{N_X} \frac{1}{N_Y} \sum_{i=1}^{N_X} \sum_{j=1}^{N_Y} \frac{\text{HM}(\mathbf{S}_i^{(x)} \rightarrow \mathbf{S}_j^{(y)}) + \text{HM}(\mathbf{S}_j^{(y)} \rightarrow \mathbf{S}_i^{(x)})}{2}. \quad (4.45)$$

From which naturally follows the estimate of its upper bound:

$$\hat{C}_{XX}^* = \frac{1}{N_X(N_X - 1)} \sum_{i=1}^{N_X} \sum_{j=i+1}^{N_X} \text{HM}(\mathbf{S}_i^{(x)} \rightarrow \mathbf{S}_j^{(x)}) + \text{HM}(\mathbf{S}_j^{(x)} \rightarrow \mathbf{S}_i^{(x)}). \quad (4.46)$$

Which yields a measure similar to the distance-based M_D :

$$\text{HM}^* = \frac{\hat{C}_{XY}}{\frac{1}{2}(\hat{C}_{XX}^* + \hat{C}_{YY}^*)}. \quad (4.47)$$

Through case studies, we will show in Sect. 4.3.4 that VP^* , CF2^* , HM^* correctly distinguish differences between point processes.

4.4 Examples

In this section, we will show that the sample bias can cause spike-train measures to be unsuitable to distinguish between point-processes despite the fact that the firing rate or temporal characteristics of the two point processes are different by a factor of two or three.




Figure 4.5: Discriminability profiles for three standard similarity measures: the Hunter and Milton measure (HM; a,d,g), the coincidence factor without replacement (CF2; b,e,h) and the scaled Victor and Purpura metric (VP; c,f,i). The discriminability is plotted (dash line) for different values of either the firing intensity of a Gamma process (first row; a, b, c), the spike time jitter (second row; d, e, f), or the average latency (third row; g, h, i). The discriminability profiles for the analogous measures HM*, CF2* and VP* are shown for all cases with a full line. The target value of the firing intensity, the jitter or the latency is indicated with a dot line with a circle at $\mathcal{D} = 0$. The area where a similarity measure shows a negative discriminability is indicated in gray.

4.4.1 A Framework for Comparing The Different Measures

Suppose we have a set of spike trains X that were generated with process x – for instance, N_X spike trains could come from the measurement of N_X neuronal responses to repeated presentation of a fixed (but potentially time-dependent) stimulus. Furthermore, we have a second set of spike trains Y generated with the process y . This second process could have been generated by a model. In order to measure whether the model process y is similar to the neuronal process x , we would like a similarity function between zero and one, $M(X, Y)$, to measure the match between X and Y . The letter M stands for *match* (or *similarity*).

In statistical terms we want to evaluate if we can reject the hypothesis that x is the same as y using the restricted data contained in X and Y . To do this we consider the null hypothesis that Y is no more different to X than X is to another instantiation of x that we call X' . Let X' be a another set of N_Y spike trains taken from process x . Our confidence in *rejecting* the hypothesis that x is the same as y grows with the probability of observing $M(X, X') > M(X, Y)$. With the test statistics:

$$\mathcal{D} = M(X, X') - M(X, Y), \quad (4.48)$$

this probability is:

$$p = \int_0^{\infty} p(\mathcal{D}) d\mathcal{D} \quad (4.49)$$

where $p(\mathcal{D})$ is the probability density of \mathcal{D} and the integral runs over positive values $\mathcal{D} > 0$.

In addition to p , another useful quantity is the mean of \mathcal{D} , defined as $\overline{\mathcal{D}} = E[\mathcal{D}] = \int_{-1}^1 \mathcal{D} p(\mathcal{D}) d\mathcal{D}$. One would expect that whenever $E[M(X, Y)] < E[M(X, X')]$, $\overline{\mathcal{D}} > 0$ indicating that the match between process x and process y is worse than the match between process x and itself. Suppose that we construct an artificial data set for which we know *a priori* that x is different from y . Thus we know that $\overline{\mathcal{D}}$ should be positive. If we find nevertheless a value $\overline{\mathcal{D}} < 0$, the only remaining possibility is that the measure of match M on which we have based our statistical variable $\overline{\mathcal{D}}$ is ill-suited for assessing the similarity between collections of spike trains. We will show in this section that some measures have $\overline{\mathcal{D}} < 0$ even if process x is distinct from y .

In the following we study the hypothesis that a model y (for which we can choose the parameters) is statistically different from neuronal data x . In Section 4.5 the neuronal data comes from recordings, but here our ‘neuronal data’ is artificially generated. For a given data set X , the self-match $M(X, X')$ is fixed and close to one while the match $M(X, Y)$ depends on the choice of parameters of the model process y . Hence maximizing the match between model and data is the same as minimizing the statistical discriminability $\overline{\mathcal{D}}$, see Eq. 4.48. Our aim is to shed light on the flaws of existing measures by showing that the mean discriminability $\overline{\mathcal{D}}$ can go negative – which should never occur for sensible measures of match.

4.4.2 Numerical Methods

In the following sections we will compute $\overline{\mathcal{D}}$ in different scenarios to investigate whether the metric has a bias indicated by $\overline{\mathcal{D}} < 0$. We calculate $\overline{\mathcal{D}}$ numerically for the following measures of match: angle-based M_a , bias-corrected M_a^* , distance-based M_D , bias-corrected distance M_D^* , re-scaled Victor-Purpura VP, Coincidence factor without replacement CF2 and Hunter and Milton HM as well as our novel de-biased measures VP^* , HM^* and $CF2^*$. We generate N spike trains of the process y (data set Y), and $2N$ spike trains of the process x (data sets X and X'). To calculate the expected value we repeated the simulation of the $3N$ spike trains between 1000 to 10000 times to obtain smooth discriminability profiles.

Artificial data sets were generated with different processes. To simulate inhomogeneous Poisson processes and the Spike Response Model (characterized by the firing intensity $v(t)$ or the conditional intensity $\lambda(t)$) we discretized time with bins of $\Delta t = 0.1$ ms, and compared the probability of spiking in that bin ($P_F = 1 - e^{-v(t)\Delta t}$ or $P_F = 1 - e^{-\lambda(t)\Delta t}$) to a random number of uniform distribution. All programs for the simulations were written in Matlab. For the VP measure, we used the Matlab code by Daniel Reich and Jonathan Victor.⁹

Different measures often correspond to a different kernel shape. To calculate M_a and M_D and their bias-corrected equivalents we used $K_\Delta(s, s') = h_t(s)h_t(s')$ with the triangular kernel $h_t = (1 - |t|/2\Delta)\Theta(1 - |t|/2\Delta)$ as defined with the VP metric. In order to compare the other measures on the same basis we chose the kernel's parameter (see Sect. 4.2.6 in Eq.'s 4.17, 4.16, 4.18 and 4.21) such that the integral of the kernel was equal to 4 ms: $\int_{-\infty}^{\infty} h_e(s; \Delta) ds = \int_{-\infty}^{\infty} h_t(s; \Delta) ds = \int_{-\infty}^{\infty} h_r(s; \Delta) ds = \int_{-\infty}^{\infty} h_g(s; \Delta) ds = 4$ ms. For example, h_r intrinsic to CF2 has an integral of 2Δ , therefore we took $\Delta = 2$ ms for computing CF2. The integral of the kernel was 4 ms for all simulations and figures shown in this article, except in Sect. 4.4.7 where Δ was varied over a large range.

4.4.3 Case Study I – Match based on PSTH or Population Activity

In order to have full control over all details, we work with data sets X , X' and Y that are generated artificially. Three different point processes are studied.

Homogenous Gamma Process

We investigate how well the multiple measures discriminate between two homogeneous Gamma processes with different firing intensities. This can be seen as probing the sensitivity of a measure to distinguish two different messages in a pure rate code. For example, a neuron might respond to stimulus 'A' with spike trains that can be described as a gamma process of rate v_x . For a stimulus 'B', it is again a gamma process but of rate v_y . Can we distinguish between stimulus 'A' and 'B'?

⁹<http://www-users.med.cornell.edu/jdvicto/metricdf.html>, viewed in June 2010.

We sample from the Gamma interval distribution of order 2. The probability density of observing an interspike interval of s is therefore:

$$p(s) = \nu e^{-\nu s} (\nu s). \quad (4.50)$$

Using a total time of $T = 150$ seconds and a first spike randomly put within the first second, we generate $2N$ spike trains with $\nu = \nu_x = 10$ Hz to form X and X' and N spike trains with $\nu = \nu_y$ to form Y . ν_y can vary between 1 and 25 Hz. Sample spike trains from X with $\nu_x = 10$ Hz and for Y with $\nu_y = 1$ Hz are shown in Fig 4.3a. Changing the order of the gamma-process to one (Poisson) does not induce sizable changes in the curves of Fig. 4.4a-b and Fig. 4.5a-c (results not shown).

Fig. 4.4 shows the discriminability profiles for $N = 20, 60$ for the angle-based match (M_a , Eq. 4.25) and the distance-based match (M_D , Eq. 4.27) and the bias corrected measures (M_a^* , Eq. 4.36 and M_D^* , 4.37) with $N = 20$. Biased similarity measures of angular separation show a region where $\overline{\mathcal{D}} < 0$. This bias does not disappear if we go from a small sample of $N = 20$ to a rather large sample of $N = 60$. As discussed in Section 5.1, a value of $\overline{\mathcal{D}} < 0$ indicates that the measure of match is flawed.

When we correct for sample bias we observe an elimination of the negative $\overline{\mathcal{D}}$ areas. Three conclusions can be drawn from Fig. 4.4 a and b. First, a naive measure based on angles or distances is biased, as indicated by negative values of $\overline{\mathcal{D}}$. Second, even a bias-corrected measure based on angles is not able to discriminate between two ensembles of different firing rates. Third, the bias-corrected measure based on distance is highly sensitive to different firing rates.

Let us mention in passing a minor point. We have argued in Sect. 4.3 that the angular separation should be completely insensitive to the modifications in the amplitude of the firing intensity. However, in Fig. 4.4a the bias-corrected angle-based measure shows positive $\overline{\mathcal{D}}$ at $\nu_y = 1$ Hz. This can be explained by recalling that it is not entirely unbiased when the norm of the population activity is close to zero (Section 4.3.2). Indeed, for $\nu_y = 1$ Hz we calculate $\hat{C}_{YY}^* = 0.11$ with a distribution that is clearly not Gaussian. Under these conditions, the inverse of an unbiased estimator is not itself unbiased and this additional bias can explain why M_a^* appears sensitive to the firing rate difference between $\nu_x = 10$ Hz and $\nu_y = 1$ Hz.

Spike Time Jitter

Next we investigate how well the multiple measures discriminate features of a spike-timing code. We can think, for instance, of the first-spike latency code discussed in [Gawne et al., 1996] and in [Gollisch and Meister, 2008]. We consider the arrival of a single spike with a given average latency and at a given precision. We look at the sensitivity of the similarity measures to changes in precision. In this scenario we generate $2N$ spike trains made of a single spike shifted by a Gaussian random number having a standard deviation of $\sigma_x = 3$ ms. These spike

Chapter 4. Improved Similarity Measures for Small Sets of Spike Trains

trains form the sets X and X' . We ask whether we can discriminate the spike trains from N other spike trains made of a single spike, each shifted by a Gaussian random number having a different standard deviation of σ_y which can vary between 0 and 8 ms. Sample spike trains from X with $\sigma_x = 3$ ms and for Y with $\sigma_y = 1$ ms are shown in Fig 4.3b.

We find that for $N = 20$, naive implementations of the angle-based match M_a and the distance-based match M_D show a region where $\overline{\mathcal{D}} < 0$. The size of the region decreases with increasing size of the sets as we can see in the insets of Fig. 4.4 c and d. The region where $\overline{\mathcal{D}} < 0$ occurs for parameters $\sigma_y < \sigma_x$ which indicates the deterministic bias: a spiking process is seen as more similar to a process with a smaller jitter than to a process with the correct jitter. In contrast, the bias-corrected measure M_a^* and M_D^* can discriminate spikes arriving with any two distinct jitters.

Since each spike train consists only of a single spike we expect from the results in Sect. 4.3.2 that M_a^* and M_D^* can have an additional bias due to the inverse of the unbiased estimator. Nonetheless we do not observe a region with negative discriminability for the bias-corrected measures, indicating that in this case the additional bias due to the inverse is negligible. We conclude that for timing-based codes, the raw measures of match M_a (based on angle) and M_D (based on distance) exhibit a deterministic bias, that is, they favour a model interpretation with less jitter than the real process. Both bias-corrected matches M_D^* and M_a^* work well.

Spike Time Latency

To probe the sensitivity to average latency, we generate $2N$ spike trains made of a single spike shifted by δt_x plus a Gaussian random number with standard deviation of $\sigma = 3$ ms. We ask whether we can distinguish these spike trains from N spike trains made of a single spike coming with a different latency δt_y plus a Gaussian random number with standard deviation of $\sigma = 3$ ms. This will produce a firing intensity given by:

$$v(t) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(t - \delta t_y)^2}{2\sigma^2}\right). \quad (4.51)$$

Sample spike trains from X with $\delta t_x = 0$ ms and for Y with $\delta t_y = 4$ ms are shown in Fig 4.3c.

Fig. 4.4e and f show the discriminability profiles for $N = 20$, and 60 for M_a and M_D and the bias corrected measures M_a^* and M_D^* with $N = 20$. We observe that increasing the size of the sample increased the discriminability. We note that the bias-corrected measures with $N = 20$ show better discriminability than the raw measures for a much larger sample size of $N = 60$.

The overall conclusion of Sect. 4.4.3 is that (i) the bias-corrected measures always work better than the biased ones, (ii) the distance-based match M_D^* works on all 3 scenarios whereas the angle-based match M_a^* works only for temporal codes and is insensitive to firing rate codes.

4.4.4 Case Study II – Measures Loosely Related to an Inner Product

Using the same paradigms (gamma process, spike time jitter and spike time latency) we check the quality of those well-known similarity measures that relate only loosely to the vector space framework. In particular we consider the HM measure, the coincidence factor CF2, the Victor-Purpura metric VP (see Sect. 4.2.7), along with HM^* , VP^* and $CF2^*$. We calculate numerically the value of $\overline{\mathcal{D}}$ for each measure and plot the results in Fig. 4.5 using $N = 20$ (changing N here will only affect the standard deviation as explained in Sect. 4.3.1). We see that for a gamma-process (Fig. 4.5 a, b, c) there is an area where $\overline{\mathcal{D}} < 0$ for all three measures but for both VP and HM it is more pronounced than for CF2.

The fact that $\overline{\mathcal{D}} < 0$ for large firing rates indicates that a gamma process which in reality has a low firing rate of 10 Hz will be seen as deceptively similar to a Gamma process at higher firing rate. This can be understood by a greater number of random coincidences at higher firing rates that are not discounted in the similarity measures. On the other hand CF2 appears not to be sensitive to different firing rates. This is not too surprising because CF2 corrects for random coincidences. A comparison with Fig. 4.4a suggests therefore that we can interpret CF2 as an angular separation between spike train vectors, since both angular match and CF2 are insensitive to firing rates.

Modified versions of the three measures (HM^* , $CF2^*$ and VP^*) have $\overline{\mathcal{D}} \geq 0$ for all firing rate ν_y . We see (Fig. 4.5a) that HM^* is insensitive to changes in the overall firing rate, which would make this measure related – conceptually – to an angular separation. $CF2^*$ and VP^* on the other hand have a strong sensitivity to the overall firing rate and so are to be interpreted as distance-based measures.

When investigating the discriminability for spike-time jitter, we find that HM, CF2 and VP exhibit a region of negative discriminability (Fig. 4.5d, e, f). This means that a process generating spikes with a greater precision will be systematically overrated when compared to a process with larger jitter. This is what we called deterministic bias in Sect. 4.3.1. The modification suggested in Sect. 4.3.4 removes this deterministic bias in the new measures HM^* , $CF2^*$ and VP^* which show $\overline{\mathcal{D}} \geq 0$ for all jitter σ_y .

All measures retained a high and positive discriminability for average latency (Fig. 4.5 g, h, i). This indicates that the three measures HM, VP and CF2 are equally useful for the detection of spike latency codes. The discriminability was greater when using HM^* instead of HM, and when using $CF2^*$ instead of CF2, but not when using VP^* instead of VP.

We conclude that HM, CF2 and VP all fail to discriminate the level of precision and sometimes the overall firing rate. The measures HM^* , $CF2^*$ and VP^* have $\overline{\mathcal{D}} \geq 0$ everywhere. We note that VP^* and $CF2^*$ were sensitive to change of firing statistics from all scenarios considered, with $CF2^*$ showing a slightly greater sensitivity.




Figure 4.6: Raster plots and firing rates illustrating the phase coding . Two sets of $N=10$ spike trains (examples of X and Y) are shown for comparing mixtures of random spike times and precisely locked spike times, one with $\alpha_x = 0.5$ (top) the other with $\alpha_y = 0$ (bottom). The shaded area shows the instantaneous firing rate of the process with its axis on the right.

4.4.5 Case Study III - Phase code

We can explore the deterministic bias further by considering an inhomogeneous Poisson process with a tuneable proportion of spikes that arrive at specific times. Let us consider an inhomogeneous Poisson process with a time-varying intensity made of 50 equally spaced Gaussian bumps superposed on a constant baseline. The amplitude of the baseline is chosen to complement to the amplitude of the gaussian such that the expected number of spikes is always 50 for a single trial. The amplitudes are regulated by a parameter α :

$$v(t) = 50 \frac{\alpha}{T} + \frac{1 - \alpha}{\sqrt{2\pi\sigma^2}} \sum_{k=1}^{50} \exp\left(-\frac{(t - t_0 - k\delta t)^2}{2\sigma^2}\right), \quad (4.52)$$

with the total time $T = 5$ s, the inter-bump interval $\delta t = 100$ ms and the latency of the first bump $t_0 = 50$ ms. The parameter α tunes the proportion of spikes that will arise from the Gaussians versus the number of spikes that will arrive at random times (see Fig. 4.6). When $\alpha = 1$, the spike train is Poisson and conversely when $\alpha = 0$ spiking is entirely periodic with a jitter σ of 3 ms. We call α the random fraction. This type of spike-generating process can be related to a phase code riding on top of a random background [Lisman, 2005, Kayser et al., 2009].

The inhomogeneous Poisson process is simulated using Eq. 4.52. We generate $2N$ spike trains with 50% of phase-locked spikes ($\alpha_x = 0.5$). We then compare those spike trains with N spike trains generated with α_y and calculate $\overline{\mathcal{D}}$ numerically.

Fig. 4.7 shows the discriminability profile for M_a^* , M_D^* , HM, CF2 and VP with $N = 10$. The thick solid line gives the mean value of \mathcal{D} averaged over many samples with $N = 10$ each. We can observe again that only the measures M_a^* and M_D^* do not have a region where $\overline{\mathcal{D}} < 0$ when

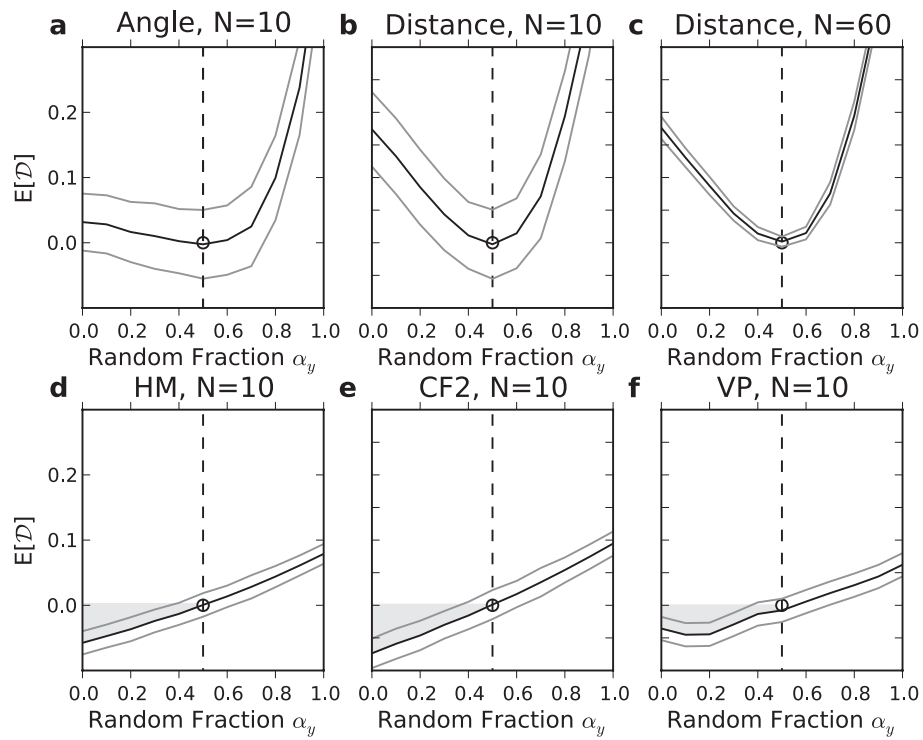


Figure 4.7: Phase coding: Discriminability profiles for: M_a^* (**a**, $N=10$), M_d^* (**b**, $N=10$; **c**, $N=60$), the Hunter and Milton measure (**d**, $N=10$), the coincidence factor without replacement (**e**, $N=10$) and the scaled Victor and Purpura metric, VP, (**f**, $N=10$). The discriminability is plotted (full, black line) for different values of the fraction of spikes that arrive at a fixed time. The target value of the fixed fraction is indicated with a dash line with a circle at $\overline{\mathcal{D}} = 0$. The area where a similarity measure shows a negative discriminability is indicated in gray.

$\alpha_y < \alpha_x$. All the commonly used measures such as VP, HM and CF2 have a strong bias. If the real data has a random fraction α_x and the model α_y , model parameters optimized with VP, HM, or CF2 will give a best value of $\alpha_y = 0$. That is, a smaller random fraction increases the value of the similarity measures HM, CF2, and VP despite the fact that the real process for generating the ensemble X has a stochastic fraction $\alpha_x = 0.5$. This observation confirms a strong deterministic bias that is to be expected in view of our theoretical results in Sect. 4.3.

The standard deviations in Fig. 4.7 are much larger for the measures M_a^* and M_D^* than for HM, CF2, VP. This comes from the estimation of the variance required for correcting the bias and which is less precise at small sample sizes. The question arises whether the large standard deviation is a disadvantage of M_a^* or M_D^* as compared to their biased versions M_a and M_D .

We compared the biased and bias-corrected distance and looked at the variance of \mathcal{D} . Fig. 4.8a shows again that the standard deviation is greater for bias-corrected M_D^* than for biased M_D by superposing the curves on the same plot. Keeping the random fraction fixed at $\alpha_y = 0.2$ and changing the size of the sample used to construct the population activity, we see that the standard deviation decreases rapidly as we increase N (Fig. 4.8c). The point where the lower line of the standard deviation crosses zero is $N = 6$ for the bias-corrected measure and $N = 30$ for the biased measure. Moreover, it takes roughly $N = 400$ samples for the biased measure to achieve the same discriminability as the bias-corrected.

The probability of observing $\mathcal{D} < 0$ is the confidence level of observing \mathcal{D} within the confidence interval of negative discriminability. In Fig. 4.8b, we show that the probability of observing $\mathcal{D} < 0$ as a function of the sample size N is always smaller for the bias-corrected M_a^* .

In summary, the bias-corrected match M_D^* works extremely well for phase-coded input. It is significantly better than the angle-based match M_a^* . The classical measures such as HM, CF2 and VP fail completely.

4.4.6 Case Study IV - Spike Response Model

We evaluate the discriminability of the similarity measures on more realistic data by using the Spike-Response Model with escape noise to generate the spike trains (SRM; [Plesser and Gerstner, 2000b],[Gerstner, 2008]). The SRM considered here is a generalized linear model (GLM) for which the probability of spiking depends on all the spiking history [Gerstner et al., 1996b, Truccolo et al., 2005, Pillow et al., 2005]. We use a simple exponential for both the input filter (having a resistance R and a time constant τ_m : $\kappa(t) = \frac{R}{\tau_m} e^{-t/\tau_m}$) and for the spike-history kernel (having an amplitude A and a time constant τ_A : $\eta(t) = -Ae^{-t/\tau_A}$). The conditional firing intensity for this model is:

$$\lambda(t|I, S_t) = f\left(\frac{[\kappa * I](t) + [\eta * S_t](t)}{\delta V} - \theta_0\right) \quad (4.53)$$




Figure 4.8: The similarity measure M_D^* increases discriminability for all sample sizes greater than one. **a** The discriminability for the scaled distance between population activities with $N=10$ (M_D , dash line) and the same measure corrected for small sample bias (M_D^* , full, thick line). The locked fraction of 0.2 is chosen for plots **b** and **c**. **b** Probability of observing a negative discriminability is smaller for the bias-corrected (M_D^* , full, thick line) than without bias correction (M_D , dash line) at all sample sizes above one. **c** Discriminability is plotted against the sample size with (M_D^* , full, thick line) and without (M_D , dash line) bias correction. The area within one standard deviation of the expected value is shown in pale gray.




Figure 4.9: Illustrating the effect of the model parameters on the PSTH. **a** Shows the spike times of $N = 100$ repeated presentation of the same stimulus to a Spike Response Model with $\delta V = 5$ mV, $\tau_A = 100$ ms and $\tau_m = 10$ ms. The PSTHs (thick lines) in **b**, **c**, and **d** are produced by averaging the number of spikes in each 0.1 ms time bin over $N = 10000$ repetitions, using the same paradigm as in **a**. Using the same time-dependent input, three PSTHs are produced for different values of **b** the scaling factor (δV), **c** the adaptation time constant (τ_A), and **d** the time constant of the input filter (τ_m).




Figure 4.10: The discriminability of the Victor and Purpura measure (dash line) and M_D^* ($N=10$, full, thick line) is shown for varying model parameters: **a** the scaling factor (δV), **b** the adaptation time constant (τ_A), and **c** the time constant of the input filter (τ_m). The target value of the fixed fraction is indicated with a dash line with a circle at $\overline{\mathcal{D}} = 0$. The area where a similarity measure displays a negative discriminability is indicated in gray.

where δV scales the stochasticity of the model (also called the temperature), θ_0 relates to the firing threshold and regulates the offset, f is the link function (here an exponential $f(x) = \lambda_0 e^x$), and S_t is the spike train of all times before t , and I is the stimulus current. We model the stimulus as an Ornstein-Uhlenbeck process (mean: 0 pA, std: 20 pA, correlation time: 3 ms) representing current arriving from a large number of synapses. For all the simulations presented we fixed $A = 2$ mV, $R = 0.1$ M Ω , and $\theta_0 = 2$. We used spike trains of one second to calculate numerically the discriminability $\overline{\mathcal{D}}$ of M_D^* and VP .

Fig. 4.9 shows that changing the parameters δV , τ_m , τ_A has a non-trivial effect on the instantaneous firing intensity. Let us now assume that the ‘real’ neuron has parameters $\delta V = 5$ mV, $\tau_A = 100$ ms and $\tau_m = 10$ ms. We use this neuron model to generate the data set X with $N = 10$ spike trains. We now try to extract the *correct* parameters by simulating the same model but with variable parameters.

Naively, one would expect that we find the correct set of parameters by minimizing a discriminability \mathcal{D} between our simulations Y and reference sets X , X' . Fig. 4.10 shows that this is not the case if we base the distance measure on the VP metric. The bias in the VP measure (and other measures in the literature) gives rise to a pseudo-optimal set of parameters with $\delta V = 2$, $\tau_m = 8$ ms and $\tau_A = 80$ ms. On the other hand, if we use the unbiased distance measure M_D^* we retrieve the correct parameters for δV , τ_A and τ_m . We conclude that the bias in firing rate and spike jitter that we had seen earlier in the VP measure (Fig. 4.5) gives rise to a strong mismatch of parameters.

4.4.7 Case Study V - Classification

Spike-train similarity measures have been used extensively for the classification of spike trains. The examples of previous subsections have shown that 2-category classification requires methods such as M_d^* , VP^* , HM^* or $CF2^*$ to distinguish between processes that differ in their level of stochasticity. Here, we compare classification performance using the classical D_{spk} defined in [Victor and Purpura, 1997] and the new D_{spk}^* defined in Table 4.2.

We follow the leave-one-out classification method of [Victor and Purpura, 1997]. We use the phase coding of Sect. 4.4.5 to conceive 5 different classes of spike trains according to 5 different values of the random fraction ($\alpha = [0, .25, .5, .75, 1]$). We draw 25 spike trains from each of the 5 different classes, which makes a total of 125 spike trains. We will estimate the classification performance by randomly choosing one of the 125 spike trains and estimating with a spike-train similarity measure to which of the 5 classes it should belong. To do this, we measure the average D_{spk} of the singled-out spike train $S_i^{(v)}$ with $S_j^{(\mu)}$ the 25 or 24 spike trains of each class μ :

$$d(S_i^{(v)}, \mu) = \left[\frac{1}{N} \sum_{j=1}^N D_{spk} \left(S_i^{(v)}, S_j^{(\mu)} \right)^z \right]^{1/z} \quad (4.54)$$

where the exponent z is introduced to span various topological situations [Victor and Purpura, 1997]. The singled-out spike train is assigned to the class with the minimum average distance. We then repeat the operation N_{tot} times and build the matrix $Z_{v\mu}$ made of the number of times a spike trains from class v was attributed to class μ . The classification performance gives an estimate of the transmitted information H in bits [Victor and Purpura, 1997]:

$$H = \frac{1}{N_{tot}} \sum_{v,\mu} Z_{v\mu} \left[\log_2 Z_{v\mu} - \log_2 \sum_v Z_{v,\mu} - \log_2 \sum_m Z_{v,m} + \log_2 N_{tot} \right]. \quad (4.55)$$

We repeat the classification 5 times for multiple values of $q = 2/\Delta$, and then repeat the whole procedure again with D_{spk}^* .

Using D_{spk}^* instead of d almost doubles the amount of transmitted information H recovered when q is sufficiently large (Fig. 4.11). The increase in transmitted information is robust to changes in the exponent z (Fig. 4.11, right). At small q the classification is almost random. This can be explained by the fact that for small cost q of moving a spike corresponds to a large time scale on which the metric is sensitive. Therefore, increasing q makes the metric gradually more sensitive to the presence of precisely timed spikes coming with a Gaussian jitter of ± 2 ms. The advantage of using D_{spk}^* instead of D_{spk} starts to be significant at $q = 8$ cost per second and increases further until a maximum is reached at $q = 256$ cost per second. The optimum of D_{spk}^* at $q = 256$ cost per second corresponds to a maximal temporal shift of $\Delta = 4$ ms.

We conclude that the improved similarity measure can increase the information recovered with spike train classification. Since the classification performance depends heavily on the




Figure 4.11: Classification of spike trains coming from different point processes is improved by using D_{spk}^* instead of the measure of Eq. 4.54 d . *Left*: The transmitted information H is shown as a function of the cost of moving one spike q for classifications performed using D_{spk}^* (thick black line, circle markers) or the original Victor-Purpura Distance d with $z = 2$ according to Eq. 4.54 (gray line, square markers). Leave-one-out classification was used on 5 different classes of spike trains. Perfect classification in a scenario with five equally possible classes would have resulted in $H = \log_2(5) = 2.3$ bits shown with the dash line. The maximum likelihood classifier under the assumption that the generative model is known would be able to achieve $H = 2.1$ bits (dot-dashed). Error bars indicate one standard deviation. *Right*: The transmitted information H is shown as a function of the cost of moving one spike q for classifications performed using D_{spk}^* (circle markers with error bars) or the original Victor-Purpura Distance d with various z exponents ($z = -2$ black dotted, $z = 4$ full gray line, $z = -4$ dashed gray line, $z = 8$ full black line, $z = -8$ dashed black line).

statistics of the spike train classes, further investigations is required to assess the extent to which this increase is significant and general.

4.5 Fitting Cortical Neurons

In 2009, the INCF (International Neuroinformatics Coordinating Facility) organized the third competition in modeling single neuron activity. The aim was to provide common benchmark data and evaluation criteria to compare the performance of different neural models [Gerstner and Naud, 2009]. We use a subset of this publicly available data to demonstrate the impact of biased performance measures on the model selection process (see Section 4.3.1).

4.5.1 Experimental details

A L5 pyramidal cell in a brain slice of a 14 days old rat was stimulated with a time-dependent, in-vivo like current. Somatic recordings were made to obtain the voltage trace of the membrane potential. In “current-clamp mode”, a current was injected through the same patch pipette. The current waveform was generated using a superposition of Poissonian excitatory and inhibitory spike trains, convolved with two exponential kernels. In total, there were 13 repetitions of 60 seconds each. Full experimental details can be found elsewhere¹⁰.

The training set consists of the first 38 seconds for which both the injected current and the voltage trace are provided. For the test set, the injected current for the last 22 seconds is provided and the spike times have to be predicted.

To compare the recorded spike trains with the predicted ones, the coincidence factor without replacement (CF2) with a window width of $\Delta = 4$ ms was used (see Section 4.2.7). The question we want to address here is whether the choice of CF2 to quantify model performance could have had a strong influence on the final ranking of the submissions.

4.5.2 Neural Model and Fitting

We modeled the spiking activity of the cell with a SRM having an exponential link function ($f(x) = \lambda_0 e^x$ in Eq. 6.3). As in Section 4.4.6, the instantaneous firing rate $\lambda(t|H_t)$ is conditioned on a filtered version of the injected current $[\kappa * I](t)$, the filtered past activity of the neuron $[\eta * S_t](t)$ and a constant offset θ_0 to model baseline activity (see Equation 6.3).

The spike-history filter η is described as a sum of $m_h = 9$ basis splines (B-splines of order 3). Knot points are spaced on a roughly logarithmic scale up to 320 ms. Let the j th spline be of

¹⁰INCF web site: <http://www.incf.org/Events/competitions/spike-time-prediction/2009>

shape $A_j(\Delta t)$, then the contributing term to Equation 6.3 above is:

$$\frac{[\eta * S_t](t)}{\delta V} = \sum_j^{m_h} \theta_{h,j} [A_j * S](t), \quad (4.56)$$

where $\theta_{h,j}$ are parameters of the spike-history filter. Similarly, the current is filtered with a kernel that is made up of $m_c = 14$ spline basis functions $B_j(\Delta t)$ with knot points roughly logarithmically spaced up to 250 ms. The contribution from the current to the linear sum of Equation 6.3 is thus given by:

$$\frac{[\kappa * I](t)}{\delta V} = \sum_j^{m_c} \theta_{c,j} [B_j * I](t), \quad (4.57)$$

with parameters $\theta_{c,j}$ for the current filter.

Spike times are extracted from the voltage trace by searching for zero-crossings from below. The resulting series of spikes is transformed into a binary sequence with a bin size of 1 ms. The subthreshold fluctuations of the voltage trace are ignored for the fitting procedure because the aim of the competition organized by the INCF was a prediction of spike times rather than the subthreshold membrane potential. All parameters are determined using standard IRLS (iteratively reweighted least squares; [McCullagh and Nelder, 1998]) method to obtain a maximum-likelihood estimate of the 24 parameters in $\vec{\theta}$.

The estimated spike-history filter η captures mainly refractory effects. The kernel used for the current shows an approximately exponential decay with a positive influence of the injected current on spiking probability. Applying the time-rescaling theorem [Brown et al., 2002] suggests a good fit on the training data. Samples of spike trains are generated using the current trace of the test set to generate the model's prediction.

4.5.3 Deterministic Bias Penalizes the Maximum Likelihood Solution

Since the model is stochastic, we generated 13 different realizations of a spike train, conditioned on the same current, *i. e.* the 22 seconds of current given in the test set of the INCF competition. These 13 realizations are then compared with the experimentally observed spike trains. This procedure gives a first value of the CF2 score. The mean CF2 score, averaged over $N = 100$ repetitions, is $CF2^{(ML)} = 0.36$. Here, the index (ML) indicates this solution is found with a stochastic model using an optimization based on maximum-likelihood as an objective. We now vary the stochasticity of the model to see whether the CF2 exhibits a bias towards deterministic models. The level of stochasticity is controlled by the parameter δV in Equation 6.3. The maximum-likelihood solution corresponds to $\delta V = 1$ mV. For $\delta V \rightarrow \infty$, the model tends to assign a constant spiking probability for each bin — a homogeneous Poisson process,




Figure 4.12: Fitting cortical neurons. The model parameters obtained by maximum-likelihood estimation ($\delta V = 1$ mV) were scaled by a temperature-like factor and the performance of these modified models were evaluated using the CF2 (solid line). The vertical line indicates the performance at the maximum-likelihood estimate. For each δV , $N = 100$ predictions were generated with 13 trials each. Points denote mean performance; error bars denote standard deviation. The dashed line shows the bias-corrected distance measure based on the coincidence factor (M_D^* with the kernel K_Δ defined for CF in Table 4.1).

while for $\delta V \rightarrow 0$, the model becomes more deterministic, the transfer function approaches a step function and the neuron fires as soon as the right-hand side of Eq. 6.3 becomes positive. We vary the parameter δV in a wide range and adjust the constant offset θ_0 numerically to keep the firing rate approximately equal to the firing rate observed at $\delta V = 1$ mV.

For each value of δV , we generated 13 realizations from this modified model and calculated the CF2 performance. We repeat this procedure several times so as to find the mean CF2 values as a function of δV (Figure 4.12, solid line). As expected, the performances approach 0 (chance level) for large values of δV since the model turns into a homogeneous Poisson process. However, for values lower than $\delta V = 1$ mV, a clear increase in performance is visible: The maximum-likelihood solution does not yield the highest CF2 score, instead the purely deterministic model ($\delta V \rightarrow 0$) exhibits the peak performance of around 0.45 for this model setup.

We re-evaluated the same spike time predictions using a debiased version of the coincidence factor, M_D^* using the kernel K_Δ defined for CF in Table 4.1 (see Figure 4.12, dashed line). The peak performance in this metric now almost coincides with the maximum-likelihood solution and hence the bias towards a deterministic threshold-model has disappeared. These results also suggest that the maximum-likelihood estimate is optimal in the sense that the model

performance cannot be trivially enhanced by moving away from the maximum-likelihood estimate (choosing $\delta V \neq 1$ mV). We note that changing δV is equivalent to rescaling all model parameters by a constant factor.

Dividing performance measure CF2 by the intrinsic reliability of the neuron (which is a constant for a given neuron, here it was 0.782) as in the original competition will scale CF2 uniformly, independent of δV . This scaling gives an insight into of the performance of the model in comparison to the intrinsic noise in the data but will not remove the deterministic bias.

4.5.4 New Ranking in INCF Competition

The INCF competition accepted both stochastic models and deterministic models. Based on our theoretical results in Sect. 4.3 and the simulations of Sect. 4.4, we expect that the CF2 measure penalized stochastic models, and that re-evaluating the submitted predictions with bias-correction is important to make a fair ranking of how well the models match with the data.

Table 4.3 shows the ranking of all participations in Challenge A in the INCF competition 2009 for both CF2 and the new CF2*. The two measures are therefore equivalent and differ mainly in how we normalize the number of coincident spikes. We hide the name and affiliation of participants but we classify the models in three types. Type I: deterministic models with a single predicted spike train for all repetitions. Type II: deterministic models fitted independently on each repetition such that there is a different predicted spike train each repetition. Type III: stochastic models.

The first striking observation in Table 4.3 is that all deterministic models move several ranks down the list once deterministic bias is removed from the similarity measures. The second observation that catches the eye is that a stochastic model (# 11, type III) that was ranked as one of the three worst models when using CF2 is now in the top three when we use the new measure for model ranking. With CF2*, the top three models are either stochastic (type III) or deterministic models fitted independently on each repetition (type II). Fitting independently on each repetition introduces variability in the predicted spike trains which goes towards the variability intrinsic to real neurons. Some of the stochastic models performed badly with CF2 and still perform badly with CF2*. This indicates that the bias-correction does not simply give an advantage to stochastic models but requires the correct type of stochasticity.

We observe that most of the submissions were made using deterministic models and that this could be explained by the fact that these models had an unfair advantage against the stochastic models, because CF2 was used to quantify the respective performances. It could be said that CF2 does not adequately take into account the stochastic nature of the neurons. Following the theoretical results in Sect. 4.3 and the case studies in Sect. 4.4, we propose that bias-corrected measures based on distance between population activities such as M_D^* or the

Chapter 4. Improved Similarity Measures for Small Sets of Spike Trains

Participant	CF2/R (%)	Participant	CF2* (%)
# 1 (type II)	76.2	# 7 (type II)	74.3
# 2 (type I)	75.4	# 6 (type II)	72.3
# 3 (type I)	73.4	# 11 (type III)	68.7
# 4 (type I)	71.4	# 8 (type II)	68.4
# 5 (type I)	71.0	# 1 (type II)	66.4
# 6 (type II)	68.9	# 2 (type I)	66.1
# 7 (type II)	68.3	# 3 (type I)	64.5
# 8 (type II)	67.0	# 4 (type I)	62.7
# 9 (type I)	63.3	# 5 (type I)	62.7
# 10 (type III)	58.4	# 10 (type III)	59.6
# 11 (type III)	58.3	# 9 (type I)	57.0
# 12 (type III)	56.0	# 12 (type III)	49.6

Table 4.3: Columns on the left: Evaluation of CF2 normalized by the intrinsic reliability of the data, as published on the INCF web page. Right: re-ordering with the bias-corrected measure CF2*. We find that all deterministic models move several ranks down when the corrected CF2* measure is used, *i. e.* after removal of the deterministic bias. Submissions are labeled according to types of models. Type I: deterministic models with a single predicted spike train for all repetitions. Type II: deterministic models fitted independently on each repetition such that there are different predicted spike trains for each repetition. Type III: stochastic models.

analogous D_p^* and CF2* are more appropriate to quantify whether neuron models match real neurons.

4.6 Conclusions

Spikes as Vectors in a Vector Space

In the formalism proposed by [Carnell and Richardson, 2005], the inner product between spike trains yields natural definitions of distance or angular separation. The principle encompasses existing similarity measures such as those proposed by [van Rossum, 2001] or [Schreiber et al., 2003]. We noted that the coincidence factor of [Kistler et al., 1997] as well as the VP measure (for large q values) can also be related to an inner product and thus to vector space measures. In a direct extension of the formalism we introduced the population activity, that is, the vector made by averaging the spike-train vectors within a sample and showed that the scalar product between two population activities reduces to the commonly used correlation factor for PSTHs [Eggermont et al., 1983, Gawne et al., 1991, Paninski et al., 2005, David and Gallant, 2005, Petersen et al., 2008].

When spike trains, or pools of spike trains, are seen as vectors in a vector space, it is natural to consider quantities such as distance or angular separation to compare these spike trains. We have shown that angular separation can be interpreted as a similarity measure sensitive to the timing of events only, while the distance is a similarity measure sensitive to the timing of

events and the overall firing rate (Sect. 4.3). We confirmed this difference with surrogate spike trains in Sect. 4.4.3.

Kolmogorov-Smirnov Test

The improved similarity measures we have presented all boil down to a distance between PSTHs. As mentioned in the introduction, this is just one way of taking into account the intrinsic variability. Another option would be to apply a Kolmogorov-Smirnov test or use Kullback-Leibler divergence to assess if the distances across spike trains within one sample have the same probability distribution as the distances across spike trains from two different samples. While the Kullback-Leibler divergence is notoriously hard to estimate for small samples sizes [Akaike, 1973], the Kolmogorov-Smirnov (K-S) test could provide a measure of distance between two sets of spike trains. The basic idea is to compare the distribution of pairwise distances $D(S_i^{(x)}, S_j^{(x)})$ within a set to that of the pairwise distances $D(S_i^{(x)}, S_j^{(y)})$ across the set X and the set Y , where D can be calculated with any of the commonly used spike train metrics. We will pursue the direct analysis of this alternative method elsewhere. Preliminary results indicate that a K-S test on sets of spike trains can efficiently discriminate point processes with different level of stochasticity and that there would be no sample bias (results not shown).

Beyond the PSTH

Consider two classes of perfectly periodic spike trains. The first class contains doublets of spikes with a period of $2T$, the second class is made of single spikes with a period of T , but for both processes the phase changes randomly between one trial and the next so that both processes have the same constant firing intensity. Because the PSTH is identical for both classes the two classes can not be distinguished with the methods presented in this paper, although they would appear immediately distinguishable with parametric methods [Paiva et al., 2009a] or with the spike pattern metric of [Victor and Purpura, 1996] under the K-S test described in the previous paragraph.

Small-Sample Bias

We studied multiple types of spike-train similarity measures in terms of their potential to discriminate between different probabilistic point processes. We avoided the variability introduced by the size of the coincidence kernel [Shimazaki and Shinomoto, 2007, Paiva et al., 2009b] by considering coincidence kernels normalized by a condition on the integral. We have shown that many of these measures perform badly on simple tasks and that this can be remedied by taking into account a small-sample bias in the comparison of firing intensities. This bias is relevant when the sample consists of less than approximately 400 spike trains (see Fig. 4.8) and it can lead to dramatic failures when the available sample contains less than 30

spike trains.

From our case studies we conclude that even if the spike-time metrics are valid measures for comparing spike-trains, one cannot correctly discriminate between the underlying – stochastic – point processes by simply looking at their element-by-element average. For each similarity measure we proposed an improved version and we showed that the novel variant is indeed sensitive to the statistical features of the point processes generating the spike trains. The relevance of the small-sample bias has been previously shown to be important in the comparison of PSTH [David and Gallant, 2005, Petersen et al., 2008].

Implication for the Classification of Neuronal Responses

Our improved similarity measures for classification should partially improve the estimates of mutual information, even though there will always remain a sample bias because higher-order statistics are difficult to estimate from a limited number of elements [Treves and Panzeri, 1995, Panzeri et al., 2007]. The extent to which our improved measures will increase classification performance depends strongly on how the different classes differ in terms of the statistical properties of the point processes. For instance, we expect that if the different classes differ in their spike timings, there should be a smaller advantage than when the different classes differ in their spike-time jitter or overall firing rate (Fig. 4.4). We have shown an example of the improved classification using the variable locked fraction scenario in Sect. 4.4.7. In this example, the improved similarity measure increased the amount of information recovered. However, further investigations would be required to assess evaluate how significant is this increase with real spike trains.

Implication for the Measures of Intrinsic Reliability

Measures of intrinsic reliability such as R_X in Eq. 4.15 should be free of sample bias because they do not depend on $\|\hat{\mathbf{v}}\|$. Still, R_X can be biased when L_x is small and has a high variance, since the inverse of an unbiased estimator is not itself unbiased [Voinov, 1985]. The methods presented herein do not correct for this specific type of bias. On the other hand, our analytical results introduced a geometrical interpretation of the intrinsic reliability R_X : it corresponds to the variance of spike trains vectors around the population activity (that is, their center-of-mass, or PSTH). We also noted that intrinsic variability [Croner et al., 1993] corresponds to intrinsic reliability [Schreiber et al., 2003] under different normalizations.

Implications for Fitting: The Deterministic Bias

Why should none of the standard similarity measures be used for fitting neuronal processes? We have shown that a simple element-by-element average induces a preference for more deterministic models when a spike-train similarity measure is used as a cost function for fitting a mathematical model to data (Sect. 4.3.1 and 4.5.3). Moreover, we have indicated that

the parameters obtained when optimizing a biased similarity measure are wrong. When the available sample size is small, we have shown that it is possible to use bias-corrected measures, otherwise the estimated parameters such as the level of stochasticity in the model will be inadequately estimated. In general, element by element comparison yields an underestimate of the similarity between a stochastic model and a real neuron. This deterministic bias turns out to be responsible for shifting the maximum of the similarity measures such as coincidence factor away from the maximum likelihood solution, which explains a previously reported discrepancy [Shinomoto, 2010]. Perhaps the most striking example of the importance of the deterministic bias was the re-ordering of the results of the single-neuron prediction challenge (Sect. 4.5.4).

5 Parameter Extraction and Classification of Three Neuron Types Reveals two Different Adaptation Mechanisms

Cortical information processing originates from the exchange of action potentials between many cell types.¹ In order to capture the essence of these interactions it is of critical importance to build mathematical models that reflect the characteristic features of spike generation in individual neurons. We propose a framework to automatically extract such features from current-clamp experiments, in particular the passive properties of a neuron (i.e. membrane time constant, reversal potential and capacitance), the spike-triggered adaptation currents as well as the dynamics of the action potential threshold. The stochastic model that results from our maximum likelihood approach accurately predicts the spike times, the subthreshold voltage, the firing patterns, and the type of frequency-current curve. Extracting the model parameters for three cortical cell types revealed that cell types show highly significant differences in the time course of the spike-triggered currents and moving threshold, that is, in their adaptation and refractory properties, but not in their passive properties. In particular, GABAergic fast-spiking neurons mediate weak adaptation through spike-triggered currents only, whereas regular spiking excitatory neurons mediate adaptation with both moving threshold and spike-triggered currents. GABAergic non-fast-spiking neurons combine the two distinct adaptation mechanisms with reduced strength. Differences between cell types are large enough to enable automatic classification of neurons into three different classes. Parameter extraction is performed for individual neurons so that we find not only the mean parameter values for each neuron type, but also the spread of parameters within a group of neurons, which will be useful for future large-scale computer simulations.

5.1 Introduction

Cortical neurons exhibit a variety of different firing patterns in response to step currents [Connors and Gutnick, 1990, Markram et al., 2004] which has led to intricate electrophysiology

¹Text copied from Mensi, S., Naud, R., Pozzorini, C., Avermann, C., Petersen, C., H. and Gerstner, W., under review (full citation in the References).

Chapter 5. Parameter Extraction and Classification of Three Neuron Types Reveals two Different Adaptation Mechanisms

ical characterization schemes for three main neuronal cell-types (regular-spiking excitatory neurons (Exc), GABAergic fast-spiking neurons (FS) and GABAergic non-fast-spiking neurons (NFS)). Characterization is often done manually by visual observation of the firing patterns or, more systematically, by automatic extraction of a few parameters such as the ratio of the first to the second interspike interval, the minimum voltage of the spike afterpotential, the width of an action potential etc. The definition of these parameters is usually made arbitrarily. Such a characterization scheme is relevant for a biological classification of cell-types since it reflects the biochemical composition of the cells [Toledo-Rodriguez et al., 2004]. On the other hand, parameters such as the spike width or spike height have little to do with the computation made by the neuron when transforming its inputs into an output spike train. The question then arises whether characterization and classification of the cell types is possible on the properties affecting the conversion of synaptic inputs into a spike.

In the community of computational neuroscience it has been established over the last 20 years that simplified spiking neuron models are capable of reproducing the variety of firing patterns that have been found in experimental preparations, including delayed spike onset, bursting, strong or weak adaptation, refractoriness, etc [Gerstner et al., 1996b, Izhikevich, 2004, 2007, Brette and Gerstner, 2005, Naud et al., 2008, Touboul and Brette, 2009]. All of these models belong to the family of generalized integrate-and-fire models, but vary in the way the standard leaky integrate-and-fire model [Lapicque, 1907, Stein, 1967b] is generalized. Features to upgrade the simple integrate-and-fire model include spike after-currents [Baldissera et al., 1976, Gerstner et al., 1996b, Benda and Herz, 2003, Izhikevich, 2007, Paninski et al., 2004, Brette and Gerstner, 2005], dynamic threshold [Hill, 1936, Fuortes and Mantegazzini, 1962, Chacron et al., 2003, Jolivet et al., 2006, Badel et al., 2007], smooth spike initiation [Fourcaud-Trocme et al., 2003, Latham et al., 2000, Brette and Gerstner, 2005] and linearized subthreshold currents [Richardson et al., 2003, Izhikevich, 2007]. Important questions are then: which of these features are needed for basic cortical computation? How many levels of complexity do we have to add to account for relevant features of cortical dynamics? Is the spike-frequency adaptation mediated by moving thresholds or spike-triggered currents?

To answer these questions several groups used what one could call the Turing test for single neurons [Keat et al., 2001, Pillow et al., 2005, Paninski et al., 2005, Jolivet et al., 2006, 2008a, Kobayashi et al., 2009, Gerstner and Naud, 2009]: a somatic current is injected into a cortical neuron *in-vitro* while its response is being recorded. The modeler then asks how the model response compares with the neuronal response on data that was not used for optimizing the parameters of the model. This is done both qualitatively (for instance by reproducing the firing patterns [Izhikevich, 2004, Naud et al., 2008, Touboul and Brette, 2009, Mihalas and Niebur, 2009]) and quantitatively. The quantitative test consists of predicting the correct spike-timing and subthreshold voltage of the real neuron. Obviously, a simplified neuron model is not expected to work across the whole range of stimuli that one can artificially design and apply to electrophysiological experiments. The optimal version of this test would therefore use a stimulus that is similar to the one a real neuron receives in an appropriate *in-vivo* situation [Monier et al., 2008], so that the activity of the neuron is within a range that can be expected *in-vivo* [Crochet and Petersen, 2006, Poulet and Petersen, 2008, Gentet et al., 2010]. The

optimal stimulus will appear as a noisy time-series reminiscent but not identical to the white noise previously used for characterizing the input/output relationship [Bryant and Segundo, 1976, Marmarelis and Marmarelis, 1978, de Ruyter van Stevenick and Bialek, 1988]. We say that the model is 'good enough' if a neutral expert is not able to distinguish the activity of the model from that of the real neuron. We call this the '*Turing Test*' of neuron models in analogy to the test of intelligence in computer programs suggested by Alan Turing over sixty years ago [Turing, 1950].

The aim of this paper is three-fold. Firstly, we present a systematic method for extracting the parameters defining the properties affecting the conversion of synaptic input into a spike emission at the level of the soma. The method relies on the separation of the parameters affecting the subthreshold voltage and those affecting the firing threshold and its dynamics. Our method improves previously described methods for extracting spike-triggered currents [Paninski et al., 2004, Jolivet et al., 2006, Badel et al., 2007] and dynamic threshold [Azouz and Gray, 2000, Chacron et al., 2003, Badel et al., 2007]. We use this method to fit experimental data from fast spiking GABAergic neurons (FS), non-fast spiking GABAergic neurons (NFS) and excitatory cells (Exc). For each neuron type, we find the simplest model that reproduces the activity of neurons on data that was not used for parameter optimization. The models we extract reproduces the excitability type of GABAergic FS, GABAergic NFS and Exc neurons [Connors and Gutnick, 1990, Tateno et al., 2004, Markram et al., 2004] to a high degree of accuracy so that on average above 80% of the predictable spikes can be predicted while the mismatch in subthreshold voltage prediction is less than 2 mV.

Secondly, we ask which are the features essential for the neuron model to pass the Turing Test. We consider subthreshold resonance, conductance- versus current-based adaptation as well as current- versus threshold-based adaptation. Our results reveal the importance of adaptation currents and of a moving threshold with time constants that can be as long as hundreds of milliseconds. We find that different types of neurons use moving thresholds and spike-triggered currents differently to mediate refractoriness and spike-frequency adaptation. Thirdly and finally, we show that the parameters of the adaptation currents and dynamic threshold can be used for an automatic classification of the electrophysiological traces into three well-separated classes, whereas the passive parameters alone do not contain a sufficient amount of information to do so. We observe that the three neuron types have very different threshold dynamics and that efficient classification can be done using only the parameters regulating the dynamics of the threshold.

We expect that applying the automatic fitting method on a greater database could allow an unsupervised classification of the extracted computational properties, which would open the possibility to detect the potential spread of parameters within a given class of neurons and therefore avoiding a forced classification if in reality parameters are continuous. Finally, we expect that our method of parameter extraction once combined with similar results for synapses, dendrites and connectivity patterns will enable us to build network models where neuronal parameters and firing patterns reproduce not only the mean 'typical' firing of different neuron types, but also the spread of firing characteristics within and between classes of neurons.

5.2 Material and Methods

5.2.1 Model Dynamics

We consider a family of integrate-and-fire models augmented with a spike-triggered current $\eta(t)$ [Gerstner et al., 1996b, Paninski et al., 2005] and stochastic spike emission (escape noise [Plesser and Gerstner, 2000a, Gerstner and Kistler, 2002, Pillow et al., 2008]). In these models the somatic voltage is deterministic and follows the differential equation:

$$C\dot{V}(t) = -g_l(V(t) - E_l) + \sum_{\{\hat{t}_j\}} \eta(t - \hat{t}_j) + I(t) \quad (5.1)$$

C , g_l and E_l are the passive membrane parameters of the neuron: the membrane capacitance, conductance and the leak potential, respectively. $I(t)$ is a time-dependent input current. Adaptation of the subthreshold membrane potential is mediated by an adaptation current η which is triggered at the firing time \hat{t}_j ; contributions from previous spikes in the spiking history accumulate by summation of the contributions of $\eta(t - \hat{t}_j)$ over all spike times $\{\hat{t}_j\} = \{\hat{t}_1, \hat{t}_2, \dots\}$. By convention, the current η in Eq. 5.1 is depolarizing when its amplitude is positive and hyperpolarizing otherwise. Even though the reset of the voltage after a spike is equivalent to a short hyperpolarizing pulse, we prefer to work with an explicit reset, that is discussed below. The spike-triggered current could be any function of time, but the first-order effect of a single adaptation ion channel (such as M-currents or the calcium-dependent potassium current [Benda and Herz, 2003]) yields an exponentially decaying current after each spike [Brunel, 2010]. Allowing for N such linearized ion-channels, the resulting spike-triggered current is a sum of exponentially decaying contributions with different amplitude b_i and time constant τ_i :

$$\eta^{(N)}(t) = \sum_i^N b_i \exp\left(-\frac{t}{\tau_i}\right) \Theta(t) \quad (5.2)$$

where the Heaviside step function:

$$\Theta(t) = \begin{cases} 0 & \text{if } t < 0 \\ 1 & \text{otherwise} \end{cases} \quad (5.3)$$

ensures that only past spikes can have an influence. With a short time constant τ_i this linearized ion-channel mediates refractoriness, with a long time constant the effect of multiple spikes can accumulate, and the current will lead to spike frequency adaptation.

One can see Eq. 5.2 as an expansion of $\eta(t)$ in a linear combination of exponential basis

functions. For fitting purposes we sometimes choose another set of basis and write $\eta(t)$ as:

$$\begin{aligned}\eta(t) &= \sum_{k=1}^K a_k f^{(k)}(t) \\ f^{(k)}(t) &= \text{rect}\left(\frac{t - T_k}{\delta T_k}\right) = \Theta\left(t - \left(T_k - \frac{\delta T_k}{2}\right)\right) \cdot \Theta\left(T_k + \frac{\delta T_k}{2} - t\right)\end{aligned}\quad (5.4)$$

where $f^{(k)}(t)$ are rectangular basis functions centered at T_k and of width δT_k . The index k runs from 1 to K where K is the total number of rectangular functions used for fitting. The parameters T_k , δT_k and K are meta-parameters fixed *a priori* so that the shape of the spike after-currents is controlled by the coefficient a_k (see also [Paninski et al., 2005]). Parameters T_k , δT_k and K are chosen to give sufficient freedom to span all possible η while avoiding overfitting.

Equation 5.4 will be used to extract the precise shape of the spike after-current $\eta(t)$ without assumptions with respect to its shape (single exponential, double exponential, etc.). Yet, once we have the rectangular-based $\eta(t)$ we want to test whether a simpler model made of a combination of a small number of exponential processes also works. To this end we use Eq. 5.2 with parameters τ_i optimized during the fitting procedure.

As a variant of the model defined in Eq. 5.1, we can also turn the spike-triggered current η into a spike-triggered conductance. This is done by replacing η in Eq. (5.1) by $\eta_{\text{cond}}(V - E_k)$, where η_{cond} has now units of conductance and E_k is the reversal potential.

The passive properties of a neurons, such as the reversal potential E_l , the leak conductance g_l and the capacitance C , characterize the linear response of a neuron to current injection and so form the membrane filter κ of the neurons. The membrane filter arises from the solution of Eq. 5.1 using the initial condition $V(0) = E_l$:

$$V(t) - E_l = \int \kappa(t-s)I(s)ds + \int \kappa(t-s) \sum_{\{\hat{i}\}} \eta(s - \hat{t})ds \quad (5.5)$$

where $\kappa(t)$ acts as a low-pass filter on the current I and on the spike-triggered currents:

$$\kappa(t) = \frac{1}{g_l} \exp\left(\frac{-t}{\tau_m}\right) \quad (5.6)$$

with the membrane time constant $\tau_m = \frac{g_l}{C}$. Eq. 5.5 is known as the Spike-Response Model (SRM; [Gerstner et al., 1996b]). It summarizes the dynamics of the membrane potential as made of two parts: the effect of the input current (first term on the right-hand side of Eq. 5.5) and the effect of the action potential (second term on the right-hand side of Eq. 5.5). The functions $\kappa(t)$ and $\eta(t)$ are *response kernels* which depict the membrane filter and the shape of the spike-after currents, respectively. Eq. 5.5 taken without restrictions on the shape on $\eta(t)$ and $\kappa(t)$ is very general and can take into account the effect of multiple subthreshold or spike-triggered ion-channels.

The linear filter $\kappa(t)$ may or may not consist of a single exponential decay. It is known that

Chapter 5. Parameter Extraction and Classification of Three Neuron Types Reveals two Different Adaptation Mechanisms

subthreshold currents, like I_m or I_h [Sabah and Leibovic, 1969, Mauro et al., 1970, Koch, 1984] generate subthreshold resonances or delayed spiking responses to steps [Naud et al., 2008, Izhikevich, 2004]. These subthreshold currents provide a current-to-voltage filter with a negative bump (for resonance) or with an exponential decay with two time constants (for delayed firing onset). Thus, to be sure that our assumption of an exponential filter is not too restrictive, we compute the Wiener-Hopf optimal filter [Jolivet et al., 2004] (i.e. we allow for arbitrary shape of κ in Equation 5.5) away from the spikes such that we are left with only the left term on the right-hand side of Eq. 5.5.

The deterministic voltage in Equation 5.1 is integrated so as to yield a voltage $V(t) = V(t|I, \theta, \{\hat{t}_j\})$ which depends on the input current $I(t')$ for $t' < t$; on the past firing times $\hat{t}_1, \hat{t}_2, \hat{t}_3, \dots < t$, as well as the set of parameters (denoted as θ) chosen for the model. This deterministic voltage is the central variable for the stochastic spike generation. Using the escape-rate picture [Plesser and Gerstner, 2000a, Gerstner and Kistler, 2002, Paninski et al., 2005], the conditional probability intensity of emitting a spike is a nonlinear function of the instantaneous distance from the voltage threshold. More precisely the stochastic spiking process follows an inhomogeneous point process with conditional firing intensity $\lambda(t|V, V_T)$ given by:

$$\lambda(t|V, V_T) = \lambda_0 \exp\left(\frac{V(t) - V_T(t)}{\Delta V}\right) \quad (5.7)$$

where λ_0 has units of s^{-1} , so that $\lambda(t)$ is in Hz , V is the membrane potential, $V_T(t)$ is the voltage threshold and ΔV describes the sharpness of the exponential function. In principle, any function of $V - V_T$ would be possible, but it was shown previously that the exponential function explains experimental data well [Jolivet et al., 2006]. Note that the value of λ_0 can be chosen arbitrarily since a scaling $\lambda_0 = a\lambda_1 = \lambda_1 \exp(\log a)$ can be compensated by a shift of $\frac{V_T}{\Delta V}$ by $\log(a)$.

The voltage threshold $V_T(t)$ can be either a constant threshold, with $V_T(t) = V_0$ or a time-dependent sliding threshold that implements an additional source of refractoriness [Gerstner and Kistler, 2002, Chacron et al., 2003, Badel et al., 2007]. In our model, the dynamic threshold is a cumulative function:

$$\begin{aligned} V_T(t) &= V_0 + \sum_{\{\hat{t}_j\}} \gamma(t - \hat{t}_j) \\ \gamma(t) &= \sum_{p=1}^P c_p f^{(p)}(t) \\ f^{(p)}(t) &= \text{rect}\left(\frac{t - T_p}{\delta T_p}\right) \end{aligned} \quad (5.8)$$

where V_0 is the threshold baseline, $\{\hat{t}_j\}$ the set of spike times that were emitted before t and $\gamma(t)$ is the spike-triggered motion in voltage threshold. The shape of γ is defined by the coefficients c_p which control the amplitude of a set of the rectangular basis $f^{(p)}(t)$.

When a spike is emitted the numerical integration is stopped and after a short absolute

refractory period T_{refr} , V is reset to E_{reset} . This voltage reset is typical for a large class of Integrate-and-Fire (IF) models.

We call IF models with N exponential adaptation currents, as defined in Eq. 5.2, IF_N . Models with the estimation of adaptation currents in terms of bins (i.e. rectangular basis functions), as in Eq. 5.4, are called IF_η . And if the model's adaptation is based on conductance η_{cond} we call it $\text{IF}_{\eta_{\text{cond}}}$. When a dynamic threshold $V_T(t)$, as specified in Eq. 5.8, is used instead of the constant threshold, we denote it with '+dyn'. For example, $\text{IF}_2 + \text{dyn}$ is an IF model with two spike-triggered currents and the dynamic threshold defined by $\gamma(t)$.

5.2.2 *In vitro* Two-photon Microscopy and Whole-cell Recordings

Brains of 17-22 day old GAD67-GFP knock-mice [Tamamaki et al., 2003] were removed and quickly placed into an ice cold modified ACSF [Bureau et al., 2006] (containing in mM: 110 choline chloride, 25 NaHCO_3 , 25 D-glucose, 11.6 sodium ascorbate, 7 MgCl_2 , 3.1 sodium pyruvate, 2.5 KCl, 1.25 NaH_2PO_4 and 0.5 CaCl_2). 300 μm oblique slices (parasagittal 35° away from vertical) were cut with a vibrating slicer (Leica VT1000S, Germany) and subsequently transferred into standard ACSF (containing in mM: 125 NaCl, 25 NaHCO_3 , 25 D-glucose, 2.5 KCl, 1.25 NaH_2PO_4 , 2 CaCl_2 and 1 MgCl_2) aerated with 95% O_2 / 5% CO_2 at 35°C for 15 minutes. Afterwards slices were maintained at room temperature for at least 30 minutes prior to use to allow for recovery from the slicing procedure. GFP-expressing neurons were visualized using a two-photon microscope (Prairie Technologies, USA). Infrared two-photon excitation light of 880 nm was generated by a MaiTai laser (Spectra-Physics, France) and focused into the slice tissue through a 40x 0.8NA water immersion objective (Olympus, Japan). Detection of bandpass-filtered green fluorescence (525 \pm 35 nm) was achieved using PMTs above the objective and below the condenser. Infrared light was passed through a Dodt contrast element (Luigs & Neumann, Germany) and detected by an additional PMT to allow creation of a high contrast view of the brain tissue. Whole-cell recordings were carried out at 33°C in standard ACSF. Borosilicate pipettes of 5-7 M Ω resistance were used. The patch-pipette intracellular solution contained (in mM): 135 K-gluconate, 4 KCl, 4 Mg-ATP, 10 Na_2 -phosphocreatine, 0.3 Na-GTP, 10 HEPES (pH 7.3, 280 mOsm). A Multiclamp 700A amplifier (Axon Instruments, USA) was used for whole-cell recordings. Data were low-pass Bessel filtered at 10 kHz and digitized at 20 kHz with an ITC-18 acquisition interface (HEKA Electronics, Germany). The measured membrane potential was not corrected for the liquid junction potential.

Cortical layers 2/3 GABAergic inhibitory neurons were distinguished from excitatory neurons (Exc) by their expression of GFP [Gentet et al., 2010]. GFP-expressing GABAergic neurons were further classified into fast-spiking (FS) and non-fast-spiking (NFS) neurons with respect to their action-potential (AP) kinetics upon somatic current pulse injection. An AP half-width lower than 0.75 ms was used as a selection criterion for GABAergic FS neurons.

Chapter 5. Parameter Extraction and Classification of Three Neuron Types Reveals two Different Adaptation Mechanisms




Figure 5.1: **Stimulation protocol and dynamics of a IF model.** *A*: Six spike trains are generated by an inhomogeneous Poisson process with piecewise constant rates. Rates are switched simultaneously for all neurons. The rate changes at intervals chosen randomly between 300 - 500 ms. *B*: The excitatory spike trains are convolved with an excitatory postsynaptic current (EPSC) kernel that has a time constant of 2 ms to mimic AMPA-receptor dynamics and the inhibitory spike trains are convolved with an inhibitory postsynaptic current (IPSC) kernel with a time constant of 10 ms to mimic GABA-receptor dynamics. *C*: The six current traces are combined in a weighted sum, with weight w_1 to w_6 to produce the final input current. The weights are chosen to drive the neuron to a mean firing frequency between 5 and 15 Hz. *D*: An example of a model voltage trace (bottom) generated by the injection of the input current (top). Here the model has $N = 2$ spike-triggered currents (IF_2).

5.2.3 Stimulation Protocol: Synaptic-like current

We construct input currents $I_{syn}(t)$ as a weighted sum of six impinging spike trains constructed from independent inhomogeneous Poisson processes (see also [Poliakov et al., 1997]). Three spike-trains are convolved with a mono-exponential filter with time constant of 2 ms mimicking spike arrival at excitatory synapses and summed with weight $w_{1,2,3}$, while the three remaining ones are convolved with a 10 ms mono-exponential filter for the inhibitory spike trains and summed with weight $w_{4,5,6}$. The weights are chosen to drive the neuron to a mean firing frequency between 5 and 15 Hz. All the Poisson processes shared the same time-dependent intensity (firing rate) which is a concatenation of blocks of 300-500 ms duration each with a constant intensity chosen randomly from a uniform distribution between 0 and 50 Hz. The duration of the blocks was drawn from a uniform distribution (Figure 5.1 A, B, C). We call the time-dependent input $I_{syn}(t)$ constructed by the above procedure a synaptic-like current.

One minute of this synaptic-like current is injected in the cortical neurons repeatedly interleaving the repetitions with silent periods of 10 seconds. Note that in each repetition the cell received exactly the same time course of synaptic-like current which enables us to study peristimuli time histograms (PSTH) and reliability of the neuronal spiking. In order to compare

experimental results with our models, we used the first 30 seconds for fitting and reserved the last 30 seconds of each repetition for testing the performance of the models. These two subsets of the data are called training and test sets, respectively.

To assess the robustness of our fitting procedure on surrogate data, we also test it on 4 different stimulation paradigms. First we use a gradually increasing ramp of current of 30 seconds. Second we construct an input current as a series of 300 ms step current interleaving with 200 ms of silence, with increasing intensities. Third we used as a stimulation current a white noise with 0 mean that lasts for 30 seconds. Finally we construct a 30 seconds long input current made of colored noise, generated according to an Ornstein-Uhlenbeck process with correlation time constant of 4 ms. All the parameters of these currents are chosen so that the resulting input current produced an averaged firing frequency of approximately 15 Hz.

5.2.4 Performance Measurements

Since we are assessing the performance on the test set, the performance of the model will not increase by merely increasing the number of parameters, because over-fitting would occur on the training set and lower the performance on the test set. Two distinct criteria are used to evaluate the performance of our IF models: (i) the precise spike time² and (ii) the subthreshold voltage prediction.

Neurons, as well as stochastic models, have some trial-to-trial variations due to intrinsic noise. In neurons, this intrinsic noise is mainly due to channel noise and spontaneous synaptic events [Faisal et al., 2008]. To quantify the spike time prediction we used a method that corrects the bias due to the small number of available repetitions [Naud et al., To appear; 2012]³. We first use our optimal model to generate $R_m = 1000$ spike-time predictions. The next step is to compute a quantity between zero and one that measures how well the set of model spike trains R_m matches with the set of R_r experimental spike trains, where R_r is the number of repetitions available for the specific cell (lower index r for real neuron). To do so, we first count the total number of spikes of the model that fall within ± 4 ms of a spike in the a recorded spike train and average the count over all repetitions. We call the resulting raw measure of comparison between model (lower index m) and real neuron (lower index r): n_{mr} . Next we count the average number of coincidences between distinct repetitions of the model spike trains and average the results across all available repetitions. We call this quantity n_{mm} . Similarly, n_{rr} is the analogous quantity evaluated over the available repetitions recorded in the real neuron. Finally, we combine these numbers into the measure:

$$M_d^* = \frac{n_{mr}}{\frac{1}{2}(n_{mm} + n_{rr})} \quad (5.9)$$

which is an estimate of the percentage of predictable spikes that are correctly predicted with a precision of ± 4 ms. It can be shown that this quantity corresponds to a normalized distance

²Readers who have read Chapter 4 will be familiar with M_d^* and can skip the next paragraph to go directly to Eq. 5.10.

³Chapter 4

Chapter 5. Parameter Extraction and Classification of Three Neuron Types Reveals two Different Adaptation Mechanisms

between Peri-Stimulus Time Histograms (PSTHs) [Naud et al., To appear; 2012]⁴.

The subthreshold voltage prediction is evaluated using the root mean square error on the voltage (RMSE) defined by

$$\text{RMSE} = \sqrt{\frac{1}{T_1} \int_{\psi_1} (V_{\text{ref}}(t|I) - V_{\text{predicted}}(t|I, \theta, \{\hat{t}_{\text{ref}}\}))^2 dt} \quad (5.10)$$

where $V_{\text{ref}}(t|I)$ is the recorded voltage given an input current I , $V_{\text{predicted}}(t|I, \theta, \{\hat{t}_{\text{ref}}\})$ is the voltage predicted by a model with the set of parameters θ and the same input current. The index $\{\hat{t}_{\text{ref}}\}$ indicates that for the voltage comparison we force spikes at exactly the same time as in the recorded voltage trace. In other words, we ask: how close is the voltage of the model to that of the data, given the input and the firing times in the recent past. The squared error is integrated over a subset ψ_1 of the available data defined in the next section. T_1 is the total time for which the squared error is considered, $T_1 = \int_{\psi_1} dt$. Note that the RMSE does not take into account the experimentally recorded variability induced by intrinsic noise such as channel noise. Thus, it is impossible to know with precision the trial-to-trial subthreshold voltage variations. However one can find an approximation of this intrinsic subthreshold voltage fluctuations by computing RMSE_{ψ_4} , the root-mean-squared error between repetitions only on the subset of data ψ_4 (defined in the next section). RMSE_{ψ_4} takes only into account voltage recordings sufficiently far away from the spikes so as to avoid the effect of adaptation currents due to spikes emitted at different times across repetitions. RMSE_{ψ_4} is clearly an overestimation of the true subthreshold fluctuations because of long-lasting spike-triggered effects. Nevertheless it gives some insight on how good our models have to predict the subthreshold voltage. To quantify the ability of the models to replicate subthreshold voltage fluctuations we used the normalized 'Root Mean Square Error Ratio': $\text{RMSER} = \frac{\text{RMSE}}{\text{RMSE}_{\psi_4}}$, so that $\text{RMSER} \approx 1$ for a good reproduction of the subthreshold voltage and $\text{RMSER} > 1$ for inaccurate reproduction. All results are shown in terms of this normalized RMSER.

5.2.5 Fitting Procedure

To extract all the parameters of the models a four-step procedure is applied. The main steps of the method are illustrated in Figure 5.2. For the whole fitting procedure it is convenient to define appropriate subsets of the available data.

Since we do not want to model the exact shape of the action potential we cut out, around each spike time \hat{t} , a small segment of the data which we ignore. In the experimental data the spike time \hat{t} is defined to be the time when the membrane potential crosses a given voltage V_{detect} from below, here we set the detection threshold V_{detect} to 0 mV.

Let us, at each moment t , refer to the last previous spike timing as \hat{t}^{last} and to the next spike

⁴ *Idem*

time in the future as \hat{t}^{next} . The subsets of the data that we use for fitting are:

$$\psi_1 = \{t \mid t > \hat{t}^{(\text{last})} + T_{\text{refr}} \text{ and } t < \hat{t}^{(\text{next})}\} \quad (5.11)$$

$$\psi_2 = \{t \mid t > \hat{t}^{(\text{last})} + T_{\text{refr}} \text{ and } t < \hat{t}^{(\text{next})} - 2 \text{ ms}\} \quad (5.12)$$

with T_{refr} the absolute refractory period. Thus ψ_1 represents a set of recordings times where the voltage is subthreshold and outside the absolute refractory period. ψ_2 further removes 2 ms of data before each spike. We will also use two other subsets of the data:

$$\psi_3 = \{t \mid t > \hat{t}^{(\text{last})} + T_{\text{adapt}} \text{ and } t < \hat{t}^{(\text{next})} - 2 \text{ ms}\} \quad (5.13)$$

$$\psi_4 = \{t \mid \psi_3^{(1)} \cap \psi_3^{(2)} \dots \cap \psi_3^{(R)}\} \quad (5.14)$$

Thus ψ_3 removes a period of time T_{adapt} after the spikes where $T_{\text{adapt}} > T_{\text{refr}}$. We use for T_{adapt} a period of 200 ms when the recording has a high average firing frequency (> 5 Hz) but we use $T_{\text{adapt}} = 500$ ms otherwise. For our recordings made of multiple repetitions, we only consider the subset of times that are separated by a period of at least T_{adapt} from any previous spike, from any repetition. This subset is therefore the intersection between the subsets ψ_3 of each repetition. Each subset $\psi_1, \psi_2, \psi_3, \psi_4$ will be used in different steps of the fitting procedure, which we describe now:

Step 1. We measure the average spike shape by computing the spike-triggered average voltage [Jolivet et al., 2004]. The refractory period T_{refr} and reset potential E_{reset} are defined by the value of the minimum of the after hyper-polarization (AHP), see Figure 5.2 B, left panel. For Excitatory neurons the AHP trace does not have any local minimum, and we arbitrarily chose $T_{\text{refr}} = 4$ ms, and fixed E_{reset} to the corresponding voltage at this moment. As long as it remains short, the exact value of T_{refr} has no impact on the quality of the fit.

Step 2. We extract from the experimental voltage traces the first-order estimate of the time derivative using the finite difference: $\dot{V}_t = (V_{t+1} - V_t)/dt$. Here, V_t denotes the binned voltage time-series as obtained from the recordings, using a bin size dt of 0.05 ms. We use the derivative on the data set ψ_2 so as to optimize the parameters by minimizing the sum square error (SSE) between the observed voltage time-derivative \dot{V}_t and that of the model. Rewriting Eq. 5.1, it is clear that the parameters $\vec{\theta}_2 = \{-\frac{g_l}{C}, \frac{g_l E_l}{C}, \frac{1}{C}, \frac{a_1}{C}, \dots, \frac{a_K}{C}\}$ act linearly on the model's voltage time-derivative $\dot{V}_t^{(\text{mod})}$:

$$\dot{V}_t^{(\text{mod})} = -\frac{g_l}{C}(V_t - E_l) + \frac{1}{C}I_t + \sum_{k=1}^K \frac{a_k}{C} \sum_{\{\hat{t}_j\}} f_{t-\hat{t}_j}^{(k)} \quad (5.15)$$

For notational convenience, we will write the above equation as a matrix equation, defining $\dot{\mathbf{V}}^{(\text{mod})}$ as the vector of the binned voltage time-derivatives and \mathbf{X} as a matrix

Chapter 5. Parameter Extraction and Classification of Three Neuron Types Reveals two Different Adaptation Mechanisms




Figure 5.2: **Fitting protocol.** *A*: Experimental data set, injected current $I(t)$ (top panel), recorded voltage trace $V(t)$ (middle panel) and extracted spike times $y(t)$ (bottom panel). *B*: Example of extracted parameters. (i) All the spikes are aligned (black) and averaged (red) to obtain the spike shape. (ii) Adaptation current η (black) and a double exponential fit (blue). Inset, membrane filter extracted with a Wiener-Hopf optimal filter method (black) and single exponential fit (red). (iii) Dynamic threshold γ (black) and a double exponential fit (green), the grey area represent one standard deviation. *C*: An example of model voltage trace ($IF_2 + \text{dyn}$) produced by the injection of the input current $I(t)$ and the corresponding experimental voltage. Experimental data (black), predicted voltage trace (red), voltage threshold (green) and adaptation current (blue). The inset shows the subthreshold voltage over 100 ms in a region far away from any spikes.

with the vector of voltage \vec{V}_t (binned as a function of time) in the first column, a vector $\vec{1}$ of all ones in the second column, the vector of input current \vec{I}_t (binned as a function of time) in the third column, and the value $\sum_{\hat{t}_j} f_{t-\hat{t}_j}^{(k)}$ for $1 \leq k \leq K$ evaluated at the known spike times \hat{t} in the remaining K columns, such that the differential equation

5.15 becomes:

$$\dot{V} = \mathbf{X}\vec{\theta}_2 \quad (5.16)$$

The parameters can be estimated by minimizing the sum of squared errors $SSE(\vec{\theta}_2) = \|\dot{V}_{\psi_2} - \mathbf{X}_{\psi_2}\vec{\theta}_2\|^2$ between the voltage derivative of the experimental trace V_{ψ_2} and that of the model, $\mathbf{X}_{\psi_2}\vec{\theta}_2$, averaged over all points in the data set ψ_2 that comprise the voltage trace in the subthreshold regime. According to multi-linear regression theory [Weisberg, 2005], the optimal set of parameters is then given by

$$\hat{\vec{\theta}}_2 = (\mathbf{X}_{\psi_2}^T \mathbf{X}_{\psi_2})^{-1} \mathbf{X}_{\psi_2}^T \dot{V}_{\psi_2} \quad (5.17)$$

This method was used in [Paninski et al., 2005] as a linear method to maximize the likelihood of observing the measured \dot{V} time series. This step gives the passive parameters of the neurons and the adaptation current η (figure 5.2 B, middle panel, inset). Note that here we voluntarily discard the spikes from the data (because we only consider the subset ψ_2), but it is straightforward to apply the same linear regression method on the whole recordings (including the spikes). In this case the first bins in η model the shape of the spike, and an explicit reset of the voltage would no longer be necessary.

When a model IF_N is considered, the time constants τ_i and the amplitude b_i of the adaptation currents w_i are extracted from $\eta(t) = \sum_k a_k f^{(k)}(t)$ by the fit of N exponential functions. When a model $\text{IF}_{\eta_{\text{cond}}}$ is considered, we observe that if E_k is known a priori, then exactly the same protocol can be applied. So we perform the linear regression defined by Eq. 5.17 iteratively for a set of $\{E_k\}$ and chose the optimal \hat{E}_k to be the one that minimizes the SSE of the regression (see Figure 5.7 A and B).

Step 3. To extract the cumulative dynamic threshold γ from the data, we maximize the likelihood of generating the experimental spike train by our model. The log-likelihood for a spike train can be written in terms of the probability of observing a spike in each experimental time bin p_t by using Bayes theorem recursively in time [Paninski et al., 2004]:

$$\log L(\vec{\theta}_3) = \sum_{\{\hat{t}_j\}} \log(1 - p_t) - \sum_{\tilde{\psi}_2} \log(p_t) \quad (5.18)$$

where the set of spike times $\{\hat{t}_j\}$ is taken to be 0.5 ms before the peak of the spike and $\tilde{\psi}_2$ is the set of times ψ_2 with the spike times removed. The probability of having no spike in a time bin $[t, t + \Delta_T)$ is [Gerstner and Kistler, 2002]:

$$p_t = \exp\left(-\int_t^{t+\Delta} \lambda(t') dt'\right) \approx e^{-\lambda_t \Delta_T} \quad (5.19)$$

where the approximation holds for small Δ_T (here $\Delta_T = 0.05$ ms). λ_t is the discretized

Chapter 5. Parameter Extraction and Classification of Three Neuron Types Reveals two Different Adaptation Mechanisms

version of $\lambda(t)$. Using Eq. 5.7 we have:

$$\lambda_t = \lambda_0 \exp(\vec{X}_t \vec{\theta}_3) \quad (5.20)$$

where the vector \vec{X}_t is made of V_t , one, and $\sum_{\{\hat{i}_j\}} f_t^{(p)}$, and the parameters are $\vec{\theta}_3 = [\frac{1}{\Delta V}, -\frac{V_0}{\Delta V}, -\frac{c_1}{\Delta V}, \dots, -\frac{c_p}{\Delta V}]^T$. Note that to avoid correlations between the parameters for subthreshold voltage and the parameters that describe the spike emission process, we use the recorded voltage V_t and not the modelled voltage, as it would be done with a purely generative model.

Now, using λ_t defined in Eq. 5.20 with $\lambda_0 = 1/\Delta_T$ and using the fact that $\lambda_t \Delta_T$ is small, then the optimal set of parameter $\hat{\theta}_3$ is simply given by:

$$\hat{\theta}_3 = \arg \max_{\vec{\theta}_3} \left(\sum_{\{\hat{i}_j\}} \vec{X}_t \vec{\theta}_3 - \sum_{\psi_2} \exp(\vec{X}_t \vec{\theta}_3) \right) \quad (5.21)$$

With the exponential link function in Eq. 5.20 we are sure that the log-likelihood is a convex function of the parameters $\vec{\theta}_3$ [Paninski et al., 2005]. To perform the gradient-ascent of the log-likelihood function, the simplest method is perhaps to use a pre-programmed script (for instance `fminunc.m` in Matlab), but we used the iteratively re-weighted least-square method, also called Fischer's scoring method [McCullagh and Nelder, 1998].

Step 4. At this point all the parameters have been extracted, but in order to obtain an optimal spike time prediction in terms of M_d^* , we recompute the baseline threshold V_0 so that it maximizes $M_d^*(V_T)$. To do so we find the parameter V_0 that maximizes M_d^* through an exhaustive search over a large range of parameters V_0 .

All numerical computations have been done in Matlab (The Mathworks, Natwick, MA) on a Desktop computer. In practice our fitting procedure is straightforward and fast; it takes only few minutes on a desktop computer to extract all the parameters of a model from 10 seconds of voltage recordings and current injection producing a firing frequency of 10 Hz.

5.3 Results

Before turning to experimental data obtained from cortical excitatory neurons, GABAergic fast-spiking and GABAergic non-fast-spiking neurons we checked the consistency of our parameter extraction method on artificially generated data.

5.3.1 Efficiency and Accuracy of the Fitting Method on Surrogate Data

To test efficiency and accuracy of the fitting method, we first generated 60 seconds of surrogate data from an IF model augmented with 2 spike-triggered currents and a dynamic threshold

(IF₂ + dyn) and used our fitting method to retrieve its parameters. Thus we used an IF₂ + dyn model to fit data from another IF₂ + dyn model. The error in the estimated parameters and thus the prediction performance of the model depends on the amount of data used for fitting. Figure 5.3 shows how the fitting quality evolves as a function of the amount of data used for parameter extraction. We observe that the voltage prediction of the fitted model becomes better when a larger amount of data is used (see Figure 5.3 A and C). But even with a small amount of data (1 second), the predicted voltage is relatively accurate, producing a RMSE of 0.43 mV (the RMSE obtained with 15 seconds of data is 0.26 mV). The spike time prediction depends also on the size of the training set, leading to a $M_d^* = 0.79$ with 1 second of data while $M_d^* = 0.99$ when 15 seconds are used. Furthermore the relative error ϵ in the parameter estimate (computed according to $\epsilon = \langle |\frac{\theta - \hat{\theta}}{\theta}| \rangle$, with θ the true parameters and $\hat{\theta}$ the extracted parameters) also shows a strong decrease as a function of training set size, from $\epsilon(1s) = 0.13$ to $\epsilon(15s) = 0.03$. Thus with less than 20 seconds of data our method enables us to retrieve the reference parameters to a high degree of accuracy (Fig. 5.3).

To be sure that parameters have reached their steady-state, we systematically compare RMSE, M_d^* and ϵ at each training set size. We find that if less than 14 seconds of recordings are used (Figure 5.3 C, shaded area) parameters are significantly different from their steady-state value as quantified by ϵ (two-sample t-test, $\alpha = 0.05$). So, when a reference model is fitted to itself, 14 seconds of recordings with 10 Hz firing frequency are sufficient to obtain a good fit in terms of RMSE, M_d^* and ϵ .

Since fitting a model to data (model extraction) is more difficult than fitting to itself (model identification) [Druckmann et al., 2008], we use for the fits to real neurons (as done in the next subsection) always a training set of at least 20 seconds and a separated test set that contains at least 100 spikes with firing frequencies from 0.1 to 40 Hz.

We checked that the fitting procedure works for other stimulation protocols. To study this we used a reference IF₂ + dyn stimulated with various kinds of widely used currents: series of steps, ramp, white noise and colored noise. We tested the fitted models on a test set made with the synaptic-like current. The results are shown in Table 5.1. We conclude that step current or synaptic-like currents are more informative and enable a more reliable parameter extraction than ramp current and white or colored noise injection. But across all protocols, extracting IF parameters is effective (error less than 15%) and we always obtain a $M_d^* > 0.8$ with a relatively low RMSE.

5.3.2 Accuracy of Fitted Models on FS, NFS and Exc

The accuracy of the fitting method is summarized in Fig. 5.4 for exemplars of FS, NFS and Exc neurons. With voltage prediction consistently below the intrinsic RMSE of the data, and spike time prediction above $M_d^* = 0.78$ we conclude that the simple neuron models, in combination with our fitting method accurately models the three different cell types.

Chapter 5. Parameter Extraction and Classification of Three Neuron Types Reveals two Different Adaptation Mechanisms




Figure 5.3: Assessment of the fitting procedure by fitting a model to a model. *A*: Upper panel: Example of injected current used as part of the test data set. Lower Panel: Corresponding voltage trace of the reference data (black) and a modeled voltage trace (red), the upper traces are obtained with a training set of 1 second and the lower traces with a training set of 15 seconds. *B*: Boxplot of the normalized parameters extracted with a training set of 15 seconds. A perfect fit of the reference parameters would lead to normalized parameters equal to units. All the parameters are within 5% of target value. Crosses denote outliers. *C*: M_d^* (red) as a quality measure of spike time prediction, RMSE (blue) as a quality measure of subthreshold voltage prediction, and the relative error in the parameters estimates ϵ (black) as a function of the size of the training set, error bars are one standard deviation. The shaded area corresponds to a training set size insufficient to reliably extract the optimal parameters. Stars denote last significant difference between M_d^* and ϵ at a given training set size and the same values obtained with the full training set (two sample t-test, $\alpha = 0.05$).

Our neuron models extracted by the above procedure are able to reproduce the typical behavior of the different cell types in terms of their firing patterns and frequency-input curve. To show this we reproduce an experiment done by Tateno et al. [2004], where the authors stimulate GABAergic FS and Exc neurons with step currents of 1 second at various but constant current strength and classify cells as a function of their frequency-intensity curves (f-I Curve). Using this method the authors conclude that GABAergic FS neurons have a step in the f-I curve (Type-II excitability) whereas Exc neurons exhibit a smooth f-I curve (Type-I excitability). Figure 5.5 *A*, *B* and *C*, which have been generated using our optimal models for GABAergic FS, GABAergic NFS and Exc neurons, are analogous to Figure 4 of the paper of Tateno et al.

Type of Stimuli	RMSE	M_d^*	Relative Error
	[mV]	[-]	[-]
ramp	0.15	0.81	0.14
step	0.42	0.97	0.04
white noise	1.77	0.90	0.11
colored noise	0.26	0.88	0.12
synaptic-like noise	0.29	0.98	0.04

Table 5.1: **Evaluation of the fitting procedure on different common types of stimulation.**

We evaluate the ability of our methods to extract parameters from various types of stimulation in the training set. After parameter extraction the model is evaluated on a test set consisting of synaptic-like current as defined in section ‘Stimulation Protocol: Synaptic-like current’. The evaluation criteria are: the RMSE on the voltage, the M_d^* factor as a quality measure for spike time prediction, and the relative error in the parameter estimates. Step current and synaptic-like current are the most informative stimulation protocols.

[2004]. Our FS models are typical Type-II with a minimal critical frequency (f_c) for the first ISI of 15.5 Hz and a steady state critical frequency of 5.45 Hz. When constant currents are used, smaller frequencies are impossible so that the f-I curve exhibits a finite jump at f_c . The critical frequency can be traced back to the first zero-crossing of the effective spike-triggered adaptation. The emergence of the Type-II behavior from our FS model comes from the depolarizing segment in the spike-triggered current $\eta(t)$ shown in the inset of Figure 5.5 A. If the amplitude of the stimulation current is sufficient to evoke a spike, then the depolarizing part of $\eta(t)$ ensures that repetitive firing follows as long as the same stimulating current is maintained. The minimal firing frequency is roughly set by the time to the maximum of the depolarization.

For NFS and Exc models we obtain a smooth transition between silence and firing, which is the behavior of the Type-I neurons consistent with experiments of Tateno et al. [2004]. Moreover Figure 5.5 D-F provides examples of firing patterns from our optimal models for each cell type, analogous to the examples provided in Figure 5 of Tateno et al. [2004].

5.3.3 Essential features for Subthreshold Voltage Prediction

The membrane filter for each neuron type is well approximated by a single exponential function (Figure 5.6A-C), which shows that voltage-dependent subthreshold currents are small and can be neglected. Moreover it follows from this finding that, for the neuron types studied here, subthreshold resonance is not the most important factor for accurate prediction of each of the neuron types. Specific spike-triggered currents, however, are necessary to reproduce the subthreshold voltage.

The spike-triggered current $\eta(t)$ corresponds to the stereotypical current that flows into the neurons after a spike. After the onset of the spike, the dominant features of the spike-triggered currents consist of (i) the current that produces the spike (i.e. the spike shape), (ii) a short refractory current that follows the end of the AP (just a few milliseconds after the spike onset)

Chapter 5. Parameter Extraction and Classification of Three Neuron Types Reveals two Different Adaptation Mechanisms




Figure 5.4: Examples of voltage traces and spike train predictions for the 3 neuron classes. *A*: Injected current. The waveform of the injected current is the same for the 3 neurons but we rescale it to achieve an average firing frequency between 5 and 15 Hz, see ‘Material and Methods’ for details. Thus the scale bar is different for each neuron type. *B*: Upper panel: left, experimental voltage trace (black) and modeled voltage trace (red) for a GABAergic FS neuron. Inset: zoom on 150 ms of subthreshold voltage. Subthreshold voltage prediction is quantified by the measure RMSE. Lower panel: Raster plot of the emitted spikes during 11 repetitions for the experimental data (black) and for the model prediction (red). Spikes that occur reliably in the neuron are reliably reproduced by the model. Spike timing prediction is quantified by the measure M_d^* . *C*, as in *B* but for a GABAergic NFS neuron. *D*, as in *B* but for a Excitatory neuron.

and (iii) a long cumulative current that can adapt the spike-frequency of the neurons. Here, we consider only parts ii) and iii) since these currents are the most important part for the processing of information done by the neuron [Koch, 1999, Hille, 1992, Jolivet et al., 2008a]. To investigate the shape of the spike-triggered currents, we measure the cumulative adaptation current $\eta(t)$ for each neuron and group these by cell type. Figure 5.6 *D-F* shows the mean




Figure 5.5: **Instantaneous firing frequency of models as a function of the intensity of a step current, (f-I curve).** *A*: f-I curve for a model of GABAergic FS neuron: black trace, 1st interspike interval: red trace, steady state. Since there is a non-zero onset firing frequency f_c (critical frequency) at the threshold value of the constant step current, the FS model has a type-II excitability. Inset: the effective spike-triggered adaptation as extracted by our method and fully described in section ‘Effective Spike-triggered Adaptation’. The horizontal axis is at 0 mV. *B*: As in *A* but for a model of GABAergic NFS neuron. Here there is a smooth transition between silent and spiking activity, so the NFS model is type-I. *C*: As in *B* but for a model of Exc neuron. Here again the Exc model is a type-I model. *D-F*: Examples of different firing pattern of the three models, FS (*D*), NFS (*E*) and Exc (*F*) neurons, for three different intensities of the step current.

adaptation current $\eta(t)$ (black lines), one for each neuron class.

We also study the ability of the models to predict the subthreshold voltage (RMSER) as a function of the number N of exponential terms used to model the extracted spike-triggered current $\eta(t)$. Note that, as mentioned in Materials and Methods, the lower the RMSER, the better the prediction and a RMSER < 1 is possible, because we overestimate the intrinsic RMSE. Here we investigate models with static thresholds, since the dynamic threshold has no impact on the predicted voltage when the spikes are forced.

Figure 5.6 *D* shows the average spike-triggered current of the 9 GABAergic FS neurons. We observe that this spike-triggered current has two main parts, a strong and fast hyperpolarization that prevents repetitive firing during the first 40 ms, followed by a weaker but longer depolar-

Chapter 5. Parameter Extraction and Classification of Three Neuron Types Reveals two Different Adaptation Mechanisms

ization that lasts for 350 ms. This resonance is distinct from strictly subthreshold (resonating) membrane currents [Hutcheon and Yarom, 2000] since no resonance was observed in the membrane filter $\kappa(t)$. This current can be well approximated by a double exponential decay (IF₂ with $b_1 = -111.61$ pA, $\tau_1 = 36.86$ ms and $b_2 = 72.64$ pA, $\tau_2 = 61.76$ ms). The membrane potential prediction also shows that $N = 2$ is necessary and sufficient for optimal RMSER (Fig. 5.6 G).

The NFS GABAergic neurons show a simpler spike-triggered current that only mediates hyperpolarization, and which can be fitted with a single exponential function (IF₁ with amplitude and time constant $b_1 = -29.02$ pA, $\tau_1 = 34.58$ ms (figure 5.6 E, blue trace)). This current produces the relatively weak adapting behavior of the GABAergic NFS neurons, characteristic of their firing patterns [Kawaguchi and Hama, 1987]. The membrane potential prediction also shows that $N = 1$ is necessary and sufficient for optimal RMSER (Fig. 5.6 H).

Excitatory neurons have a stronger and longer adaptation current $\eta(t)$ (Fig. 5.6 F) than the GABAergic NFS and the GABAergic FS cells, which mediates the regular spiking (accommodating) behavior of these cells. Again this current is well approximated with a mono-exponential function, with $b_1 = -48.35$ pA, $\tau_1 = 44.89$ ms. Moreover we observe more variability in $\eta(t)$ across individual cells for Exc than for the two other groups. The membrane potential prediction also shows that $N = 1$ is necessary and sufficient for optimal RMSER (Fig. 5.6 I).

From these results, we conclude that the shape and dimensions of $\eta(t)$ is cell-type specific. Moreover, we observe that $\eta(t)$ in GABAergic NFS and Exc differs only by their time scale and amplitude whereas $\eta(t)$ in GABAergic NFS and Exc have a distinct shape from $\eta(t)$ in GABAergic FS. We also notice that whatever the neuron type, spike-triggered currents extend for a few hundred milliseconds, so that due to the cumulative effect, the spike-triggered current can strongly influence the spike-frequency adaptation of the neurons on long timescales. This suggests that adaptation currents might mediate some aspect of cell-type specific behavior (i.e. firing patterns).

In summary it appears that (i) the adaptation currents of different neuron classes reflect their typical firing behavior (see Figure 5.5), (ii) adaptation currents act on multiple timescales (FS: 37 and 62 ms, NFS: 35 ms, Exc: 45 ms) and (iii) NFS and Exc have strictly hyperpolarizing currents (leading to spike-frequency adaptation), but FS have both hyperpolarizing and depolarizing currents. From these results it is clear that spike-triggered currents can cause different types of spike-frequency adaptation, however some others mechanisms can also contribute such as a fatigue of action potential (AP) initiation mechanisms [Kobayashi et al., 2009, Benda et al., 2010]. This is discussed in the next section.

In the previous paragraphs, we discussed the importance of spike-triggered currents. However we know that these mechanisms are mediated by a spike-triggered change in conductance rather than current [Schwindt et al., 1988b,a]. To address this issue we fit spike-triggered conductances instead of spike-triggered currents and look at the magnitude of the improvements that follow.

In Figure 5.7 A, we observe the movement of the cumulative change in conductance following each spike η_{cond} for excitatory neurons. The $RMSE$ depends on the reversal potential E_k and

shows a minimum at $\hat{E}_k = -51.89$ mV (Figure 5.7 B).

The shape of the spike-triggered conductances η_{cond} is more difficult to interpret than the standard spike-triggered current η , because of its dependency on the reversal potential E_k . In fact, the effect of η_{cond} on voltage depends on the instantaneous difference between the actual membrane potential and the reversal potential.

Figure 5.7 C and D shows the averaged RMSER and M_d^* , for the three cell types and, for IF_η and $\text{IF}_{\eta_{\text{dyn}}}$. One can observe that conductance-based adaptation does not lead to any significant improvements in terms of M_d^* and RMSER (two sample t-test, $p > 0.2$ for M_d^* and $p > 0.07$ for RMSER). Thus we conclude that conductance- or current-based spike-triggered events make equally valid models. The choice of using one or the other depends on the preference of the modeler.

5.3.4 Essential Features for Spike Time Prediction

To explore the spike time prediction of our models, we compute M_d^* (see Materials and Methods) averaged across all neurons of a given type. For example, as mentioned in the Methods section, a value of $M_d^* = 0.8$ indicates that 80% of the predictable spikes are correctly predicted by the model. The notion of ‘predictable spikes’ takes into account that a spike time that occurs only once or twice across 10 repetitions is harder to predict than a spike that occurs in all 10 repetitions at the same moment (± 4 ms). In order to compare the effects of adaptation current η and sliding threshold γ , we show the spike time prediction for models with and without spike-triggered currents as well as with or without sliding threshold. Figure 5.8D-F, show our results for GABAergic FS, GABAergic NFS and Exc neurons, respectively. The dynamic thresholds were in each case described as a double exponential, consistent with the results of Figure 5.8 A-C (green traces).

For the GABAergic FS cells (Figure 5.8 A), we find that there is almost no movement of the threshold. The small fluctuations in the extracted moving thresholds are presumably due to noise in the estimation. This implies that all the adaptive behavior of GABAergic FS cells is mediated by the spike-triggered current and not by any changes in the AP threshold. We did not find any significant difference across models in terms of spike-time prediction ($p > 0.07$ for all pairs, with 2 sample t-test with $\alpha = 0.05$). Nevertheless when IF_0 is augmented with bin-based spike-triggered currents η , we obtain a minor gain in spike time predictions ($\Delta M_d^* = 0.05$). There is negligible increase in M_d^* when augmenting the models with sliding threshold $\Delta M_d^* = 0.007$. The optimal model for GABAergic FS neurons is $\text{IF}_\eta + \text{dyn}$ producing a M_d^* of 0.87 ± 0.06 , but the gain compared to other model variants is marginal.

The moving threshold $\gamma(t)$ of GABAergic NFS neurons follows a double exponential decay, with parameters $b_1 = 3.64$ mV, $\tau_1 = 21.88$ ms for the fast component, and $b_2 = 1.24$ mV, $\tau_2 = 336.50$ ms for the late component (5.8 B, green trace). After half a second to a second, $\gamma(t)$ is weak but by cumulating over multiple spikes the late component can contribute to spike-frequency adaptation. When comparing spike time prediction M_d^* , IF_0 performs always worse than models augmented with spike-triggered currents η or dynamic AP threshold γ . More precisely, adding a sliding threshold produces a highly significant gain in the spike

Chapter 5. Parameter Extraction and Classification of Three Neuron Types Reveals two Different Adaptation Mechanisms

prediction measure of $\Delta M_d^* = 0.10$. Adding a spike-triggered current produces a net gain of $\Delta M_d^* = 0.07$ compared to IF_0 . This leads us to the optimal model for GABAergic NFS cells being $IF_\eta + \text{dyn}$ with a value of spike time prediction $M_d^*(IF_\eta + \text{dyn})$ that indicates that more than 90% of the predictable spikes are indeed predicted by the model.

The excitatory neurons also show a double exponential decay with $b_1 = 12.45$ mV, $\tau_1 = 37.22$ ms and $b_2 = 1.98$ mV, $\tau_2 = 499.80$ ms (5.8 C, green trace) but at least twice as strong as for GABAergic NFS cells. The effect of the spikes on the threshold last for more than 1 second. The case of the excitatory neurons is special because they show a voltage reset to a value above the threshold baseline ($E_{\text{reset}} > V_T$). Therefore all models with static threshold produce repetitive firing at very high and unrealistic frequencies, which leads to a very low M_d^* , as one can observe on Figure 5.8 F. Thus models upgraded with a sliding threshold always generate a significantly higher M_d^* than similar models with static threshold ($\Delta M_d^* = 0.48$). We also observe a small increase of $\Delta M_d^* = 0.07$ for models upgraded with spike-triggered currents. These results are in agreement with the structure of the effective adaptation process of Exc neurons, where the majority of the adaptation process comes from the sliding threshold (see Figure 5.10 C). The optimal model for excitatory neurons is an $IF_\eta + \text{dyn}$ that produces $M_d^* = 0.81 \pm 0.04$. The ISI distributions of the data (Fig. 5.8 G-I) agree with the ones coming from the optimal models, but not for the simplest models without spike-triggered adaptation.

Finally we summarize results from Figure 5.6 and 5.8 by observing that the minimal optimal model for GABAergic FS cells must have two spike-triggered currents and a static threshold (IF_2), whereas for GABAergic NFS and Exc neurons the minimal model must have only one spike-triggered current, but a dynamic threshold ($IF_1 + \text{dyn}$). These minimal models are able to reproduce the same subthreshold voltage fluctuations than the intrinsic ones ($\text{RMSE} \leq 1$) and can reproduce at least 80% of the predictable spikes.

We observe that GABAergic NFS and Exc neurons have AP threshold dynamics that extend from milliseconds to more than 500 milliseconds. Moreover, due to cumulative property, the moving threshold can tune the neuron's firing frequencies, and so the PSTH, on timescales beyond 1 second. Finally AP threshold dynamics are only present in some cell types and – when present – act on very long timescales. We also note that it is the effect of γ combined with the spike-triggered current η that produce the effective adaptation behavior of a given neuron.

5.3.5 Cell-type Classification

In the last sections we showed that the membrane filter, the time course of adaptation and the AP dynamics strongly depend on the neuron type. Here we ask whether we can characterize cell types solely on their extracted parameters. We classify cell types based on (i) their passive parameters C , g_l and E_l (ii) their adaptation current η described by $\{a_k\}$ and the value of voltage reset E_{reset} (note that here we use a combination between η and E_{reset} , because the extracted spike-triggered current depends on the voltage reset, and in this view, E_{reset} is a part of the adaptation process), (iii) the shape of the dynamic threshold γ described by $\{c_p\}$, or (iv)

all the parameters. To do so we perform standard Principal Component Analysis (PCA) using (i) only the passive parameters for each cells, (ii) only the parameters of adaptation current, (iii) only the parameters of the dynamic threshold and (iv) using all the parameters.

Using any subset of parameters is sufficient for classification with the first two principal components (Figure 5.9 *B-D*) except when passive parameters only are used (Figure 5.9 *A*). Classification based on passive parameters fails because neurons do not differ in a significant way in their capacitance, their leak conductance or their reverse potential. Thus when PCA is applied on the passive parameters we are not able to distinguish between different types and representation of Exc, GABAergic NFS and GABAergic FS neurons, producing no clear cluster. However, if PCA is applied on the parameters that characterize the adaptation current and/or the dynamic threshold, we can successfully classify neurons. Moreover we also observe that the variance between the different cells are mainly explained by the reset value E_{reset} and the dynamic threshold (5.9 *D*, right panel).

5.4 Discussion

5.4.1 Automatic Fitting Method

There is a rich history of fitting neuron models to intra-cellular recordings of real neurons [Vanier and Bower, 1999, Keren et al., 2005, Achard and De Schutter, 2006, Huys et al., 2006, Jolivet et al., 2006, Kobayashi and Shinomoto, 2007, Druckmann et al., 2007, Badel et al., 2008b, Kobayashi et al., 2009]. The variety of approaches comes from the choice of neuron model and the fitting method. Still, not all methods yield models that can predict the spike times and membrane potential with high accuracy [Jolivet et al., 2008a, Gerstner and Naud, 2009]. To predict the membrane potential with Hodgkin-and-Huxley compartmental models, one needs prior knowledge on the dynamics of the ion channels present in the recorded cell [Huys et al., 2006, Druckmann et al., 2008]. Without knowledge of the ion channels, fitting Hodgkin-and-Huxley compartmental models becomes plagued with local minima, and there are often multiple parameters that share the same fitting quality [Achard and De Schutter, 2006]. The only hope for a fitting method that can easily be applied to multiple systems for which we have insufficient knowledge of the ion channel dynamics is to use convex fitting methods and IF models [Paninski et al., 2004]. Earlier work [Jolivet et al., 2006] had an efficient fitting method for the subthreshold voltage but the black-box fitting of the adaptive threshold was not convex. The method of [Badel et al., 2008b] may have been convex, but it applied only to models without spike-frequency adaptation. The method of [Paninski et al., 2004] was convex but did not use the information contained in the voltage trace while the method of [Paninski et al., 2005] was convex and used the voltage trace but lacked the moving threshold required for efficient spike prediction. In this paper we use a method that improves upon the earlier multi-linear regression method [Paninski et al., 2005] by adding a second fitting step for the moving threshold taken from the literature on generalized linear models [McCullagh and Nelder, 1998]. The method is sure to find only one set of optimal parameters because it is made of two convex fitting methods (multiple linear regression and generalized linear

Chapter 5. Parameter Extraction and Classification of Three Neuron Types Reveals two Different Adaptation Mechanisms

model with Poisson or Bernoulli probability distribution, but see constraints for the convexity in [Paninski et al., 2004]). We expect the method to generalize well to many cell types because the total time of the spike-triggered current, number and size of the basis functions is not expected to depend on the cell types.

The notion of a dynamic threshold affecting neuronal computation also has a long history [Hill, 1936, Azouz and Gray, 1999, 2000, Gerstner and Kistler, 2002, Chacron et al., 2003]. So far, the methods for fitting the dynamics of the firing threshold have relied on the measurement of the effective threshold for each spike [Azouz and Gray, 2000, Chacron et al., 2003]. Instead, the method presented in this article uses the whole voltage trace, providing information about the firing threshold each time a transient increase in the membrane potential is not followed by a spike. We expect this method to be more precise since the number of data points used to constrain the moving threshold is not proportional to the number of spikes but to the number of data points constituting the subthreshold voltage trace in the regime close to threshold.

Since the choice of model will affect the prediction performance the question arises, why we omit the non-linearity responsible for spike initiation in IF-type models [Fourcaud-Trocme et al., 2003, Badel et al., 2008b]. An exponential non-linearity in integrate-and-fire models was shown to be crucial for accurate processing of the inputs at high frequency [Fourcaud-Trocme et al., 2003]. Such a non-linear term can interact with subthreshold currents to produce a variety of firing patterns [Izhikevich, 2004, Izhikevich and Edelman, 2008, Naud et al., 2008]⁵. The IF models used here assume strictly linear voltage dynamics and are fitted away from the spikes so that the increased non-linearity close to a spike does not bias the parameter estimation. It is not trivial to generalize the present convex method to fit exponential non-linearity for spike initiation. However, in the present model the spike initiation is stochastic with an exponential non-linearity for the probability of spiking as a function of voltage as extracted from experimental data [Jolivet et al., 2006]. This exponential non-linearity should not be confused with the exponential non-linearity in the AdEx model [Brette and Gerstner, 2005] or the model of [Badel et al., 2008b]. Nevertheless, the escape-noise IF model can efficiently treat high-frequency input and reproduce all the main firing patterns (results not shown). The exponential non-linearity for spike initiation has also been shown to be dynamic and to depend on the time since the last spike [Badel et al., 2008b]. We have checked (results not shown) that a spike-triggered current can be responsible for the apparent dynamics of the exponential IV-curve observed using the method of [Badel et al., 2008b].

Another feature of some neurons that is not present in the model is the subthreshold resonance. Strong subthreshold resonance has been observed in the dendrites of large excitatory neurons [Cook et al., 2007] and in ‘mes V’ neurons [Izhikevich, 2007]. Subthreshold resonance is thought to be mediated by an additional current linearly coupled with the membrane potential. Similarly, delayed spiking upon step current injection has also been attributed to an additional current linearly coupled with the membrane potential [Naud et al., 2008]. One can check for the presence of such additional effect by ensuring that the passive membrane filter $\kappa(t)$ is a single exponential; resonance will yield a filter with a negative undershoot, while delayed onset yields a filter with a double exponential decay. If the membrane filter is not a single

⁵Chapter 2

exponential decay, it becomes crucial to fit an additional current having linear coupling with the soma. We have tested a method which involves adding the term $ae^{-t/\tau_w} * V(t)$ to Eq. 5.1. Multi-linear regression can still be applied to determine the strength of the coupling a when a time constant τ_w is assumed. Iterating through a large range of possible τ_w by repeating the multi-linear regression usually yields a convex function of the MSE as a function of τ_w . This method enables us to fit the parameters mediating subthreshold resonance or delayed spiking.

5.4.2 Effective Spike-triggered Adaptation

We have seen that the effects of spike-triggered currents and dynamic threshold merge to produce spike-frequency adaptation. The effect of $\eta(t)$ on the voltage is simply given by the convolution of η by the membrane filter κ ($\eta_v(t) = -\kappa * \eta$). The effective adaptation $\zeta(t)$ is then the combination of the threshold dynamics $\gamma(t)$ and the effect in voltage of the spike-triggered current:

$$\zeta(t) = -[\kappa * \eta](t) + \gamma(t). \quad (5.22)$$

To answer the question of whether the spike-frequency adaptation is mediated by spike-triggered currents or moving threshold, we can look at the respective contribution of $\gamma(t)$ and $\eta_v(t)$ towards the effective adaptation $\zeta(t)$. One can see in Figure 5.10 A, that the effective spike-triggered adaptation of GABAergic FS cells are clearly dominated by spike-triggered currents. For the Exc cells, $\zeta(t)$ is dominated by the moving threshold (see figure 5.10 C), whereas for the GABAergic NFS cells the effective adaptation process $\zeta(t)$ (Figure 5.10 B) is clearly a combination of η and γ , where either spike-triggered mechanism mediates approximately half of the effective spike-triggered adaptation. Therefore, we conclude that the adaptation is mediated by different processes in different cell types: mostly by the moving threshold for the Exc, entirely spike-triggered current for the FS, an equal mix of threshold and current for the NFS.

The effective adaptation for GABAergic FS cells has the particularity of crossing zero after 30-80 ms and remaining negative (*i. e.* facilitating) thereafter. The zero-crossing then determines the type of excitability: Under constant current injection, when the neuron model fires its first spike, the effective threshold is pushed away from the membrane potential and decays to the point where the effective threshold is below the membrane potential, and hence forcing a spike after a period equivalent to the first zero-crossing in the function $\zeta(t)$. One can see that under constant current injection this type of neuron will not fire at a frequency smaller than the inverse of the first-zero crossing of $\zeta(t)$: $f_c \approx 12 - 33$ Hz, in agreement with Fig. 5.5. The facilitating tail of the effective adaptation must lead to spike-frequency facilitation as observed in the firing patterns of Fig. 5.5. Moreover, the peak of the facilitating part in $\eta(t)$ will indicate a preferred frequency around 10 Hz. On the other hand, strictly decaying functions $\zeta(t)$ as in NFS and Exc will produce adapting firing patterns and type I excitability.

5.4.3 Interpretation of Model Parameters

We extract a set of parameters for a typical models of each neuron class. The parameters are related to the underlying biophysics (density, distribution and dynamics of ion channels, membrane capacitance, resistance), but the exact relationship is unclear. For instance, the extracted membrane time constants (τ_m ranged from 5 ms to 20 ms) are slightly shorter than the typical membrane time constant measured in experiments with voltage recording of the response to subthreshold step current injection at rest. This discrepancy can be explained by the fact that we measure passive parameters of neurons when they are stimulated instead of at rest. The different working regime of membrane potential will activate differently the subthreshold conductances (spike-triggered or not) and thus modify the effective membrane time constant [Richardson et al., 2003, Kobayashi et al., 2009]. The presence of the electrode may also bias our estimate of the membrane capacitance as discussed in [Badel et al., 2008b]. Similarly, the reversal potential, voltage threshold and voltage reset may depend on the bath solution, and the intra-pipette solution. The amplitude and time scale of the spike-triggered adaptation, should not be affected by the electrode or the bath solution, but the temperature at which the experiment was performed can affect the dynamics of the underlying ion channels.

The effect of the spike-triggered current on the voltage $\eta_v(t)$ is closely related to the spike-after potential (SAP; [Sah, 1996]). The spike-after potential may differ from $\eta_v(t)$ since it is measured around the resting potential while $\eta(t)$ is an average of the spike-triggered current under synaptic-like current injection. Furthermore, the spike-after potential is measured after a spike that was artificially triggered by a large and short current injection. The amount of charge that was injected to produce the spike will leak out of the membrane on a time scale given by the membrane time constant. The spike-after current extracted by standard experimental protocols [Sah, 1996] is thus biased by the current used for stimulating the spike. Yet, the close relationship between the SAP and $\eta_v(t)$ indicate that $\eta(t)$ should be mediated by the same ion channels mediating the SAP. Namely: the M-type current I_M [Adams et al., 1982], the after-hyperpolarization current I_{AHP} [Madison and Nicoll, 1984] or any other calcium-dependent ion channels [Hille, 1992, Sah, 1996, Koch, 1999, Wang et al., 2003]. Moreover, spike-triggered events in the dendrites can also shape the spike-triggered current [Doiron et al., 2007].

The movement of the threshold after a spike has been proposed to depend on sodium channel de-inactivation [Fleidervish et al., 1996, Kobayashi et al., 2009]. Following a spike, a portion of the Na^+ channels responsible for the spike initiation stay inactivated, which leads to a higher effective threshold. The sodium channels then de-inactivate which results in a gradual decay of the spiking threshold. It has been proposed that only a subtype of sodium channels are inactivating [Martina and Jonas, 1997]. Our results corroborate this hypothesis since only the GABAergic NFS and Exc types have a moving threshold, which suggest that the GABAergic FS neurons do not express the inactivating sodium channels [Martina and Jonas, 1997].

5.4.4 Classification

Classification of neuron types can be done on multiple features: firing pattern, spike shape, morphology, expression of molecular markers [Markram et al., 2004]. Here we classify based on the computational properties of the neurons, that is, the parameters regulating how the neuron encodes the incoming current into spike trains. These computational properties are determined by the expression of ion channels and lead to firing patterns that depend on the neuron type. We have shown that the classification of neuron types relates to a classification of the computational properties beyond the classification of firing patterns (GABAergic NFS and Exc cell-types are both regular spiking, and accommodating neurons). In other terms classification is possible even if the shape of the spike used traditionally for the distinction between GABAergic NFS and Exc neuron is not taken into account, since the different neuron types encode the incoming current differently. Note that the shape of the AP is not communicated to the post-synaptic neuron whereas its firing pattern is

We found that the passive properties of the neurons (capacitance, input resistance, membrane time-constant) are not sufficient to efficiently distinguish between the neurons types. It is the adaptation properties (voltage reset, spike-triggered current, moving threshold) that distinguishes the different cell types. These results would indicate that the strength and time-scale of adaptation are important parameters of cortical network computation, indicating a direction for investigating the importance of micro-circuitry in cortical networks.

We have studied only three types of neurons, but two types of excitatory neurons [Connors and Gutnick, 1990] and various types of GABAergic neurons [Markram et al., 2004] have been described. Further work is needed to check that such extensive classification can be done on the properties affecting the conversion of synaptic inputs to a spike. Classification on a greater pool of neurons would also enable to study how distinct the computational properties of different neuron types are and if a finer classification can be inferred. We expect that one would need on the order of one hundred recorded neurons under synaptic-like current injection to study in detail the classification of cortical neurons, and perform unsupervised clustering on the computational properties.

Neurons recorded *in vitro* can show very different properties from their alter ego in the intact, awake and behaving animal. In the awake animal it is not yet possible to know the input a neuron receives from its synaptic connection, but the somatic voltage can be recorded [Crochet and Petersen, 2006]. In such an experiment it is not possible to apply this part of the fitting method for the spike-triggered current and passive properties because the method requires the knowledge of the stimulating current arriving at the soma. The moving threshold, on the other hand, can be extracted since our fitting method only requires the voltage trace and the time of the spike. It remains to be tested if moving thresholds can effectively be extracted from *in vivo* recordings. It is an interesting avenue for further research since this would allow to study the correspondence of *in vitro* and *in vivo* threshold dynamics and its classification across cell types.

Chapter 5. Parameter Extraction and Classification of Three Neuron Types Reveals two Different Adaptation Mechanisms

5.4.5 How Good is Good?

Depending on the neuron type, the optimal IF model is able to predict between 81 and 91% of the predictable spikes and reproduce the subthreshold voltage fluctuations with a precision in the range of millivolts. One can then ask: Why can't we achieve a perfect prediction (i.e. $M_d^* = 100\%$ and a *RMSE* close to 0 mV)? What is missing? We can think of three possible explanations.

First, the experimental data suffers from some unavoidable drifts that are not due to the neuron itself and that we do not model. These drifts are presumably due to some additional currents that flow out of the neuron near the patch junction and so affect the recorded membrane potential in a non-systematic manner. The experimental drifts can greatly limit the maximum prediction performance.

Second, it is known that injecting current through the same electrode used for recording the voltage corrupts the recorded voltage [Brette et al., 2008]. This artifact shows as a high-frequency component of the recorded voltage that is correlated with the current being injected. Since we are fitting on the voltage trace, the artifactual component of the voltage will bias the estimated parameters. Mainly, the electrode artifact will affect the estimation of the membrane time-constant [Brette et al., 2008]. We can also speculate that the average current triggering a spike (the so-called spike-triggered averaged current, or STA) will contribute erroneously to the measured spike-triggered current. However we expect these effects to be small and to affect only very small time scales since the time-constant of the electrode contribution was measured to be below the millisecond range [Badel et al., 2008b]. Similarly, the erroneous high-frequency component of the voltage can bias our estimation of the threshold and its dynamics. The extent to which the bias in membrane time-constant and spike-triggered adaptation affects the prediction performance would have to be studied in details, but we have preliminary results showing that the effect is negligible. In any case, the artifact due to simultaneous current injection can be greatly reduced by the use of active electrode compensation methods [Brette et al., 2008, Badel et al., 2008b].

Third, by modeling neurons with simple IF models, we neglect some non-linearities that are present in the neurons. For instance, the voltage dynamics close to a spike become strongly non-linear due to the activation of sodium channels [Badel et al., 2008b]. Saturation in the open/close fraction of the ion-channels can cause the spike-triggered current of spikes in a burst to differ from the spike-triggered current of isolated spikes, leading to higher-order dependencies on the spiking history. Furthermore, we know that ion channels mediating the spike-triggered currents have time-constants which depend non-linearly on the voltage, while our formalism imposes voltage-independent time-constants. Similarly, the escape-noise formalism is only an approximation to the full dynamics entailed by stochastic activation of a limited number of voltage-dependent ion channels. For injection of synaptic-like current into the soma of a cortical neuron all these approximation prove to be very good since the spike-time prediction is high and would be even higher if experimental drifts and electrode artifacts could be completely removed.




Figure 5.6: Essential Model Features for Subthreshold Voltage Prediction in GABAergic FS, GABAergic NFS and Exc. *A:* Membrane filter κ extracted – with the Wiener-Hopf method – individually for each GABAergic FS cell (9 cells) are averaged to obtain the mean membrane filter of the GABAergic FS neurons (black) and its standard deviation (grey area). A single exponential (red) is fitted on the mean κ . A single exponential is sufficient to approximate the membrane filter and no resonance is present in any of the 9 GABAergic FS cells or the 8 GABAergic NFS (*B*) or the 9 Exc neurons (*C*). *D:* The adaptation current η extracted individually for each Fast Spiking GABAergic neuron (9 FS cells) is averaged to obtain the mean η current (black). The error bars represent one standard deviation. Note that a positive current represents a depolarizing drive and negative current a hyperpolarizing drive. A double exponential (blue) is fitted to this mean η current. Inset: zoom on the first 300 ms. *E* and *F:* same Figures with the 8 non-Fast Spiking GABAergic neurons (NFS) and the 9 Exc neurons (Exc), respectively. Note that for the NFS and the Exc, a single exponential fit (blue trace) is sufficient to approximate the mean η current. *G:* Mean RMSE for the 9 GABAergic FS cells as a function of the number of exponential function used to describe the adaptation current η (0 to 4), and for an arbitrary shaped spike-triggered current η (last column). Error bars are one standard deviation. Black stars denote a significant difference (two sample t-test). *H:* and *I:* same Figures with the the 8 non-Fast Spiking GABAergic neurons (NFS) and the 9 Exc neurons (Exc), respectively.

Chapter 5. Parameter Extraction and Classification of Three Neuron Types Reveals two Different Adaptation Mechanisms




Figure 5.7: Conductance vs current-based adaptation. *A:* Conductance-based adaptation extracted individually for each Excitatory neuron (9 Exc cells) is averaged to obtain the mean η_{cond} (black). Error bars represent one standard deviation. A double exponential (blue) is fitted on η_{cond} . Inset: zoom on the first 300 ms. *B:* $RMSE$ as a function of the reversal potential E_k of η_{cond} . The optimal $RMSE$ is obtain with \hat{E}_k . *C:* Mean RMSER for the 9 GABAergic FS, 8 GABAergic NFS and 9 Exc cells for models with current-based adaptation (IF $_{\eta}$, dark grey) and similar model with conductance-based adaptation (IF $_{\eta_{cond}}$, light grey), error bars represent one standard deviation. There are no significant difference between IF $_{\eta}$ and IF $_{\eta_{cond}}$ for a given cell type (two-samples t-test). *D:* Spike times prediction measured by the mean M_d^* for the 9 GABAergic FS, 8 GABAergic NFS and 9 Exc cells for model with current-based adaptation (IF $_{\eta}$ + dyn, dark grey) and similar model with conductance-based adaptation (IF $_{\eta_{cond}}$ + dyn, light grey), error bars represent one standard deviation. There are no significant difference between current and conductance-based models for a given cell type (two-sample t-test).




Figure 5.8: **Essential Model Features for Spike Time Prediction in GABAergic FS, GABAergic NFS and Exc.** *A*: The dynamic threshold γ extracted individually for each GABAergic FS cell (9 cells) is averaged to obtain the mean threshold dynamics for GABAergic FS neurons (black) and its standard deviation (error bars). A double exponential (green) is fitted to the mean threshold. Inset: zoom on the first 300 ms. *B* and *C*: same Figures with the the 8 non-Fast Spiking GABAergic neurons (NFS) and the 9 Exc neurons (Exc), respectively. *D*: Quality of the spike timing prediction quantified by the mean M_d^* for the 9 GABAergic FS cells as a function of the model types, static threshold (dark grey), dynamic threshold (light grey). Error bars correspond to one standard deviation. Black stars represent a significant difference between different IF models, (two-sample t-test). *E* and *F*: as in *D* but for the 8 non-Fast Spiking GABAergic neurons (NFS) and the 9 Excitatory neurons, respectively. *G*: Experimental ISI distribution (grey), ISI distribution for a model without adaptation current (IF₀, red) and ISI distribution for a model with adaptation currents and a dynamic threshold (IF_η + dyn, black). Inset: ISI density of the data as a function of the ISI density of the models (IF₀, red and IF_η + dyn, black). For perfect fit of the experimental ISI distribution all dots will lie on the diagonal (black line). *H*: As in *G* but for the 8 non-Fast Spiking GABAergic neurons (NFS). *I*: As in *G* but for the 9 Excitatory neurons (Exc).

Chapter 5. Parameter Extraction and Classification of Three Neuron Types Reveals two Different Adaptation Mechanisms




Figure 5.9: **Principal components of the 3 studied cell types.** Principal components are obtained with standard principal component analysis (PCA) for all the cells. *A*: PCA applied on the passive parameters C , g_l and E_l of the modeled neurons. Left panel: projection of all the passive parameters from all the cells (blue square: FS, green circle: NFS, red losange: Exc) onto the space defined by the two first principal components. Inset: zoom on the outlined region. Right panel: Loadings of the first and the second principle component, respectively. The same methods is applied on the adaption current η , the AP dynamics γ and all the parameters of a model $IF_\eta + \text{dyn}$, *B*, *C* and *D*, respectively. In these cases black lines are lines that separate each groups, obtained by a standard linear classifier.




Figure 5.10: **Effective spike-triggered adaptation for the three neuron classes.** As discussed in the main text the effective spike-triggered adaptation is the sum of the dynamic threshold and of the adaptation current η convolved with the membrane filter. *A*: Effective spike-triggered adaptation for each GABAergic FS cells (9 cells) are averaged to obtain the mean effective adaptation (black) and its standard deviation (grey shaded area). Mean threshold dynamics (green) and mean η_v (blue). Inset: zoom on the first 300 ms. *B* and *C*: same Figures with the the 8 non-Fast Spiking GABAergic neurons (NFS) and the 9 Exc neurons (Exc), respectively.

6 Coding and Decoding with Adapting Neurons: a Population Approach

Theoretical explorations of adaptation in spiking neurons often fall flat due to a dependence on the infinitely many spiking histories possible. Basic observables such as the Peri-Stimulus Time-Histogram (PSTH) and linear filter could not be expressed as an explicit function of the model's parameters. In the present article, we use an approximation of the full dynamics which becomes exact in the limit of small population activity and/or weak refractoriness. This approximation allows us to derive an expression for encoding and decoding time-dependent stimulus in the population activity. In like manner, we derive an expression for the linear filter which shows how high-pass and band-pass properties can arise from distinct shapes of the spike after-potential. In all cases the approximation matches very well with direct simulations of large neuronal populations. An analytical expression can shed light onto previously obscure processes. Here we discover that the decoding of a population of weakly active neurons only requires two quantities: i) the instantaneous population activity and ii) an accumulation of the past history weighted by a factor that relates to the effective spike after-potential. The results presented here can be used to make mean-field theory models of neuron networks closer to experimental observations.

6.1 Introduction

How information can be encoded and decoded with populations of neurons is a fundamental question of computational neuroscience [Dayan and Abbott, 2001, Gerstner and Kistler, 2002]. A time-varying stimulus can be encoded in the active fraction of a population of neurons, a coding procedure that we will refer to as population coding. Such population coding is likely to be used at least to some degree by the nervous system [Abeles, 1991, Averbeck et al., 2006]. Given the need for fast processing of information by the brain [Thorpe et al., 1996], population coding is an efficient way to average information in noisy and unreliable neurons. Then, by adapting to the stimulus, population coding is not only fast but also energy-efficient [Laughlin and Sejnowski, 2003].

When one is presented with a constant stimuli, the population of neurons coding for this

Chapter 6. Coding and Decoding with Adapting Neurons: a Population Approach

stimulus responds first strongly but then adapts to the stimulus. Such an adapting population code is said to be ambiguous because adapting responses cannot be directly related to the instantaneous estimate of the population activity[Fairhall et al., 2001]. The fact that our perception of the stimuli does not change indicates that the adapting response can be efficiently decoded by the appropriate brain areas. But according to what rule can population activity be decoded? What elements of the population history are relevant? What are the basic principles?

Theoretical approaches to adapting neuron ensembles describe the dynamics in terms of non-linear, delayed partial differential equations[Muller et al., 2007, Richardson, 2009, Farkhooi et al., 2011], or in terms of a system of ordinary differential equations[Toyozumi et al., 2009]. These approaches are significant improvements to the computationally expensive approach of simulating individually each neuron of the population. Other approaches assume that neurons are mainly refractory and these approaches lead to self-consistent equations for the population activity such as in Wilson and Cowan [1972] or Gerstner [1995, 2000]. The refractory period can also be set aside so as to apply mean-field methods. Such mean-fields methods showed fascinating properties of neuronal networks: chaotic dynamics[Sompolinsky et al., 1988, van Vreeswijk and Sompolinsky, 1998], attractors [Amit and Tsodyks, 1991, Amit and Brunel, 1997, Brunel, 2000, Renart et al., 2010] and computations [Amit and Brunel, 1997, Mayor and Gerstner, 2003, Bertschinger and Natschläger, 2004]. But taking adaptation into account is essential to relate with experimental observations[Pillow et al., 2005, Jolivet et al., 2006, 2008a].

Spike-frequency adaptation appears in practically all neuron types of the nervous system[Benda and Herz, 2003]. Biophysical processes that can mediate spike-frequency adaptation include spike-triggered activation or inactivation of ion-channels[Storm, 1987, Schwindt et al., 1988b,a] and a spike-triggered increase in the firing threshold[Hill, 1936, Fuortes and Mantegazzini, 1962, Azouz and Gray, 2000, Mensi et al., Under Review]¹. Neurons adapt a little bit more each time they emit a spike, and it is the trailing effect of every previous spike that sets the level of adaptation [Kobayashi et al., 2009, Mensi et al., Under Review]². The effect of a single spike on future spiking probability cannot be summarized by a single time constant. Rather, the spike-triggered adaptation unfolds on multiple time scales and varies strongly across cell-types. Relating spiking dynamics to the population dynamics must therefore be easily generalizable to the different profiles of spike-triggered adaptation observed in experiments.

In the linear regime, spike-frequency adaptation can introduce a high-pass filter[Benda et al., 2005], but this may not be the case for all cell-types. Experimental studies showed that the properties of the linear system can vary across cell-types and experimental conditions [Pike et al., 2000, Arsiero et al., 2007, Lundstrom et al., 2008]. Theoretical relations have been drawn between the linear filter and properties of the relative refractory period[Gerstner, 2000], the spike-initiation[Fourcaud-Trocme et al., 2003] and subthreshold resonances[Richardson et al., 2003] and to simple adaptation dynamics [Shriki et al., 2003].

¹Chapter 5

²*Idem*

To address the problem of coding in adapting neurons, we used approximation methods that are valid at low population-activity and/or small amplitude of effective spike after-potentials. This approximation is justified by the low firing rates observed in cortical neurons *in vivo* [de Kock and Sakmann, 2009, Crochet et al., 2011] and by the small after-potential amplitudes associated with the long time constants in cortical neurons [Mensi et al., Under Review]³. It is also applicable when neurons are embedded in a large, weakly connected, homogeneous network, *i.e.* when the assumptions of mean-field theory apply, as well as in the linear regime.

The paper is organized as follows: first we derive the implicit equation which describes the population activity as a function of the input and compare the theory with simulations. Then we derive an equation for decoding the population activity. We show with simulations that this decoding scheme can resolve the ambiguity of adapting responses by taking into account a weighted function of the past activity. Finally, we derive the linear filter as a function of the effective spike after-potential. Our results help to understand how input is dynamically encoded in neuronal populations and how different cell-types give rise to different population dynamics.

6.2 Results

To keep the discussion transparent, we focus on a population of unconnected neurons. Our results can be generalized to coupled or weakly coupled populations using standard theoretical methods [Brunel, 2000, Gerstner and Kistler, 2002].

6.2.1 Encoding the Population Activity

How does a population of unconnected and adapting neurons encode a given stimulating current $I(t)$? The population will respond with a set of spike trains $\{S(t)\}$, where each spike train is written mathematically as a sum of Dirac-impulses: $S(t) = \sum_{i=1}^n \delta(t - \hat{t}_i)$, where n is the total number of spikes in the train and \hat{t} denotes the time of a spike. The set of spike trains can be characterized by the time-dependent population activity $A(t)$. The population activity is defined as the expected proportion of active neurons within an infinitesimal time interval. It corresponds to the limit at large population sizes and small time interval of the number of active neurons $n_{\text{act}}(t, t + \Delta t)$ in the time interval $[t, t + \Delta t]$ divided by the total number of neurons N and the time interval Δt [Gerstner and Kistler, 2002]:

$$A(t) = \lim_{t \rightarrow 0, N \rightarrow \infty} \frac{n_{\text{act}}(t; t + \Delta t)}{N \Delta t} = \langle S(t) \rangle. \quad (6.1)$$

The angular brackets $\langle \cdot \rangle$ denote the expected value on an ensemble of identical neurons. It is equivalent to consider the expected value as averaging in time over independent repeated

³Chapter 5




Figure 6.1: **Encoding the stimulus in adapting populations.** A-C: Population activity responses (top panels; PSTH from 25 000 repeated simulations in blue, renewal theory in dashed black, first order of Eq. 6.6 in dotted black, second order in full black) to the step current input (bottom panels; black). The effective spike after-potential used was $\eta(t) = \eta_0 e^{-t/\tau_\eta}$. The model parameters were :A $\eta_0 = 1$ and $\tau_\eta = 500$ ms, B $\eta_0 = 5$ and $\tau_\eta = 500$ ms with small input step, C $\eta_0 = 5$ and $\tau_\eta = 700$ ms with large input step. The average first and average last interspike intervals were A: 170 and 342 ms, B: , C: 577 and 792 ms. D: population activity responses to an Ornstein-Uhlenbeck process with mean 20 pA, standard deviation 25 pA and correlation time constant equal to 300 ms with model parameters $\eta_0 = 1$ and $\tau_\eta = 500$ ms. The colour scheme is the same as for A-C. In all cases, the membrane filter was $\kappa(t) = \kappa_0 e^{-t/\tau}$ for $t > 0$ and zero otherwise with $\tau = 10$ ms, $\kappa_0 = 0.01$, the scaling factor was $\lambda_0 = \log(-10)$ and the PSTHs were averaged on 4 ms time bins.

presentations of a stimulus in the same single neuron. In this sense, the population activity is the time-dependent firing intensity known to experimentalists as the Peri-Stimulus Time Histogram (PSTH).

The population activity is different from the conditional firing intensity $\lambda(t|I, S)$. The condi-

tional firing intensity depends on the input $I(t)$ and on the spiking history, or past spike train, S . It is implied in the notation that the conditional firing intensity at time t depends on the spike train only up to time t . We can see $A(t)$ as the expected value of the conditional firing intensity across all possible spiking histories (see Appendix 6.5):

$$A(t) = \langle \lambda(t|I, S) \rangle. \quad (6.2)$$

An accurate model of the conditional firing intensity would depend on the whole history of spiking. We use the Spike-Response Model⁴ with escape noise (SRM; [Plesser and Gerstner, 2000a, Gerstner, 2000, 2008]) which is also referred to as the generalized linear model for spiking neurons [Paninski, 2004]. The validity of this model has been repeatedly verified in various neuron types and various experimental protocols [Pillow et al., 2005, Jolivet et al., 2006, Pillow et al., 2008, Mensi et al., Under Review]⁵. The SRM generates spikes according to the filtered input and the effect of each previous spike:

$$\lambda(t|I, S_t) = \lambda_0 f([\kappa * I](t) + [\eta * S](t)), \quad (6.3)$$

where ‘*’ denotes the convolution, $\eta(t)$ encodes the effect of each spike on the ulterior probability of spiking, λ_0 is a scaling constant related to the reversal potential with units of inverse time and $\kappa = \kappa_0 e^{-t/\tau}$ for $t > 0$ and zero otherwise is the membrane filter with amplitude κ_0 and time-constant τ . The link function f can take different shapes depending on the underlying noise process. Here we will use an exponential link function since it was shown to match with the noisy adaptive-exponential-integrate-and-fire [Mensi et al., 2011] and experimental data [Jolivet et al., 2006, Pillow et al., 2008, Mensi et al., Under Review]⁶.

The effect of each spike on the ulterior spiking probability ($\eta(t)$) can be mediated by a dynamic threshold as well as a spike-triggered current or spike after-potential [Mensi et al., Under Review]⁷. Throughout the remainder of the text we will refer to $\eta(t)$ as the effective spike after-potential. We will use this nomenclature to preserve the connections with experimental quantities, but it is important to clarify that $\eta(t)$ has no units because it relates to a scaled version of the spike after-potential. Another designation could be the spike-triggered adaptation, but the effective spike after-potential can be depolarizing which yields a positive feedback attributable to facilitation rather than adaptation. A depolarizing (facilitating) spike after-potential is associated with $\eta(t) > 0$, while a hyperpolarizing (adapting) spike after-potential is associated with $\eta(t) < 0$. The effect of a spike on future time must decrease such that $\lim_{t \rightarrow \infty} \eta(t) = 0$. We also require causality in the SRM by having $\eta(t) = 0$ for $t \in (-\infty, 0]$. The dependence of the conditional firing intensity upon the spiking history is represented by the convolution of the spike train with $\eta(t)$. Since the spike train is defined as a sum of delayed delta-functions, the

⁴Readers having read Chapter 1, 4 or 5 should be familiar with the SRM and can go directly to Eq. 6.4 since the notation is preserved.

⁵Chapter 5

⁶*Idem*

⁷*Idem*

Chapter 6. Coding and Decoding with Adapting Neurons: a Population Approach

convolution boils down to a sum of delayed spike after-potentials: $[\eta * S](t) = \sum_{i=1}^n \eta(t - \hat{t}_i)$. In the renewal scenario, one would replace the sum by the contribution from the last spike in Eq. 6.3: $\eta * S \rightarrow \eta(t - \hat{t}_{\text{last}})$.

To find a closed-form equation of the population activity using Eq. 6.2, we take the expected value of the SRM's firing intensity, Eq. 6.3. The exponential link-function allows us to separate Eq. 6.3 into two factors:

$$A(t) = A_I(t) \langle e^{\eta * S} \rangle. \quad (6.4)$$

The first factor forms the time-dependent population activity of neurons without adaptation or refractory effects as in the Linear-Nonlinear Poisson model (LNP):

$$A_I(t) = \lambda_0 \exp([\kappa * I](t)). \quad (6.5)$$

In the second factor, we can recognize the moment generating functional. This functional can be written in terms of distribution functions of the spike trains such as $\langle S(t)S(t') \rangle$, or in terms of the correlation functions such as $\langle [S(t_1) - A(t_1)][S(t_2) - A(t_2)] \rangle$. Both expansions are known [van Kampen, 1992]. We use the expansion in correlation functions for reasons that will be made clear shortly. The correlation functions are labeled $g_n(t_1, t_2, \dots, t_n)$ as in van Kampen (1992) such that the first correlation function is the population activity: $g_1(t) \equiv A(t)$, the second correlation function is $g_2(t_1, t_2) = \langle [S(t_1) - A(t_1)][S(t_2) - A(t_2)] \rangle$ for $t_1 \neq t_2$, and so on. Then, the generating functional can be written [van Kampen, 1992]:

$$\langle e^{\eta * S} \rangle = \exp \left(\sum_{m=1}^{\infty} \frac{1}{m!} \int_{-\infty}^t (e^{\eta(t-s_1)} - 1) \dots (e^{\eta(t-s_m)} - 1) g_m(s_1, \dots, s_m) ds_1 \dots ds_m \right). \quad (6.6)$$

Eq. 6.6 is a generating functional because the functional derivatives evaluated at $\eta(t) = 0$ gives back the correlation functions.

It is a physical rule of thumb that the g_m decrease rapidly in magnitude when the spikes are almost independent [van Kampen, 1992]. This is the main reason for using the expansion in terms of correlation functions instead of distribution functions. Truncating at $m = 0$ would yield $A(t) = A_I(t)$, the population activity of an inhomogeneous Poisson model, or LNP. Truncating at $m = 1$ gives the first-order equation:

$$A(t) = A_I(t) \exp \left(\int_{-\infty}^t (e^{\eta(t-s_1)} - 1) A(s) ds \right). \quad (6.7)$$

Truncating the series to the first order gives a self-consistent equation for the population activity which can be solved using standard numerical methods (see Materials and Methods).

All the higher orders contain correlations functions for which we do not know the exact analytical expressions. We will therefore approximate the second term of the series with the

correlation function of a renewal SRM receiving no input $C_0(t)$ [Gerstner and Kistler, 2002]:

$$g_2(t_1, t_2) \approx A(t_1)C_0(t_2 - t_1). \quad (6.8)$$

The correlation function is made time-dependent by scaling C_0 with the instantaneous population activity at time t_1 . There are three aspects to this approximation: i) the correlation function C_0 is calculated from renewal theory [Gerstner and Kistler, 2002] and does not take into account all the spikes in the spiking history, ii) the time-dependence is implemented by $A(t_1)$, the complete time-dependence should also take into account $A(t_2)$, iii) C_0 is calculated at zero input current instead of the stationary input, this approximation is made better by the fact that adaptation dynamically discounts from the time-dependent current such that the net input is often close to zero. Overall, we expect that the approximation (Eq. 6.8) will be more suitable for slow fluctuations in the input, and that fast fluctuations will cause greater error. Eq. 6.8 is chosen such that we can still write the population activity with a single integral. Indeed, one of the integrals of the second term in the infinite series of Eq. 6.6 can be evaluated independently of $A(t)$. We find:

$$A(t) = A_I(t) \exp \left[\int_{-\infty}^t \left(1 + \frac{1}{2} \alpha(s) \right) (e^{\eta(t-s)} - 1) A(s) ds \right], \quad (6.9)$$

where $\alpha(s) = ((e^\eta - 1) * C_0)[s]$ is the convolution between the stationary correlation $C(t)$ and the factor $e^{\eta(t)}$ - 1. Eq. 6.9 is the second-order approximation to the population activity. We expect that Eq. 6.9 will be valid in the regime of low interactions between spikes. Since the level of interaction between spikes is regulated by the strength of the spike after-potential and the level of activity, we expect that Eq. 6.9 will be valid for low $A(t)$ and/or low $\eta(t)$.

Eq. 6.9 gives a good approximation to the population activity as it can be seen by comparing with direct simulation of 25 000 neurons (Fig. 6.1). This approximation can be compared to the first-order (Eq. 6.7) and to the renewal approximation (see Material and Methods). The second-order equation (Eq. 6.9) is always better than the first-order equation (Eq. 6.7). For small spike after-potentials (Fig. 6.1 A and D) the second-order equation is almost perfect. When the spike-after potential is made large (Fig. 6.1 B and C), the interaction between the spikes increases and we start to see the effect of the semi-stationary correlation function (Eq. 6.8) and of the higher order terms in Eq. 6.6. In Fig. 6.1 C the second-order equation does not capture the full complexity of the oscillating response to a step change in input. Eq. 6.9 remains valid at the steady state, indicating that the error comes primarily from the semi-stationary correlation function (Eq. 6.8). Reducing the level of activity should reduce the level of interaction between spikes and thus improve the validity of Eq. 6.9 for large spike after-potential. This is indeed what we observe by comparing with simulations (Fig. 6.1 B).

We also compare with the renewal approximation whereby the SRM is assumed to depend on the last spike only. The renewal approximation gives systematically over-estimated population activities (Fig. 6.1 A and B). Taking into account only the last spike neglects a significant amount of self-inhibition which explains the over-estimate. When the spike after-potential

is high (Fig. 6.1 C), the population activity is dominated by refractory effects at the onset of a step (dashed line matching blue line in Fig. 6.1 C). At the end of the step, when noise has de-synchronized the activity, the averaged effect of all previous spikes is still important and it is controlling the steady-state population activity (dashed line over-estimating blue line in Fig. 6.1 C).

6.2.2 Decoding the Population Activity

Insightful understanding of Eq. 6.9 is facilitated if we inverse the encoding question and ask: what stimulus corresponds to a given population activity? The encoding equation (Eq. 6.9) can be inverted to give the decoding formula:

$$[\kappa * I](t) = \log(\lambda_0^{-1} A(t)) - \int_{-\infty}^t \left(1 + \frac{1}{2} \alpha(s)\right) (1 - e^{\eta(t-s)}) A(s) ds. \quad (6.10)$$

Eq. 6.10 reveals that to recover the input, the logarithm of the population activity must be reduced by a accumulation of the past activity. The presence of the logarithm reflects the non-linearity for encoding (the link-function in Eq. 4.53). When the instantaneous population activity is zero, the stimulus is undefined. This is synonymous to the more familiar statement that a stimulus outside of the dynamic range cannot be decoded. When the population activity is low, there is a greater variability in the decoded stimulus because small errors in the population activity are amplified by the logarithm in Eq. 6.9.

Figure 6.2 shows the ability of Eq. 6.10 to recover the input from the population activity of 25 000 model neurons. As might be expected, the decoding scheme (Eq. 6.10) fails where the encoding is badly approximated by Eq. 6.9 (Fig. 6.2 C). We conclude that Eq. 6.10 is a valid decoder in the domain of applicability of Eq. 6.9.

6.2.3 How the Effective Spike After-Potential Shapes the Transfer Function

In order to analyze the input-ouput function of the adapting population, we consider a stimulus making small fluctuations $\Delta I(t)$ around its mean I_0 :

$$I(t) = I_0 + \Delta I(t). \quad (6.11)$$

This produces a population activity making small fluctuations $\Delta A(t)$ around its mean A_0 :

$$A(t) = A_0 + \Delta A(t) \quad (6.12)$$

The steady-state adaptation according to Eq. 6.7 is given by the recursive equation:

$$A_0 = \lambda_0 \exp \left(I_0 \int_0^\infty \kappa(s) ds + A_0 \int_{-\infty}^t (1 - e^{\eta(t-s)}) ds \right) \quad (6.13)$$




Figure 6.2: **Decoding the stimulus from the population activity.** From the same traces shown in Fig. 6.1, we show the original (bottom panels, black line) and decoded stimulus (bottom panels $\kappa * I$; red line; arbitrary units) recovered from the PSTH of 25 000 independent SRM neurons (top panels; blue line) using Eq. 6.10. The decoded waveform of negative input is often not defined because it entails an activity of zero and the logarithm of zero in Eq. 6.10 is not defined.




Figure 6.3: **Validity of the analytical transfer function.** The magnitude (A) and complex phase (B) of the analytical transfer function (from Eq. 6.14) is plotted as a function of frequency in black. In blue, the transfer function computed from fitting a sine on the PSTHs of Eq. 4.53 with a small sinusoidal current (see Materials and Methods). The parameter used were $\eta(t) = \eta_0 e^{-t/\tau_\eta}$ with $\eta_0 = 1$ and $\tau_\eta = 400$ ms, $\lambda_0 = \log(-10)$, $I_0 = 160$ pA, $\max(\Delta I) = 3$ pA, $\kappa_0 = 0.01$, $\tau = 10$ ms. C-D: The current fluctuations were increased to $\Delta I = 12$ pA to produce the amplitude gain and phase shift. In all four panels, the error-bars correspond to two standard deviations.

Using Eq. 6.7 and ignoring terms of second order in the perturbations, we find an expression for the transfer function:

$$\hat{H}(\omega) = \frac{A_0 \mathcal{F}\{\kappa\}(\omega)}{1 + A_0 \mathcal{F}\{1 - e^\eta\}(\omega)}, \quad (6.14)$$

where $\mathcal{F}\{\cdot\}$ denotes the Fourier transform and ω is the frequency. The transfer function is the filter that relates the input fluctuations to the output fluctuations in Fourier space: $\mathcal{F}\{\Delta A\}(\omega) = \hat{H}(\omega) \mathcal{F}\{\Delta I\}(\omega)$. The transfer function of pyramidal neurons has been analyzed experimentally by measuring the averaged response to small sinusoidal current [Kondgen et al., 2008, Lundstrom et al., 2008]. Within such an experimental framework, the magnitude of the transfer function $|\hat{H}(\omega)|$ is the amplitude gain relating the amplitude of the stimulating current to the amplitude of the measured PSTH. The phase shift between the input sinusoidal and the output sinusoidal is the arc-tangent of the imaginary part of $\hat{H}(\omega)$ over its real part.

The theoretical transfer function (Eq. 6.14) provides a good fit to numerical simulations of a population of neurons firing according to Eq. 6.3 (Fig. 6.3). The phase shift predicted from Eq. 6.14 reproduces with high precision the phase shift observed in the response of a population of SRM neurons (Fig. 6.3B; see also Materials and Methods). The amplitude gain is also very well captured by the theory. The consistent over-estimation of the amplitude gain increases with the current amplitude (compare Fig. 6.3A with Fig. 6.3C). Therefore, the mismatch is expected to vanish at very small current amplitudes. A similar argument explains the small under-estimation of the phase shift in Fig. 6.3B and Fig. 6.3B. We conclude that the theoretical expression for the transfer function (Eq. 6.14) is a valid description at small current fluctuations and small population activity fluctuations.

It is now possible to study how the functional shape of the effective spike after-potential $\eta(t)$ influences the transfer function $\hat{H}(\omega)$. In simple models, the spike-triggered adaptation is often assumed to be a single exponential: $\eta(t) = \eta_0 e^{-t/\tau_\eta}$. The single exponential $\eta(t)$ entails a high-pass filter by attenuating the low frequencies. The amplitude of the low frequency attenuation depends on both the strength of the spike after-potential η_0 and its time-scale τ_η (Fig. 6.4 A, B, D and E). The critical frequency of the attenuation is set mostly by the time-scale of the spike after-potential (Fig. 6.4 D-E). The time scale, τ_η , also heavily influences the frequency of the maximal phase shift. The amplitude, η_0 , tunes the maximal phase shift without changing the frequency to which this maximum corresponds (Fig. 6.4 C and F).

The theory presented here allows us to explore shapes of the spike after-potential beyond the single exponential spike-triggered adaptation (Figure 6.4). As an proof of principles, we consider a spike after-potential made of a sum of two exponential. A sum of two exponentials separated in two types of qualitatively different scenarios: i) $\eta(t)$ is said to be over-damped when the amplitude corresponding to each exponential are both negative and thus $\eta(t) < 0$ everywhere. Such spike after-potentials correspond to regular spiking neurons either pyramidal or GABAergic [Mensi et al., Under Review]⁸. The effect of over-damped $\eta(t)$ yield transfer

⁸Chapter 5




Figure 6.4: **The relation between the shape of the spike-triggered adaptation and the transfer function.** Different shapes of the effective spike after-potential, $\eta(t)$, are plotted in A, D, G. The resulting filters are displayed with the amplitude gain (B, E, H) and phase spectrum (C, F, I). The top row shows the effect of modifying the amplitude of spike-triggered adaptation decaying with a single exponential: $\eta_0 = 5$ mV (black lines), 10 mV (blue lines) and 20 mV (red lines). The middle row shows the effect of changing the time constant for the spike-triggered adaptation decaying with a single exponential having a fixed amplitude: $\tau_\eta = 50$ ms (black lines), 100 ms (blue lines) and 200 ms (red lines). The bottom row shows the effect of changing the resonance frequency for resonating $\eta(t)$: $\omega_\eta = 0.2$ Hz (black lines), 1 Hz (blue lines) and 2 Hz (red lines).

functions qualitatively similar to those presented for the single exponential. ii) $\eta(t)$ is said to be resonating (or under-damped) when one of the exponentials has a positive amplitude which often lead to $\eta(t) > 0$ for a restricted range of time. Resonating spike after potentials are characterized in terms of their resonating frequency ω_η , and amplitude η_0 such that: $\eta(t) = \eta_0 e^{-t/\tau_\eta} \cos(2\pi\omega_\eta t)$. The resonance is responsible for the strongly selective band-pass of the transfer function (Fig. 6.4 A-C). The frequency that is selected by the band pass depends on the resonating frequency ω_η of the effective spike after-potential $\eta(t)$. Resonating combination of exponential were observed in the spike-after currents of fast-spiking GABAergic neurons of the cortex[Mensi et al., Under Review]⁹.

6.3 Discussion

We have derived a self-consistent formula for the population activity of independent neurons (Eq. 6.9). The equation is valid at low coupling between spikes which can be observed in real neurons whenever i) the firing rate is low, ii) the spike-after potentials have small amplitudes or iii) both the firing rate is low and the spike after-potentials have small amplitudes. Recordings from cortical neurons [Mensi et al., Under Review]¹⁰ show that the spike after-potential is made of short and long time constants. The short (refractory) time constants are associated with larger after-potential amplitudes (τ_η 80 ms has η_0 10) while long (adaptation) time constants are associated with smaller after-potential amplitudes (τ_η 400 ms has η_0 1.5). The population activity of such neurons should therefore contain a fast transient predicted by renewal theory [Gerstner, 2000, Gerstner and Kistler, 2002] and a dynamical response on longer time scales predicted by Eq. 6.9.

The decoding scheme presented in this paper (Eq. 6.10) reveals a fundamental aspect of population coding with adapting neurons. Namely, the ambiguity introduced by the adaptation can be resolved by considering a well-tuned accumulator of past activity. The neural code of adapting populations is ambiguous because it is possible to observe a range of population activities in response to a given instantaneous input (see Fig. 6.1). We have shown that resolving the ambiguity requires the knowledge of the activity but to a good approximation does not require knowing which neuron was active. At high population activity with large spike after-potentials, however, characteristics of the individual spike trains such as the inter-spike interval distributions may be required to resolve the ambiguity (as proposed by Fairhall *et al.* [Fairhall et al., 2001]). Indeed, this can be deduced from Eq. 6.6 where the second-order term relate to the auto-correlation function, a quantity closely related to the inter-spike interval distribution. The results suggest a strictly post-synaptic decoding, although presynaptic decoding such as decoding synapses[Sen et al., 1996, Pfister et al., 2010] is not incompatible and may be required to decode large fluctuations in the population activity or fast transients where renewal theory rules. Plausible biophysical processes for post-synaptic decoding include lateral inhibition[Carandini and Heeger, 1994] and non-linear dendrites[Polsky et al., 2009].

⁹ *Idem*

¹⁰ *Idem*

Our analytical expression states that the effective shape of the spike after-potential defines the weighted sum of the past activity. Therefore, one prediction of the present study is that the decoding mechanisms should involve an estimation of the past activity matching the type of cells that are decoded.

In the small input regime, we were able to derive an equation for the linear filter corresponding to a population of adapting neurons. Eq. 6.14 makes explicit the relation between the effective shape of the spike after-potential and the transfer function of signal processing theory. The effective spike after-potential appears as an important mechanism for shaping the transfer function of low input frequencies. As implied an earlier work[Benda et al., 2005] hyperpolarizing spike after-potentials introduces a high-pass in the amplitude gain. This high-pass in the amplitude gain along with the dependence of the transfer function on the mean population activity is consistent with experimental measurements in cortical neurons[Arsiero et al., 2007]. The amplitude gain can become a sharp bandpass for resonating spike after-potentials. Phase advance is possible when the spike after-potential has a depolarizing effect, possibly explaining the basic separation between lagged and non-lagged cells[Dong and Atick, 1995]. Tuning the transfer function may be crucial for optimal coding in neurons. For instance the cells may tune their transfer function to remove correlations from the stimulus[Srinivasan et al., 1982], to counter-balance the noise spectrum[Chacron et al., 2005b] or to achieve fractional differentiation[Lundstrom et al., 2008]. Eq. 6.14 makes explicit predictions of what type of adaptation dynamics can give a particular transfer function.

The scope of the present investigation was restricted to unconnected neurons. It is straightforward to extend the results to population of weakly connected neurons [Gerstner, 2000]. For instance, a population made of inter-connected neurons of M cell-types would correspond to the self-consistent system of equation:

$$A_k(t) = \lambda_0 \exp \left[(\kappa * I)[t] + \int_{-\infty}^t \left(1 + \frac{1}{2} \alpha(s, \eta_k) \right) (e^{\eta_k(t-s)} - 1) A_k(s) ds \right] \left(1 + \sum_{m=1}^M (\epsilon_{mk} * A_m)[t] \right), \quad (6.15)$$

where ϵ_{mk} is the scaled post-synaptic potential kernel from cell-type m to cell-type k (following the formalism of Gerstner and Kistler [Gerstner and Kistler, 2002]), each subpopulation is characterized by its population activity $A_k(t)$ and its specific spike after potential $\eta_k(t)$. The linear-filter results can also be readily be extended to connected neuron networks following Spiridon and Gerstner [Spiridon and Gerstner, 1999]. The problem of coding and decoding, however, becomes intricate with the correlations introduced by synaptic connection and possible correlations in the signals received by different subpopulations [Averbeck et al., 2006]. Since the spike-after potential is one of the most important parameter for distinguishing between cell classes [Mensi et al., Under Review]¹¹, the approach presented here allows one to build simple network models that take into account the neuronal cell-types beyond the sign of the synaptic connection.

¹¹ *Idem*

6.4 Materials and Methods

All simulations were performed on a desktop computer with 4 cores (Intel Core i7, 2.6 GHz, 24 GB RAM) using Matlab (The Mathworks, Natwick, MA). Eq. 6.3 was simulated by first discretizing time ($dt = 0.05$ ms) and then deciding at each time step whether a spike is emitted by comparing the probability to spike in a time-bin ($1 - e^{-\lambda(t)dt}$) to a random number of uniform distribution. The transfer function was estimated by injecting a very long sinusoidal stimulus and binning the spike times according to their phase. We stopped the calculation of the PSTH for Fig. 6.3 after 25 000 spikes. The sinusoidal PSTH thus generated was fitted with a sine using the native Matlab function *nlinfit* to extract the amplitude, phase, baseline and their estimation error. The average of the fitted baseline was then used as A_0 in Eq. 6.14.

The self consistent equations were solved by the Gauss-Newton algorithm for the discretized population activity. First we assume that there were no activity before $t = 0$ so that the lower bound in the integral equations is zero. Then, we start with the initial guess given by $A_I(t)$ and the iteratively update the estimate of $A(t)$ with the difference between the right and left hand side of the self-consistent equation scaled by the current estimate of $A(t)$. Convergence required 100 to 5000 iterations (1 to 60 seconds of simulation for 10 seconds of model time) which can be made numerically efficient by the use of FFT-based convolutions for long time-series. This is a simple but crude algorithm and time-efficiency could be gained by using more involved algorithms for solving integral equations [Press et al., 2007].

To compute the second order equation (Eq. 6.9) we need to find $C_0(\tau)$. For a neuron population at a steady state, the stationary correlation function $C(\tau) = A^{-1} \langle [S(t) - A][S(t - \tau) - A] \rangle$ is given by [Gerstner and Kistler, 2002]:

$$C(\tau) = \mathcal{F}^{-1} \left\{ \Re \left\{ \frac{1 + \hat{P}_0(\omega)}{1 - \hat{P}_0(\omega)} \right\} \right\}, \quad (6.16)$$

where \Re is the real part, \mathcal{F}^{-1} is the inverse Fourier transform and $\hat{P}_0(\omega)$ is the Fourier transformed interspike interval distribution given by [Gerstner and Kistler, 2002]:

$$P_0(\tau) = \lambda_0 \exp \left(\eta(\tau) - \int_0^\tau \lambda_0 e^{\eta(x)} dx \right). \quad (6.17)$$

6.5 Appendix: Derivation of Eq. 6.2

The probability of the train of n spikes S_n is given by [Pfister et al., 2006]:

$$P(S_n) = \prod_{\hat{t}_i \in S_n} \lambda(\hat{t}_i | S_n) e^{-\int \lambda(x | S_n) dx}, \quad (6.18)$$

where we omit writing the dependence on the input I for notational convenience. By definition, the population activity is the expected value of a spike train: $A(t) = \int S(t) P(S) DS$. Following van Kampen [van Kampen, 1992] we can integrate over all possible spike times in

Chapter 6. Coding and Decoding with Adapting Neurons: a Population Approach

an ordered or non-ordered fashion. The non-ordered approach integrates the spike times over all possible times (here the interval $(-\infty, t]$). Additionally the non-ordered integral introduces a $1/n!$ factor due to the permutation of identical events. We obtain:

$$A(t) = \sum_{n=1}^{\infty} \frac{1}{n!} \int \lambda(\hat{t}_1|\hat{t}_1, \dots, \hat{t}_n) \dots \lambda(\hat{t}_n|\hat{t}_1, \dots, \hat{t}_n) e^{-\int \lambda(x|\hat{t}_1, \dots, \hat{t}_n) dx} \sum_{j=1}^n \delta(t - \hat{t}_j) d\hat{t}_1 \dots d\hat{t}_n, \quad (6.19)$$

where the term $n = 0$ has been eliminated by the fact that $S_{n=0} = 0$. The sum of Dirac-delta functions entails a sum of n identical integrals:

$$A(t) = \sum_{n=1}^{\infty} \frac{n}{n!} \int \lambda(\hat{t}_1|\hat{t}_1, \dots, \hat{t}_{n-1}, t) \dots \lambda(\hat{t}_{n-1}|\hat{t}_1, \dots, \hat{t}_{n-1}, t) e^{-\int \lambda(x|\hat{t}_1, \dots, \hat{t}_{n-1}, t) dx} \lambda(t|\hat{t}_1, \dots, \hat{t}_{n-1}, t) d\hat{t}_1 \dots d\hat{t}_{n-1}. \quad (6.20)$$

Since the current time t is necessarily greater or equal to all spike times in the history, we can use causality arguments to remove any conditional dependence on t . Re-labeling the infinite sum with $k = n - 1$, one readily sees that we recover the weighting factor $P(S_n)$ of a specific spike train with k spikes (Eq. 6.18) in front of the momentary firing intensity $\lambda(t)$:

$$A(t) = \sum_{k=0}^{\infty} \frac{1}{k!} \int \lambda(\hat{t}_1|\hat{t}_1, \dots, \hat{t}_k) \dots \lambda(\hat{t}_k|\hat{t}_1, \dots, \hat{t}_k) e^{-\int \lambda(x|\hat{t}_1, \dots, \hat{t}_k) dx} \lambda(t|\hat{t}_1, \dots, \hat{t}_k) d\hat{t}_1 \dots d\hat{t}_k = \langle \lambda(t|S) \rangle. \quad (6.21)$$

Therefore we have shown Eq. 6.2.

7 Work in Progress

In this Chapter, I report on two projects in which I have invested considerable amount of time. The first is the extension of Chapter 5 to longer time scales in the spike-triggered threshold and current. The second aims at predicting the spike times in a neuron with active dendrites.

7.1 Power-law Adaptation

The idea is to investigate the spike triggered adaptation at long time scales. Extended spike after-potentials are suggested by experimental evidences, functionality arguments and other theoretical frameworks. First, there is a rich literature reporting evidences that spike-frequency adaptation operates on time scales longer than one second [Schwindt et al., 1988b, Greene et al., 1994, Fairhall et al., 2001, Xu et al., 1996, Ulanovsky et al., 2004, Drew and Abbott, 2006, Arganda et al., 2007, Lundstrom et al., 2008, 2010]. Also, if adaptation is to be energy efficient in terms of redundancy reduction, predictive coding or noise shaping [Barlow, 1961, Srinivasan et al., 1982, Atick and Redlich, 1990, Chacron et al., 2005c] in an environment having inherent long-range correlations [Teich, 1989, Linkenkaer-Hansen et al., 2001, Dan et al., 1996], then it must contain some slow components. There are other theoretical arguments which somewhat indirectly suggest spike-triggered adaptation at long time scales. For instance, we know there is a theoretical link between optimal plasticity rules such as spike-timing dependent plasticity [Gerstner et al., 1996a, Markram et al., 1997, Pfister et al., 2006, Toyozumi et al., 2007, Sjöström and Gerstner, 2010] and spike-triggered adaptation [Hennequin et al., 2010]. From this perspective, the long time scales of plasticity processes [Fusi et al., 2005] may be linked to long time scales in spike-triggered adaptation. Another theoretical framework suggesting long time-scales relates to scale-free dynamics of ion-channels [Marom and Abbott, 1994, Toib et al., 1998, Marom, 2010].

The profiles of spike-triggered adaptation presented in Chapter 5 do not show time scales longer than 400 ms. In fact, time bins beyond 800 ms are not well constrained by the data; they have a large relative error. We hypothesized that greater precision at longer time scales




Figure 7.1: **Dynamic threshold and spike-triggered current on long time scales.** A: The dynamic threshold of individual cells (gray lines) or averaged across all recorded cells (blue and red) fit an exponential with time constant of 22 ms (dashed black line). The blue line was calculated using the recorded voltage for the maximization of the likelihood, the red line used the modelled voltage. B: The spike-triggered current of individual cells (gray lines) or averaged across cells (red line) fit a power-law with exponent minus one (dashed black line).

can be obtained by fitting on longer experiments. We used methods similar to those described in Chapter 5 applied on longer intra-cellular recordings (10 minutes compared to 1 minute in Chapter 5) to measure the spike-triggered currents at long time-scales (see Lundstrom et al. [2008] for details on the experiments).

We find that the dynamic threshold can be modelled as a single exponential, whereas the spike-triggered current can be modelled by a power-law (Fig. 7.1). The power-law extends from 2-5 ms to 10-20 seconds. After 20 seconds our measurements become imprecise and we cannot draw conclusions on the longer time scales. The effect of a single spike can still be felt, and measured, 20 seconds after it has been emitted. We found that replacing the effective spike after-potential by a truncated power-law can predict the spike-times as well as the firing rates on long time-scales.

7.2 Active Dendrites

A complete model for single neurons must correctly reproduce the firing of spikes and bursts. Yet it must remain simple enough to be tractable in large-scale simulations. We present a study of a simplified model of Layer V pyramidal cells of the cortex with active dendrites. We hypothesized that we can model the soma and its apical trunk with only two compartments, without significant loss in the accuracy of spike-timing predictions. The model is based on experimentally measurable impulse-response functions [Segev et al., 1995], which transfer the effect of current injected in one compartment to current reaching the other. Each com-

partment was modeled with a pair of non-linear differential equations with a small number of parameters that approximate the Hodgkin-and-Huxley equations. The predictive power of this model was tested on electrophysiological experiments where noisy current was injected in both the soma and the apical dendrite simultaneously [Larkum et al., 2004].




Figure 7.2: Schematic representation of the two-compartment model. **A** Somatic and dendritic compartments communicate through passive and active propagation. The somatic compartment has spike-triggered adaptation and a moving threshold. The dendritic compartment has an activation current and recovery current. **B** Associated experimental protocol with current injection both in soma and apical dendrite of layer 5 pyramidal cells of the rat somato-sensory cortex.

7.2.1 Description of the Model

The detailed two-compartment model follows the system of differential equations:

$$C_s \frac{dV_s}{dt} = -g_s(V_s - E_s) + \alpha m + I_s + \sum_{\{\hat{t}_i\}} I_A(t - \hat{t}_i) + \epsilon_{ds} * I_d \quad (7.1)$$

$$C_d \frac{dV_d}{dt} = -g_d(V_d - E_d) + g_1 m + g_2 x + I_d + \sum_{\{\hat{t}_i\}} I_{BAP}(t - \hat{t}_i) + \epsilon_{sd} * I_s \quad (7.2)$$

$$\tau_m \frac{dm}{dt} = \frac{1}{1 + \exp\left(-\frac{V - E_m}{D_m}\right)} - m \quad (7.3)$$

$$\tau_x \frac{dx}{dt} = m - x \quad (7.4)$$

$$\tau_T \frac{dV_T}{dt} = -(V_T - E_T) + D_T \sum_{\{\hat{t}_i\}} \delta(t - \hat{t}_i) \quad (7.5)$$

$$(7.6)$$

where I_s is the current injected in the soma, I_d the current injected in the dendrites, V_s is the somatic voltage, V_d is the dendritic voltage, m is the level of activation of a putative calcium current, x is the level of activation of a putative calcium-activated potassium current (K(Ca)), and V_T is the dynamic threshold for firing somatic spikes, I_A is a spike-triggered current mediating adaptaion, I_{BAP} is the the current associated with the back-propagating action potential, ϵ_{sd} is the filter relating the current injected in the soma to the current arriving in the dendrite and ϵ_{ds} is the filter relating the current injected in the dendrite to the current arriving in the soma. The spikes are emitted if $V_s(t) > V_T(t)$ which results in $\hat{t}_{(last)} = t$ while $V_s \rightarrow E_r$ and $t \rightarrow t + \tau_R$. The parameters are listed in Table 7.1.

The generalized passive model follows:

$$\lambda(t) = \lambda_0 \exp\left(\kappa_s * I_s + \kappa_{ds} * I_d + \sum_{\{\hat{t}_i\}} \eta_A(t - \hat{t}_i)\right) \quad (7.7)$$

where λ_0 is a constant related to the reversal potential, κ_s somatic membrane filter, κ_{ds} is the filter relating the current injected in the dendrite to the voltage change in the soma, and η_A is the effective spike-triggered adaptation. Each kernel (κ , η , I_A , I_{BAP}) is expressed as a linear combination of nonlinear basis (i.e. $\kappa(t) = \sum_i a_i f_i(t)$). The rectangular function was chosen as the nonlinear basis.

7.2.2 Experimental Protocol

Parasagittal brain slices of the somato-sensory cortex (300-350 m thick) were prepared from 28-35 day-old Wistar rats. Slices were cut in ice-cold extracellular solution (ACSF), incubated at 34°C for 20 min and stored at room temperature. During experiments, slices were superfused

Variable		Value	Units
Somatic leak conductance	g_s	22	nS
Somatic capacitance	C_s	379	pF
Somatic reversal potential	E_s	-73	mV
Threshold baseline	E_T	-53	mV
Spike-triggered jump in threshold	D_T	2.0	mV
Time-constant of dynamic threshold	τ_T	27	ms
Maximum 'Ca' current	g_1	567	pA
Maximum effect of 'Ca' current in soma	α	337	n.u.
Dendritic leak conductance	g_d	22	nS
Dendritic capacitance	C_d	86	pF
Dendritic reversal potential	E_d	-53	mV
Time-constant for variable m	τ_m	6.7	ms
Time-constant for variable x	τ_x	49.9	ms
Sensitivity of 'Ca' Current	D_m	5.5	ms
Maximum 'K(Ca)' Current	g_2	-207	pA
Half-activation potential of 'Ca' current	E_m	-0.6	mV

Table 7.1: List of parameters and their fitted value for the two-compartment model.

with in ACSF at 34°C. The ACSF contained (in mM) 125 NaCl, 25 NaHCO₃, 25 Glucose, 3 KCl, 1.25 NaH₂PO₄, 2 CaCl₂, 1 MgCl₂, pH 7.4, and was continuously bubbled with 5 % CO₂ / 95 % O₂. The intracellular solution contained (in mM) 115 K⁺-gluconate, 20 KCl, 2 Mg-ATP, 2 Na₂-ATP, 10 Na₂-phosphocreatine, 0.3 GTP, 10 HEPES, 0.1, 0.01 Alexa 594 and biocytin (0.2%), pH 7.2.

Recording electrodes were pulled from thick-walled (0.25 mm) borosilicate glass capillaries and used without further modification (pipette tip resistance 5-10 MΩ for soma and 20-30MΩ for dendrites). Whole-cell voltage recordings were performed at the soma of a layer V pyramidal cell. After opening of the cellular membrane a fluorescent dye, Alexa 594 could diffuse in the entire neuron allowing to perform patch clamp recordings on the apical dendrite 600-700 μm from the soma. Both recordings were obtained using Axoclamp Dagan BVC-700A amplifiers (Dagan Corporation). Data was acquired with an ITC-16 board (Instrutech) at 10 kHz driven by routines written in the Igor software (Wavemetrics).

The injection waveform consisted of 6 blocks of 12 seconds. Each block is made of three parts: 1) one second of low-variance colored noise injected only in the soma, 2) one second of low-variance colored noise injected only in the dendritic injection site, 3) ten seconds of high-variance colored noise whose injection site depends on the block: In the first block, the 10-second stimulus is injected only in the dendritic site, the second block delivers the 10-second stimulus in the soma only, and the four remaining blocks deliver simultaneous injections in the soma and the dendrites. The colored noise was simulated with MATLAB as an Ornstein-Uhlenbeck process with a correlation time of 3 ms. The six blocks make a 72 seconds

stimulus that was injected repeatedly without redrawing the colored noise (frozen-noise). Twenty repetitions of the 72-second stimulus were carried out, separated by periods of 2-120 seconds. Out of the twenty repetitions, a set of seven successive repetitions were selected on the basis of high intrinsic reliability.

7.2.3 Fitting Methods

For the two-compartment model, we use a combination of regression methods and exhaustive search to maximize the mean square-error in voltage derivative:

- Fit of the dendritic compartment, knowing the injected currents and the somatic spiking history:
 1. Compute the first-order estimate of dV_d/dt (see Methods section of Chapter 5);
 2. Find the best estimates of the dendritic parameters linear in dV_d/dt given a set of nonlinear parameters (τ_m, D_m, E_m, τ_x). The best estimates are chosen through linear regression to minimize the mean square error of dV_d/dt .
 3. Compute iteratively step 2. on a grid of the nonlinear parameters (exhaustive search) and find the nonlinear parameters that yield the minimum mean square error of dV_d/dt .
- Fit of the somatic compartment using the fitted dendritic compartment as given.
 4. Compute the first order estimate of dV_s/dt .
 5. Find the best estimates of the somatic parameters linear in dV_s/dt given a set of nonlinear parameters (D_T, τ_T, E_T). The best estimates are chosen through linear regression to minimize the mean square error of dV_d/dt (see Methods section of Chapter 5).
 6. Compute iteratively step 5. on a grid of the nonlinear parameters (exhaustive search) and simulate the model with each set of nonlinear parameters. Take the parameters that yield the maximum Γ coincidence factor (see Sect. 3.2)

For the generalized linear model, we use maximum likelihood methods:

- Maximize the likelihood of observing the measured spikes given the input currents. The maximization is performed via the Fisher scoring method which is appropriate for generalized linear models (see Sect. 5.2.5)

7.2.4 Results

Dual patch-clamp recordings were performed in L5 Pyramidal cells of Wistar rats (see Experimental Methods). A simplified two compartment model (see Model Description) was




Figure 7.3: The two-compartment model fits qualitatively and quantitatively the electrophysiological recordings. A, B Overlay of the model (red) and experimental (black) somatic voltage trace. The dash box indicates the region that is stretched below. C, D The overlay of the model (red) and the experimental (blue) dendritic voltage is shown for its associated stretched somatic voltage time-series. Left and right show two different injection regimes contrasting by the amount of dendritic activity which is high for A, C and medium for B, D. E Residuals from the linear regression are shown for the somatic (black) and dendritic (blue) compartment. F For each repetition the Γ Coincidence factor is plotted against the intrinsic reliability of the cell. Grey points show the performance of the model on the test set and black points show the performance of the model on the training set. G Comparison of the inter-spike interval histogram for the model (red) and the experiment (black). H Comparison of the generalized passive (Pas), and the full two-compartment model (Full) with the intrinsic reliability (R) of the neuron in terms of the Γ coincidence factor. The averaged Γ factor is shown for the training set (black) and test set (Gray)

fitted on the first 36 seconds of stimulation for all repetitions and the rest of the data (36 sec) was reserved to evaluate the model's predictive power. The predictive power of the two-compartment model with active dendrites was then compared to a model without activity in the dendrites (see Sect. 7.2.1).




Figure 7.4: Fitted kernels of the two-compartment model. A The kernel for spike-triggered adaptation is negative and increases monotonically between 6 and 600 ms. B The back-propagating current reaching the dendrites is a short (2ms) and strong (900 pA) pulse. C The convolution kernel linking the current injected in the dendrite to the current reaching the soma. D The convolution kernel linking the current injected in the soma to the current reaching the dendrite.

Figure 7.3 summarizes the predictive power of the two compartment model. The somatic and dendritic voltage traces are well captured (Fig. 7.3 A-D). The main cause for erroneous prediction of the somatic voltage trace is extra or missed spikes (Fig. 7.3 A and B lower panels). The dendritic voltage trace of the model follows the recorded trace both in a low dendritic-input regime (Fig. 7.3 C) and in a high dendritic-input regime with dendritic spikes (Fig. 7.3 D). The greater spread of voltage-prediction-error (Fig. 7.3) is mainly explained by the larger range of voltages in the dendrites (somatic voltage prediction is strictly subthreshold whereas dendritic voltage prediction ranges from -70 mV to +40 mV). The interspike interval distribution is well predicted by the model (Fig. 7.3 G).




Figure 7.5: The model reproduces the qualitative features of active dendrites reported in Larkum et al. [1999] and Larkum et al. [2004]. **A** Dendritic non-linearity is triggered by somatic spiking at a critical frequency. Somatic spike-trains of 5 spikes are forced in the soma of the mathematical model at different firing frequencies. The normalized integral of the dendritic voltage is shown as a function of the somatic spiking frequency. **B** Dendritic injection modulates the slope of the somatic spiking-frequency vs. current curve. The slope of the frequency vs somatic mean current as measured between 5 and 50 Hz is plotted as a function of the dendritic mean current. Both somatic and dendritic currents injected are Ornstein-Uhlenbeck processes with a correlation time of 3 ms and a standard deviation of 300 pA. **C** Spike triggered average of the current injected in the soma (black) and in the dendrites (blue). **D** Burst-triggered average of the current injected in the soma (black) and in the dendrites (blue). For a 50-second stimulation with an Ornstein-Uhlenbeck process.

The generalized passive model does not predict as many spike times (Fig. 7.3 F and H). The intrinsic variability in the test set was 68% and the 2 compartment model predicted 50%. The prediction falls to 36 % in the absence of a dendritic non-linearity (Fig. 7.3 H). We can see that there is a greater amount of over-fitting in the two compartment model while overfitting does not appear for the passive model.

The fitted kernels indicate that the adaptation is mediated by a monotonically decaying current that starts very strongly and decays slowly for minimum 500 ms (Fig. 7.4 A). The back-propagating action potential is mediated by a strong pulse of current lasting 2-3 ms (Fig. 7.4 B). The coupling ϵ_{ds} from dendrite to soma has a maximal response after 2-3 ms and then decays so as to be slightly negative after 35 ms (Fig. 7.4 C). The coupling ϵ_{sd} from soma to dendrite follows qualitatively ϵ_{ds} with smaller amplitudes and slightly larger delays for the maximum and minimum peaks (Fig. 7.4 D), consistent with the larger membrane time-constant in the soma than in the dendrites.

The two-compartment model can reproduce qualitative features associated with dendritic non-linearity in the apical tuft of L5 pyramidal neurons. We study two of these features: the critical frequency Larkum et al. [1999] and the gain modulation Larkum et al. [2004]. The first relates to the critical somatic firing frequency above which a non-linear response is seen in the soma, reflecting calcium channel activation in the dendrites. To simulate the original experiment, we force 5 spike in the soma at various frequency and plot the integral of the dendritic voltage. The critical frequency for initiating a non-linear increase in summed dendritic voltage is 138 Hz (Fig. 7.5 A). Pérez-Garci et al. [2006] reported a critical frequency of 105 Hz while Larkum et al. [1999] reported 85 Hz. This appears to vary across different cells and pharmacological conditions.

The model also appears to perform gain modulation as in Larkum et al. [2004] (Fig. 7.5 B). The relation between somatic firing rate and mean somatic current depends on the dendritic excitability. The onset (or shift) but also the gain (or slope) of the somatic frequency versus somatic current curve depend on the mean dendritic current. The gain modulation is attributed to a greater presence of bursts (Fig. 7.5 B) caused by dendritic calcium-current activation at higher dendritic input. The link between burst and dendritic activity is reflected in the burst- and spike-triggered average injected current (Fig. 7.5 C-D) similar to Larkum et al. [2004]. The burst-triggered current is greater for the dendritic injection, whereas the spike-triggered current is larger for somatic injection.

Further experiments would be required to assess the reproducibility of the results and to understand what are the most important parameters regulating the dendritic non-linearities.

Contribution

In this chapter, I state clearly my contribution in each of the preceding chapters as well as my contribution to other publications not reproduced in the thesis.

Chapter 1 This publication is a collaboration between Wulfram Gerstner and myself. I wrote most of the text and generated most of the figures.

Chapter 2 This is the first work I have done when starting in computational neuroscience. The project was started initially by Nicolas Marcille and coworkers [Marcille et al., 2007] following an idea of Wulfram Gerstner. Nicolas Marcille helped with the technical aspects involving the figure on the parameter space. Also, I am responsible for the technical work and the figures. The text was written by myself with strong help from Wulfram Gerstner. The conceptual work was done in close collaboration between Wulfram Gerstner and myself.

Chapter 3, Jolivet et al. [2008a], Jolivet et al. [2008b] and Gerstner and Naud [2009] The first competition was started by Renaud Jolivet, Felix Schürmann and Wulfram Gerstner, based on experiments by Alexander Rauch, Matthew Larkum and Maria Toledo-Rodriguez. My contribution to Jolivet et al. [2008a] was to optimize an AdEx model on challenge A. For the 2008 and 2009 competitions I worked with Thomas K. Berger from the lab of Henry Markram to devise the experimental protocol which would be used for challenge A and B. In 2008 and 2009 I was the main organizer of the competition, discussing the conceptual details with Wulfram Gerstner, Thomas K. Berger and Shaul Druckmann. Some of the administrative work was then taken over by Raphael Ritz from the INCF in 2009. I wrote Chapter 3 in close collaboration with Wulfram Gerstner.

Chapter 4 The conceptual work, writing and technical work for Naud et al. [To appear; 2012] was done mainly by myself. Skander Mensi, Felipe Gerhard and Wulfram Gerstner also contributed to the conceptual work. Felipe Gerhard optimized the GLM on the challenge data and made figure 4.12.

Contribution

Chapter 5 I devised the experiments with Michael Avermann and devised the fitting method with Skander Mensi and Christian Pozzorini. All the experiments were done by Michael Avermann and all the analysis was done by Skander Mensi. The text was written by Skander Mensi, Wulfram Gerstner and myself.

Chapter 6 The conceptual work was a close collaboration between myself and Wulfram Gerstner. I wrote the text and performed the analyses.

Chapter 7 The work on power-law adaptation started from the Master's thesis of Christian Pozzorini whom I supervised. I first found indices of power-law tails in my data from Brice Bathellier for the dendritic integration project. Most of the work to measure long power-law tail in spike-triggered adaptation was done by Christian Pozzorini. The work on dendritic integration took something like two years of my Ph. D. time. I devised the experiments with Brice Bathellier and performed the analysis myself.

Mensi et al. [2011] The conceptual work is a collaboration between Wulfram Gerstner, Skander Mensi and myself. I helped with the numerical implementation and wrote most of the text. Skander Mensi did all the analysis and all the figures.

Summary of Contributions I have contributed to the following knowledge:

- How subthreshold and spike-triggered adaptation give rise to multiple firing patterns.
- The importance of adaptation and refractoriness in quantitative neuron model.
- The importance of stochasticity in quantitative neuron model.
- Cortical cell-types can be differentiated mainly according to their spike-triggered adaptation profiles.
- The shape of spike-triggered adaptation in cortical cells.
- How information can be encoded and decoded in populations of adapting neurons.

Bibliography

- M. Abeles. *Corticonics*. Cambridge University Press, Cambridge, 1991.
- Pablo Achard and Erik De Schutter. Complex parameter landscape for a complex neuron model. *PLoS Comput Biol*, 2(7):e94, 2006. ISSN 1553-7358 (Electronic).
- P Adams, D Brown, and A Constanti. M-currents and other potassium currents in bullfrog sympathetic neurones. *The Journal of Physiology*, Jan 1982.
- E. D. Adrian. *The basis of sensation*. W.W. Norton, New York, 1928.
- H. Akaike. Information theory and an extension of the maximum likelihood principle. In Petrov B. N. and Csaki F, editors, *Proc. of the 2nd Int. Symp. on Information Theory*, pages 267–281, 1973.
- D. J. Amit and N. Brunel. A model of spontaneous activity and local delay activity during delay periods in the cerebral cortex. *Cerebral Cortex*, 7:237–252, 1997.
- D. J. Amit and M. V. Tsodyks. Quantitative study of attractor neural networks retrieving at low spike rates. i: Substrate — spikes, rates, and neuronal gain. *Network*, 2:259–273, 1991.
- S Arganda, R Guantes, and GG de Polavieja. Sodium pumps adapt spike bursting to stimulus statistics. *Nature Neuroscience*, 10(11):1467–1473, 2007.
- A. Arieli, A. Sterkin, A. Grinvald, and A. Aertsen. Dynamics of ongoing activity: Explanation of the large variability in evoked cortical responses. *Science*, 273:1868–1871, 1996.
- Maura Arsiero, Hans-Rudolf Luscher, Brian Nils Lundstrom, and Michele Giugliano. The impact of input fluctuations on the frequency-current relationships of layer 5 pyramidal neurons in the rat medial prefrontal cortex. *Journal of Neuroscience*, 27(12):3274–3284, 2007.
- J.J. Atick and A.N. Redlich. Towards a theory of early visual processing. *Neural Computation*, 4: 559–572, 1990.
- B.B Averbeck, PE Latham, and A Pouget. Neural correlations, population coding and computation. *Nature Reviews Neuroscience*, 7(5):358–366, 2006.

Bibliography

- R Azouz and C M Gray. Cellular mechanisms contributing to response variability of cortical neurons in vivo. *J Neurosci*, 19(6):2209–23, Mar 1999.
- R Azouz and C M Gray. Dynamic spike threshold reveals a mechanism for synaptic coincidence detection in cortical neurons in vivo. *Proc Natl Acad Sci USA*, 97(14):8110–5, Jul 2000.
- Rony Azouz and Charles M. Gray. Adaptive coincidence detection and dynamic gain control in visual cortical neurons in vivo. *Neuron*, 37:513–523, 2003.
- L. Badel, W. Gerstner, and M.J.E. Richardson. Dependence of the spike-triggered average voltage on membrane response properties. *Neurocomputing*, 69:1062–1065, 2007.
- L Badel, S Lefort, TK Berger, CCH Petersen, W Gerstner, and MJE Richardson. Extracting non-linear integrate-and-fire models from experimental data using dynamic i–v curves. *Biological Cybernetics*, 99(4):361–370, 2008a.
- Laurent Badel, S Lefort, R Brette, CCH Petersen, Wulfram Gerstner, and Magnus J E Richardson. Dynamic iv curves are reliable predictors of naturalistic pyramidal-neuron voltage traces. *Journal of Neurophysiology*, 99(2):656–666, 2008b.
- M D Baker and H Bostock. Inactivation of macroscopic late na⁺ current and characteristics of unitary late na⁺ currents in sensory neurons. *J Neurophysiol*, 80(5):2538–49, Nov 1998.
- F Baldissera, B Gustafsson, and F Parmiggiani. A model for refractoriness accumulation and secondary range firing in spinal motoneurons. *Biological Cybernetics*, 24:61–65, 1976.
- H. B. Barlow. Possible principles underlying the transformation of sensory messages. In W. A. Rosenbluth, editor, *Sensory Communication*, pages 217–234. MIT Press, 1961.
- Bruce P Bean. The action potential in mammalian central neurons. *Nat Rev Neurosci*, 8(6):451–65, Jun 2007.
- Michael Beierlein, Jay R Gibson, and Barry W Connors. Two dynamically distinct inhibitory networks in layer 4 of the neocortex. *Journal of neurophysiology*, pages 1–14, Oct 2003.
- J Benda and A Herz. A universal model for spike-frequency adaptation. *Neural Computation*, 15:2523–2564, 2003.
- Jan Benda, A Longtin, and L Maler. Spike-frequency adaptation separates transient communication signals from background oscillations. *Journal of Neuroscience*, 25(9):2312–2321, March 2005.
- Jan Benda, Leonard Maler, and André Longtin. Linear versus nonlinear signal transmission in neuron models with adaptation currents or dynamic thresholds. *J Neurophysiol*, 104(5):2806–20, Nov 2010.
- M Berry and M Meister. Refractoriness and neural precision. *Journal of Neuroscience*, 18(6):2200–2211, Mar 1998.

- N. Bertschinger and T. Natschläger. Real-time computation at the edge of chaos in recurrent neural networks. *Neural Computation*, 16:1413–1436, 2004.
- W Bialek, F Rieke, RR de Ruyter Van Steveninck, and D Warland. Reading a neural code. *Science*, 252(5014):1854–1857, 1991.
- R Brette and Wulfram Gerstner. Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *Journal of Neurophysiology*, 94:3637–3642, 2005.
- R Brette, Z Piwkowska, C Monier, M Rudolph-Lilith, J Fournier, M Levy, Y Frégnac, T Bal, and A Destexhe. High-resolution intracellular recordings using a real-time computational model of the electrode. *Neuron*, 59(3):379–391, 2008.
- E Brown, R Barbieri, V Ventura, R Kass, and L Frank. The time-rescaling theorem and its application to neural spike train data analysis. *Neural computation*, 14:325–346, 2002.
- J C Brumberg and B S Gutkin. Cortical pyramidal cells as non-linear oscillators: Experiment and spike-generation theory. *Brain Research*, 2007.
- Nicolas Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neuron. *Journal of Computational Neuroscience*, 8:183–208, Jan 2000.
- Nicolas Brunel. *Computational Modeling Methods for Neuroscientists*, chapter Modeling Point Neurons. MIT Press, 2010.
- H L Bryant and J P Segundo. Spike initiation by transmembrane current: a white-noise analysis. *Journal of Physiology*, 260(2):279–314, 1976.
- Ingrid Bureau, Francisca von Saint Paul, and Karel Svoboda. Interdigitated paralemniscal and lemniscal pathways in the mouse barrel cortex. *PLoS Biol*, 4(12):e382, Nov 2006.
- W. Calvin and C.F Stevens. Synaptic noise and other sources of randomness in motoneuron interspike intervals. *Journal of Neurophysiology*, 31:574–587, 1968.
- Matteo Carandini. Amplification of trial-to-trial response variability by neurons in visual cortex. *PLoS Biology*, 2(9):1483–1493, 2004.
- Matteo Carandini and DJ Heeger. Summation and division by neurons in primate visual cortex. *Science*, 264(5163):1333–1336, 1994.
- Matteo Carandini, Jonathan C Horton, and Lawrence C Sincich. Thalamic filtering of retinal spike trains by postsynaptic summation. *Journal of Vision*, 7(14):20.1–11, Jan 2007.
- A Carnell and D Richardson. Linear algebra for time series of spikes. In M. Verleysen, editor, *Proceedings of the 13th European Symposium on Artificial Neural Networks*, pages 363–368. d-side publications, 2005.

Bibliography

- M Chacron, B Lindner, L Maler, A Longtin, and J Bastian. Experimental and theoretical demonstration of noise shaping by interspike interval correlations. *Proceedings of the SPIE*, 5841:150–163, 2005a.
- M Chacron, L Maler, and J Bastian. Electroreceptor neuron dynamics shape information transmission. *Nature Neuroscience*, Jan 2005b.
- Maurice J Chacron, Khashayar Pakdaman, and André Longtin. Interspike interval correlations, memory, adaptation, and refractoriness in a leaky integrate-and-fire model with threshold fatigue. *Neural computation*, 15(2):253–78, Feb 2003.
- Maurice J Chacron, André Longtin, and Leonard Maler. Delayed excitatory and inhibitory feedback shape neural information transmission. *Phys Rev E Stat Nonlin Soft Matter Phys*, 72(5 Pt 1):051917, Nov 2005c.
- T. R. Chay and J. Rinzel. Bursting, beating, and chaos in an excitable membrane model. *Biophysical Journal*, 47(3):357–366, 1985.
- Z Chi and D Margoliash. Temporal precision and temporal drift in brain and behavior of zebra finch song. *Neuron*, 32(5):899–910, Dec 2001.
- Zhiyi Chi, Wei Wu, Zach Haga, Nicholas G Hatsopoulos, and Daniel Margoliash. Template-based spike pattern identification with linear convolution and dynamic time warping. *Journal of Neurophysiology*, 97(2):1221–35, Feb 2007.
- Markus Christen, Adam Kohn, Thomas Ott, and Ruedi Stoop. Measuring spike pattern reliability with the lempel-ziv-distance. *J Neurosci Methods*, 156(1-2):342–50, Sep 2006.
- S Clemens, D Combes, P Meyrand, and J Simmers. Long-term expression of two interacting motor pattern-generating networks in the stomatogastric system of freely behaving lobster. *Journal of neurophysiology*, 79(3):1396–408, Mar 1998.
- B. W. Connors and M. J. Gutnick. Intrinsic firing patterns of diverse neocortical neurons. *Trends in Neuroscience*, 13:99–104, 1990.
- EP Cook, JA Guest, Y Liang, NY Masse, and CM Colbert. Dendrite-to-soma input/output function of continuous time-varying signals in hippocampal cal pyramidal neurons. *Journal of neurophysiology*, 98(5):2943, 2007.
- D. R. Cox. *Renewal theory*. Methuen, London, 1962.
- Sylvain Crochet and Carl C H Petersen. Correlating whisker behavior with membrane potential in barrel cortex of awake mice. *Nat Neurosci*, 9(5):608–10, May 2006.
- Sylvain Crochet, James F A Poulet, Yves Kremer, and Carl C H Petersen. Synaptic mechanisms underlying sparse coding of active touch. *Neuron*, 69(6):1160–75, Mar 2011.

- L J Croner, K Purpura, and E Kaplan. Response variability in retinal ganglion cells of primates. *Proc Natl Acad Sci USA*, 90(17):8128–30, Sep 1993.
- D. Daley and D. Vere-Jones. *An introduction to the theory of point processes*. Springer, New York, 1988.
- Yang Dan, Joseph J Atick, and R.C Reid. Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *The Journal of Neuroscience*, 16(10):3351, 1996.
- Stephen V David and Jack L Gallant. Predicting neuronal responses during natural vision. *Network*, 16(2-3):239–60, Jan 2005.
- P. Dayan and L. F. Abbott. *Theoretical Neuroscience*. MIT Press, Cambridge, 2001.
- C. P. J. de Kock and B. Sakmann. Spiking in primary somatosensory cortex during natural whisking in awake head-restrained rats is cell-type specific. *Proc. National Academy Sciences (USA)*, 106(38):16446–16450, 2009.
- R. R. de Ruyter van Stevenick and W. Bialek. Real-time performance of a movement-sensitive neuron in the blowfly visual system: coding and information transfer in short spike sequences. *Proc. R. Soc. B*, 234:379–414, 1988.
- A. Destexhe, M. Rudolph, and D. Pare. The high-conductance state of neocortical neurons in vivo. *Nature Reviews Neuroscience*, 4:739–751, 2003.
- Patricia M DiLorenzo and Jonathan D Victor. Taste response variability and temporal coding in the nucleus of the solitary tract of the rat. *Journal of Neurophysiology*, 90(3):1418–31, Sep 2003.
- B Doiron, A Oswald, and L Maler. Interval coding. ii. dendrite-dependent mechanisms. *Journal of Neurophysiology*, 97:2744–2757, 2007.
- Dawei W. Dong and Joseph J. Atick. Statistics of natural time-varying images. *Network: Computation in Neural Systems*, 6(3):345–358, 1995.
- Patrick Drew and L. Abbott. Models and properties of power-law adaptation in neural systems. *Journal of neurophysiology*, Jan 2006.
- S Druckmann, Y Banitt, A Gidon, and F Schürmann. A novel multiple objective optimization framework for constraining conductance-based neuron models by experimental data. *Frontiers in Neuroscience*, 1(1):7–18, 2007.
- S Druckmann, T Berger, S Hill, and F Schürmann. Evaluating automated parameter constraining procedures of neuron models by experimental and surrogate data. *Biological Cybernetics*, 9:371–379, 2008.

Bibliography

- J J Eggermont, A M Aertsen, and P I Johannesma. Quantitative characterisation procedure for auditory neurons based on the spectro-temporal receptive field. *Hearing Research*, 10(2): 167–90, May 1983.
- J Eichhorn, A Tolias, A Zien, M Kuss, CE Rasmussen, J Weston, N Logothetis, and B Schölkopf. Prediction on spike data using kernel algorithms. *Advances in Neural Information Processing Systems*, 16:1367–1374, 2004.
- H Eyherabide, A Rokem, A Herz, and I Samengo. Bursts generate a non-reducible spike-pattern code. *Front. Neurosci*, 3(1):8–14, May 2009.
- A. L. Fairhall, G.D Lewen, W. Bialek, and R.R.D. van Steveninck. Efficiency and ambiguity in an adaptive neural code. *Nature*, 412(6849):787–792, 2001.
- AA Faisal, LPJ Selen, and DM Wolpert. Noise in the nervous system. *Nature Reviews Neuroscience*, 9(4):292–303, 2008.
- F Farkhooi, Eilif Muller, and Martin P Nawrot. Adaptation reduces variability of the neuronal population code. *Physical Review E*, 83:050905, 2011.
- I A Fleidervish, A Friedman, and M J Gutnick. Slow inactivation of na⁺ current and slow cumulative spike adaptation in mouse and guinea-pig neocortical neurones in slices. *J Physiol*, 493 (Pt 1):83–97, May 1996.
- N Fourcaud-Trocme, D Hansel, C Van Vreeswijk, and Nicolas Brunel. How spike generation mechanisms determine the neuronal response to fluctuating inputs. *Journal of Neuroscience*, 23(37):11628–11640, 2003.
- M.G.F. Fuortes and F. Mantegazzini. Interpretation of the repetitive firing of nerve cells. *Journal General Physiology*, 45:1163–1179, 1962.
- Stefano Fusi, PJ Drew, and LF Abbott. Cascade models of synaptically stored memories. *Neuron*, 45(4):599–611, 2005.
- T J Gawne, J W McClurkin, B J Richmond, and L M Optican. Lateral geniculate neurons in behaving primates. iii. response predictions of a channel model with multiple spatial-to-temporal filters. *Journal of Neurophysiology*, 66(3):809–23, Sep 1991.
- T J Gawne, T W Kjaer, and B J Richmond. Latency: another potential code for feature binding in striate cortex. *Journal of Neurophysiology*, 76(2):1356–60, Aug 1996.
- W S Geisler, D G Albrecht, R J Salvi, and S S Saunders. Discrimination performance of single neurons: rate and temporal-pattern information. *Journal of Neurophysiology*, 66(1):334–62, Jul 1991.
- Luc J Gentet, Michael Avermann, Ferenc Matyas, Jochen F Staiger, and Carl C H Petersen. Membrane potential dynamics of gabaergic neurons in the barrel cortex of behaving mice. *Neuron*, 65(3):422–35, Feb 2010.

- GL Gerstein and NYS Kiang. An approach to the quantitative analysis of electrophysiological data from single neurons. *Biophysical Journal*, 1(1):15–28, 1960.
- W. Gerstner. Time structure of the activity in neural network models. *Physical Review E*, 51(1): 738–758, 1995.
- W. Gerstner. Spike-response model. *Scholarpedia*, 3(12):1343, 2008.
- W. Gerstner and R. Brette. Adaptive exponential integrate-and-fire model. *Scholarpedia*, 4(6): 8427, 2009.
- W. Gerstner and W.M. Kistler. *Spiking neuron models*. Cambridge University Press New York, 2002.
- W. Gerstner, R. Kempter, J.L. van Hemmen, and H. Wagner. A neuronal learning rule for sub-millisecond temporal coding. *Nature*, 383(6595):76–78, 1996a.
- W Gerstner, J van Hemmen, and J Cowan. What matters in neuronal locking? *Neural computation*, 8:1653–1676, 1996b.
- Wulfram Gerstner. Population dynamics of spiking neurons: Fast transients, asynchronous states, and locking. *Neural Computation*, 12:43–89, 2000.
- Wulfram Gerstner and R Naud. How good are neuron models? *Science*, 326:379–380, Oct 2009.
- N L Golding and D Oertel. Physiological identification of the targets of cartwheel cells in the dorsal cochlear nucleus. *Journal of neurophysiology*, 78(1):248–60, Jul 1997.
- T Gollisch and M Meister. Rapid neural coding in the retina with relative spike latencies. *Science*, 319:1108–1111, Feb 2008.
- CC Greene, PC Schwindt, and WE Crill. Properties and ionic mechanisms of a metabotropic glutamate receptor-mediated slow afterdepolarization in neocortical neurons. *Journal of Neurophysiology*, 72(2):693, 1994.
- Guillaume Hennequin, Wulfram Gerstner, and Jean-Pascal Pfister. Stp in adaptive neurons gives close to optimal information transmission. *Frontiers in Computational Neuroscience*, 4 (143):1–22, Aug 2010.
- A.V. Hill. Excitation and accomodation in nerve. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, 119:305–355, 1936.
- S. Hill and G. Tononi. Modeling sleep and wakefulness in the thalamocortical system. *J. Neurophysiol.*, 93(3):1671–1698, 2005.
- Bertil Hille. *Ionic channels of excitable membranes*. Sinauer, Sunderland, 1992.
- 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*, 117(4):500–544, 1952.

Bibliography

- A. V. Holden. *Chaos*. Princeton University Press, 1986.
- Frank C. Hoppensteadt and Eugene M. Izhikevich. *Weakly connected neural networks*. Springer, 1997.
- Conor Houghton. Studying spike trains using a van Rossum metric with a synapse-like filter. *J Comput Neurosci*, 26(1):149–55, Feb 2009.
- D. H. Hubel and T. N. Wiesel. Receptive fields of single neurons in the cat's striate cortex. *Journal of Physiology*, 148:574–591, 1959.
- J R Huguenard, O P Hamill, and D A Prince. Developmental changes in Na^+ conductances in rat neocortical neurons: appearance of a slowly inactivating component. *J Neurophysiol*, 59(3):778–95, Mar 1988.
- John D Hunter and John G Milton. Amplitude and frequency dependence of spike timing: implications for dynamic regulation. *Journal of Neurophysiology*, 90(1):387–94, Jul 2003.
- B Hutcheon and Y Yarom. Resonance, oscillation and the intrinsic frequency preferences of neurons. *Trends in Neurosciences*, 23:216–222, 2000.
- Q Huys, M Ahrens, and Liam Paninski. Efficient estimation of detailed single-neuron models. *Journal of neurophysiology*, Jan 2006.
- EM Izhikevich. Resonate-and-fire neurons. *Neural Networks*, 14(883-894), 2001.
- EM Izhikevich. Which model to use for cortical spiking neurons? *IEEE transactions on neural networks*, 15(5):1063–1070, 2004.
- E.M. Izhikevich and N.S. Desai. Relating stdp to bcm. *Neural Computation*, 15:1511–1523, 2003.
- EM Izhikevich, NS Desai, EC Walcott, and FC Hoppensteadt. Bursts as a unit of neural information: selective communication via resonance. *TRENDS in Neurosciences*, 26(3): 161–167, 2003.
- Eugene M Izhikevich. *Dynamical systems in neuroscience : the geometry of excitability and bursting*. MIT Press, Cambridge, Mass., 2007.
- Eugene M. Izhikevich and Gerald M. Edelman. Large-scale model of mammalian thalamocortical systems. *Proceedings of the National Academy of Sciences*, page 0712231105, 2008.
- P.L.M. Johannesma. Diffusion models of the stochastic activity of neurons. In *Neural Networks*, pages 116–144, Berlin, 1968. Springer.
- Roland S Johansson and Ingvars Birznieks. First spikes in ensembles of human tactile afferents code complex spatial fingertip events. *Nature Neuroscience*, 7(2):170–177, Feb 2004.

- R Jolivet. *Effective Minimal Threshold Models of Neuronal Activity*. PhD thesis, Ecole Polytechnique Federale de Lausanne, 2005.
- Renaud Jolivet, T Lewis, and Wulfram Gerstner. Generalized integrate-and-fire models of neuronal activity approximate spike trains of a detailed model to a high degree of accuracy. *Journal of Neurophysiology*, 92:959–976, 2004.
- Renaud Jolivet, A Rauch, H Lüscher, and Wulfram Gerstner. Predicting spike timing of neocortical pyramidal neurons by simple threshold models. *Journal of Computational Neuroscience*, 21:35–49, 2006.
- Renaud Jolivet, R Kobayashi, A Rauch, R Naud, S Shinomoto, and Wulfram Gerstner. A benchmark test for a quantitative assessment of simple neuron models. *Journal of Neuroscience Methods*, 169:417–424, 2008a.
- Renaud Jolivet, F Schürmann, TK Berger, R Naud, W Gerstner, and A Roth. The quantitative single-neuron modeling competition. *Biological Cybernetics*, 99(4):417–426, 2008b.
- Y Kawaguchi and K Hama. Two subtypes of non-pyramidal cells in rat hippocampal formation identified by intracellular recording and hrp injection. *Brain Res*, 411(1):190–5, May 1987.
- C Kayser, MA Montemurro, NK Logothetis, and S Panzeri. Spike-phase coding boosts and stabilizes information carried by spatial and temporal spike patterns. *Neuron*, 61(4):597–608, 2009.
- J Keat, P Reinagel, R C Reid, and M Meister. Predicting every spike a model for the responses of visual neurons. *Neuron*, 30:803–817, 2001.
- Naomi Keren, Noam Peled, and Alon Korngreen. Constraining compartmental models using multiple voltage recordings and genetic algorithms. *J Neurophysiol*, 94(6):3730–3742, 2005.
- W Kistler, Wulfram Gerstner, and J Hemmen. Reduction of the hodgkin-huxley equations to a single-variable threshold model. *Neural Computation*, 9:1015–1045, 1997.
- Ryota Kobayashi and Shigeru Shinomoto. State space method for predicting the spike times of a neuron. *Physical Review E*, 75(1):0011925, 2007.
- Ryota Kobayashi, Yasuhiro Tsubo, and Shigeru Shinomoto. Made-to-order spiking neuron model equipped with a multi-timescale adaptive threshold. *Frontiers in Computational Neuroscience*, 3:9, Jul 2009.
- C Koch. Cable theory in neurons with active, linearized membranes. *Biological Cybernetics*, 50:15–33, 1984.
- C. Koch. *Biophysics of Computation*. Oxford University Press, New York, Oxford, 1999.
- H. Kondgen, C. Geisler, S. Fusi, X. J. Wang, H. R. Luscher, and M. Giugliano. The dynamical response properties of neocortical neurons to temporally modulated noisy inputs in vitro. *Cerebral Cortex*, 18(9):2086–2097, 2008.

Bibliography

- H. Korn and P. Faure. Is there chaos in the brain? ii. experimental evidence and related models. *C R Biol*, 326(9):787–840, 2003.
- A Korngreen and B Sakmann. Voltage-gated k⁺ channels in layer 5 neocortical pyramidal neurones from young rats: subtypes and gradients. *The Journal of Physiology*, 525(3): 621–238, 2000.
- G Kreiman, R Krahe, W Metzner, C Koch, and F Gabbiani. Robustness and variability of neuronal coding by amplitude-sensitive afferents in the weakly electric fish *eigenmannia*. *Journal of Neurophysiology*, 84(1):189–204, Jul 2000.
- T Kreuz, JS Haas, A Morelli, HDI Abarbanel, and A Politi. Measuring spike train synchrony. *Journal of Neuroscience Methods*, 165(1):151–161, 2007.
- Thomas Kreuz, Daniel Chicharro, Ralph G Andrzejak, Julie S Haas, and Henry D I Abarbanel. Measuring multiple spike train synchrony. *J Neurosci Methods*, 183(2):287–99, Oct 2009.
- G La Camera, A Rauch, H Lüscher, and W Senn. Minimal models of adapted neuronal response to in vivo-like input currents. *Neural computation*, 16:2101–2124, 2004.
- G La Camera, A Rauch, D Thurbon, HR Lüscher, Walter Senn, and Stefano Fusi. Multiple time scales of temporal response in pyramidal and fast spiking cortical neurons. *Journal of Neurophysiology*, 96:3448–3464, 2006.
- L. Lapicque. Recherches quantitatives sur l’excitation électrique des nerfs traitée comme une polarization. *J. Physiol. Pathol. Gen.*, 9:620–635, 1907. Cited in H.C. Tuckwell, *Introduction to Theoretic Neurobiology*. (Cambridge Univ. Press, Cambridge, 1988).
- Matthew E Larkum, Walter Senn, and Hans-R Luscher. Top-down dendritic input increases the gain of layer 5 pyramidal neurons. *Cerebral Cortex*, 14(10):1059–1070, 2004.
- ME Larkum, J Zhu, and B Sakmann. A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature*, 398:338–341, 1999.
- PE Latham, BJ Richmond, PG Nelson, and S Nirenberg. Intrinsic dynamics in neuronal networks. i. theory. *Journal of neurophysiology*, 83(2):808, 2000.
- Simon B Laughlin and Terrence J Sejnowski. Communication in neuronal networks. *Science*, 301(5641):1870–4, Sep 2003.
- R Lestienne. Determination of the precision of spike timing in the visual cortex of anaesthetised cats. *Biological Cybernetics*, 74(1):55–61, 1995.
- Benjamin Lindner. Superposition of many independent spike trains is generally not a poisson process. *Phys. Rev. E*, 73(2):022901, Feb 2006.
- K Linkenkaer-Hansen, V V Nikouline, J M Palva, and R J Ilmoniemi. Long-range temporal correlations and scaling behavior in human brain oscillations. *Journal of Neuroscience*, 21(4):1370–7, Feb 2001.

- JE Lisman. Bursts as a unit of neural information: making unreliable synapses reliable. *TRENDS in Neurosciences*, 20(1):38–43, 1997.
- John Lisman. The theta/gamma discrete phase code occurring during the hippocampal phase precession may be a more general brain coding scheme. *Hippocampus*, 15(7):913–922, 2005.
- B Lundstrom, M Higgs, W Spain, and A Fairhall. Fractional differentiation by neocortical pyramidal neurons. *Nature Neuroscience*, 11:1335–1342, Nov 2008.
- Brian Nils Lundstrom, Adrienne L Fairhall, and Miguel Maravall. Multiple timescale encoding of slowly varying whisker stimulus envelope in cortical and thalamic neurons in vivo. *J Neurosci*, 30(14):5071–7, Apr 2010.
- J M MacPherson and J W Aldridge. A quantitative method of computer analysis of spike train data collected from behaving animals. *Brain Research*, 175(1):183–7, Oct 1979.
- D Madison and R Nicoll. Control of the repetitive discharge of rat ca 1 pyramidal neurones in vitro. *The Journal of Physiology*, Jan 1984.
- Z Mainen and Terrence J Sejnowski. Influence of dendritic structure on firing pattern in model neocortical neurons. *Nature*, Jan 1996.
- Z F Mainen, J Joerges, J R Huguenard, and T J Sejnowski. A model of spike initiation in neocortical pyramidal neurons. *Neuron*, 15(6):1427–1439, 1995.
- Zachary F Mainen and Terrence J Sejnowski. Reliability of spike timing in neocortical neurons. *Science*, 268(5216):1503–1506, 1995.
- A Manwani and C Koch. Detecting and estimating signals in noisy cable structures, i: neuronal noise sources. *Neural computation*, 11(8):1797–1829, 1999.
- N Marcille, C Clopath, R Ranjan, S Druckmann, F Schürmann, H Markram, and W Gerstner. Predicting neuronal activity with an adaptive exponential integrate-and-fire model. In *BMC Neuroscience*, page 121, 2007.
- H Markram. The blue brain project. *Nature Reviews Neuroscience*, 7:153–159, 2006.
- H Markram, M Toledo-Rodriguez, Y Wang, A Gupta, G Silberberg, and C Wu. Interneurons of the neocortical inhibitory system. *Nature Reviews Neuroscience*, 5(10):793–807, Oct 2004.
- Henry Markram, Joachim Lübke, Michael Frotscher, and Bert Sakmann. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, 275:213–215, 1997.
- P Z Marmarelis and V Z Marmarelis. *Analysis of Physiological Systems: The White-Noise Approach*. New York: Plenum Press, 1978.
- Shimon Marom. Neural timescales or lack thereof. *Prog Neurobiol*, 90(1):16–28, Jan 2010.

Bibliography

- Shimon Marom and LF Abbott. Modeling state-dependent inactivation of membrane currents. *Biophysical journal*, 67(2):515–520, 1994.
- M Martina and P Jonas. Functional differences in na⁺ channel gating between fast-spiking interneurons and principal neurons of rat hippocampus. *J Physiol*, 505 (Pt 3):593–603, Dec 1997.
- A Mauro, F Conti, F Dodge, and R Schor. Subthreshold behavior and phenomenological impedance of the squid giant axon. *J Gen Physiol*, 55(4):497–523, 1970. ISSN 0022-1295 (Print).
- J. Mayor and W. Gerstner. Online processing of multiple inputs in a sparsely-connected recurrent neural network. In O. Kaynak, E. Alpaydin, E. Oja, and L. Xu, editors, *Proceedings of the Joint International Conference ICANN/ICONIP 2003*, pages 839–845, Heidelberg, 2003. Springer-Verlag.
- P McCullagh and John A Nelder. *Generalized linear models*, volume 37. Chapman and Hall, 2nd ed edition, 1989.
- P McCullagh and John A Nelder. *Generalized linear models*, volume 37. Chapman & Hall/CRC, 2nd ed edition, 1998.
- S Mensi, R Naud, and W Gerstner. From stochastic nonlinear integrate-and-fire to generalized linear models. In *Advances in Neural Information Processing Systems*, 2011.
- S. Mensi, R. Naud, M. Avermann, C. Petersen, and W. Gerstner. Parameter extraction and classification of three neuron types reveals two different adaptation mechanisms. *Parameter Extraction and Classification of Three Neuron Types Reveals two Different Adaptation Mechanisms*, Under Review.
- Stefan Mihalas and Ernst Niebur. A generalized linear integrate-and-fire neural model produces diverse spiking behaviors. *Neural computation*, 21(3):704–18, Mar 2009.
- C Monier, J Fournier, and Y Frégnac. In vitro and in vivo measures of evoked excitatory and inhibitory conductance dynamics in sensory cortices. *J Neurosci Methods*, 169(2):323–65, Apr 2008.
- Eilif Muller, Lars Buesing, Johannes Schemmel, and Karlheinz Meier. Spike-frequency adapting neural ensembles: beyond mean adaptation and renewal theories. *Neural Computation*, 19(11):2958–3010, Nov 2007.
- R Naud and W Gerstner. *Computational Systems Neurobiology*, chapter The Performance (and limits) of Simple Neuron Models: Generalizations of the Leaky Integrate-and-Fire Model. Springer, 2012a.
- R Naud and W Gerstner. *Spike Timing: Mechanisms and Function*, chapter Can We Predict Every Spike. CRC Press, 2012b.

- R Naud, N Marcille, C Clopath, and Wulfram Gerstner. Firing patterns in the adaptive exponential integrate-and-fire model. *Biological Cybernetics*, 99:335–347, 2008.
- R Naud, B Bathellier, M Larkum, and W Gerstner. Spike-timing prediction in a neuron model with active dendrites. In *Statistical Analysis of Neuronal Data 5 - Poster presentation*, 2010.
- R. Naud, F. Gerhard, S. Mensi, and W. Gerstner. Improved similarity measures for small sets of spike trains. *Neural Computation*, To appear; 2012.
- L. M. Optican and B. J. Richmond. Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. 3. Information theoretic analysis. *J. Neurophysiol.*, 57:162–178, 1987.
- António R. C Paiva, I Park, and JC Príncipe. A reproducing kernel hilbert space framework for spike train signal processing. *Neural computation*, 21(2):424–449, 2009a.
- António R. C Paiva, I Park, and JC Príncipe. A comparison of binless spike train measures. *Neural Computing & Applications*, 19(3):1–15, 2009b.
- António R. C Paiva, I Park, and JC Príncipe. Inner products for representation and learning in the spike train domain. In Karim G. Oweiss, editor, *Statistical Signal Processing for Neuroscience and Neurotechnology*. Academic Press, New York, 2010.
- Liam Paninski. Maximum likelihood estimation of cascade point-process neural encoding models. *Network: Computation in Neural Systems*, 15:243–262, 2004.
- Liam Paninski, JW Pillow, and E Simoncelli. Maximum likelihood estimation of a stochastic integrate-and-fire neural encoding model. *Neural computation*, 16:2533–2561, 2004.
- Liam Paninski, JW Pillow, and E Simoncelli. Comparing integrate-and-fire models estimated using intracellular and extracellular data. *Neurocomputing*, 65-66:379–385, 2005.
- Stefano Panzeri, Riccardo Senatore, Marcelo A Montemurro, and Rasmus S Petersen. Correcting for the sampling bias problem in spike train information measures. *Journal of Neurophysiology*, 98(3):1064–1072, 2007.
- G Passmore, A Selyanko, M Mistry, M Al-Qatari, S J Marsh, E A Matthews, A H Dickenson, T A Brown, S A Burbidge, M Main, and David A. Brown. Kcnq/m currents in sensory neurons: significance for pain therapy. *Journal of Neuroscience*, 23(18):7227–7236, Jan 2003.
- E Pérez-Garci, M Gassmann, B Bettler, and ME Larkum. The gabab1b isoform mediates long-lasting inhibition of dendritic ca²⁺ spikes in layer 5 somatosensory pyramidal neurons. *Neuron*, 50(4):603–616, 2006.
- Rasmus S Petersen, Marco Brambilla, Michael R Bale, Andrea Alenda, Stefano Panzeri, Marcelo A Montemurro, and Miguel Maravall. Diverse and temporally precise kinetic feature selectivity in the vpm thalamic nucleus. *Neuron*, 60(5):890–903, Dec 2008.

Bibliography

- J Pfister, T Toyozumi, D Barber, and Wulfram Gerstner. Optimal spike-timing-dependent plasticity for precise action potential firing in supervised learning. *Neural Computation*, 18: 1318–1348, Jan 2006.
- Jean-Pascal Pfister, Peter Dayan, and Máté Lengyel. Synapses with short-term plasticity are optimal estimators of presynaptic membrane potentials. *Nat Neurosci*, 13(10):1271–5, Oct 2010.
- F. G. Pike, R. S. Goddard, J. M. Suckling, P. Ganter, N. Kasthuri, and O. Paulsen. Distinct frequency preferences of different types of rat hippocampal neurones in response to oscillatory input currents. *Journal of Physiology*, 529(1):205–213, 2000.
- JW Pillow, Liam Paninski, VJ Uzzell, EP Simoncelli, and EJ Chichilnisky. Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model. *Journal of Neuroscience*, 25(47):11003–11013, 2005.
- JW Pillow, J Shlens, Liam Paninski, A Sher, AM Litke, EJ Chichilnisky, and EP Simoncelli. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*, 454 (7207):995–999, 2008.
- H Plesser and Wulfram Gerstner. Noise in integrate-and-fire neurons: From stochastic input to escape rates. *Neural Computation*, 12:367–384, 2000a.
- H. E. Plesser and W. Gerstner. Escape rate models for noisy integrate-and-fire neurons. *Neuro-computing*, 32-33:219–224, 2000b.
- A. V. Poliakov, R. K. Powers, A. Sawczuk, and M. C. Binder. Effects of background noise on the response of rat and cat motoneurons to excitatory current transients. *Journal of Physiology*, 495:143–157, 1996.
- A. V. Poliakov, R. K. Powers, and M. C. Binder. Functional identification of input-output transforms of motoneurons in cat. *J. Physiology*, 504:401–424, 1997.
- Alon Polsky, Bartlett Mel, and Jackie Schiller. Encoding and decoding bursts by nmda spikes in basal dendrites of layer 5 pyramidal neurons. *Journal of Neuroscience*, 29(38):11891–903, Sep 2009.
- Martin Pospischil, Maria Toledo-Rodriguez, Cyril Monier, Zuzanna Piwkowska, Thierry Bal, Yves Frégnac, Henry Markram, and Alain Destexhe. Minimal hodgkin-huxley type models for different classes of cortical and thalamic neurons. *Biol Cybern*, 99(4-5):427–41, Nov 2008.
- James F A Poulet and Carl C H Petersen. Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature*, 454(7206):881–5, Aug 2008.
- W H Press, S A Teukolsky, B P Flannery, and W T Vetterling. *The Art of Scientific Computing*. Cambridge University Press, 3d edition, 2007.

- R Quian Quiroga, A Kraskov, T Kreuz, and P Grassberger. Performance of different synchronization measures in real data: A case study on electroencephalographic signals. *Phys. Rev. E*, 65(4):041903, Mar 2002a.
- R Quian Quiroga, T Kreuz, and P Grassberger. Event synchronization: A simple and fast method to measure synchronicity and time delay patterns. *Phys. Rev. E*, 66(4):041904, Oct 2002b.
- A Rauch, G La Camera, H Luscher, Walter Senn, and Stefano Fusi. Neocortical pyramidal cells respond as integrate-and-fire neurons to in vivo-like input currents. *Journal of neurophysiology*, 90:1598–1612, Jan 2003.
- D Reich, F Mechler, K Purpura, and Jonathan D Victor. Interspike intervals, receptive fields, and information encoding in primary visual cortex. *Journal of Neuroscience*, 20(5):1964–1974, Mar 2000.
- D S Reich, J D Victor, B W Knight, T Ozaki, and E Kaplan. Response variability and timing precision of neuronal spike trains in vivo. *Journal of Neurophysiology*, 77(5):2836–41, May 1997.
- D S Reich, F Mechler, and J D Victor. Temporal coding of contrast in primary visual cortex: when, what, and why. *Journal of Neurophysiology*, 85(3):1039–50, Mar 2001.
- Alfonso Renart, Jaime de la Rocha, Peter Bartho, Liad Hollender, Néstor Parga, Alex Reyes, and Kenneth D Harris. The asynchronous state in cortical circuits. *Science*, 327(5965):587–90, Jan 2010.
- Magnus J E Richardson. Dynamics of populations and networks of neurons with voltage-activated and calcium-activated currents. *Phys Rev E Stat Nonlin Soft Matter Phys*, 80(2 Pt 1):021928, 2009.
- Magnus J E Richardson, Nicolas Brunel, and V Hakim. From subthreshold to firing-rate resonance. *Journal of Neurophysiology*, 89:2538–2554, 2003.
- F Rieke, D. Warland, R. de Ruyter van Steveninck, and W. Bialek. *Spikes - Exploring the neural code*. MIT Press, Cambridge, MA, 1996.
- John Rinzel and G. Bart Ermentrout. Analysis of neural excitability and oscillations. In C. Koch and I. Segev, editors, *Methods in Neuronal Modeling, 2nd. ed.*, pages 251–291, Cambridge, 1998. MIT Press.
- N H Sabah and K N Leibovic. Subthreshold oscillatory responses of the hodgkin-huxley cable model for the squid giant axon. *Biophys J*, 9(10):1206–1222, 1969.
- P Sah. Ca²⁺-activated k⁺ currents in neurones: types, physiological roles and modulation. *Trends in Neurosciences*, 19(4):150–154, 1996.

Bibliography

- Roland Schaette, Tim Gollisch, and Andreas V M Herz. Spike-train variability of auditory neurons in vivo: dynamic responses follow predictions from constant stimuli. *Journal of Neurophysiology*, 93(6):3270–81, Jun 2005.
- F P Schoenberg and K E Tranbarger. Description of earthquake aftershock sequences using prototype point patterns. *Environmetrics*, 9(3):271–286, 2008.
- B Schrauwen and JV Campenhout. Linking non-binned spike train kernels to several existing spike train metrics. *Neurocomputing*, 70(7-9):1247–1253, 2007.
- S Schreiber, JM Fellous, D Whitmer, P Tiesinga, and Terrence J Sejnowski. A new correlation-based measure of spike timing reliability. *Neurocomputing*, 52(54):925–931, 2003.
- P Schwindt, W Spain, and W Crill. Long-lasting reduction of excitability by a sodium-dependent potassium current in cat neocortical neurons. *Journal of Neuroscience*, 61(2):233–244, 1989.
- PC Schwindt, WJ Spain, RC Foehring, MC Chubb, and WE Crill. Slow conductances in neurons from cat sensorimotor cortex in vitro and their role in slow excitability changes. *Journal of Neurophysiology*, 59(2):450, 1988a.
- PC Schwindt, WJ Spain, RC Foehring, CE Stafstrom, MC Chubb, and WE Crill. Multiple potassium conductances and their functions in neurons from cat sensorimotor cortex in vitro. *Journal of Neurophysiology*, 59(2):424, 1988b.
- I Segev, W Rall, and J Rinzel. *The theoretical foundation of dendritic function*. MIT Press, 1995.
- Kamal Sen, J. C. Jorge-Rivera, Eve Marder, and L F Abbott. Decoding synapses. *Journal of Neuroscience*, 16(19):6307–6318, 1996.
- H Shimazaki and Shigeru Shinomoto. A method for selecting the bin size of a time histogram. *Neural computation*, 19(6):1503–1527, 2007.
- Shigeru Shinomoto. Fitting a stochastic spiking model to neuronal current injection data. *Neural Networks*, 23(6):764–9, Aug 2010.
- Lavi Shpigelman, Yoram Singer, Rony Paz, and Eilon Vaadia. Spikernels: predicting arm movements by embedding population spike rate patterns in inner-product spaces. *Neural Computation*, 17(3):671–90, Mar 2005.
- O. Shriki, D. Hansel, and H. Sompolinsky. Rate models for conductance-based cortical neuronal networks. *Neural Computation*, 15:1809–1841, 2003.
- J Sjöström and W Gerstner. Spike-timing dependent plasticity. *Scholarpedia*, 5(2):1362, 2010.
- Sean J Slee, Matthew H Higgs, Adrienne L Fairhall, and William J Spain. Two-dimensional time coding in the auditory brainstem. *J Neurosci*, 25(43):9978–88, Oct 2005.

- Donald L Snyder, Michael I Miller, and Donald L Snyder. *Random point processes in time and space*. Springer-Verlag, New York, 2nd ed. edition, 1991.
- W. Softky and C. Koch. The highly irregular firing pattern of cortical cells is inconsistent with temporal integration of random epsps. *Journal of Neuroscience*, 13:334–350, 1993.
- H Sompolinsky, A Crisanti, and H Sommers. Chaos in random neural networks. *Physical Review Letters*, 61(3):259–262, Jul 1988.
- M. Spiridon and W. Gerstner. Noise spectrum and signal transmission through a population of spiking neurons. *Network: Computation in Neural Systems*, 10:257–272, 1999.
- M V Srinivasan, S B Laughlin, and A Dubs. Predictive coding: a fresh view of inhibition in the retina. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, 216(1205): 427–459, Nov 1982.
- R B Stein. A theoretical analysis of neuronal variability. *Biophysical Journal*, 5:174–194, Feb 1965.
- R B Stein. The frequency of nerve action potentials generated by applied currents. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, 167(1006):64–86, 1967a.
- R B Stein. Some models of neuronal variability. *Biophys J*, 7(1):37–68, 1967b.
- J. F. Storm. Action potential repolarization and a fast after-hyperpolarization in rat hippocampal pyramidal cells. *Journal of Physiology*, 385:733–759, 1987.
- S. H. Strogatz. *Nonlinear dynamical systems and chaos*. Addison Wesley, Reading MA, 1994.
- S. P. Strong, Roland Koberle, Rob R. de Ruyter van Steveninck, and William Bialek. Entropy and information in neural spike trains. *The American Physical Society, Physical Review Letters*, 80 Issue 1:197–200, 1998.
- Nobuaki Tamamaki, Yuchio Yanagawa, Ryohei Tomioka, Jun-Ichi Miyazaki, Kunihiro Obata, and Takeshi Kaneko. Green fluorescent protein expression and colocalization with calretinin, parvalbumin, and somatostatin in the gad67-gfp knock-in mouse. *J Comp Neurol*, 467(1): 60–79, Dec 2003.
- T. Tateno, A. Harsch, and H.P.C. Robinson. Threshold firing frequency - current relationships of neurons in rat somatosensory cortex: Type 1 and type 2 dynamics. *J. Neurophysiology*, 92: 2283–2294, 2004.
- M C Teich. Fractal character of the auditory neural spike train. *IEEE Trans Biomed Eng*, 36(1): 150–60, 1989.
- S. Thorpe, D. Fize, and C. Marlot. Speed of processing in the human visual system. *Nature*, 381:520–522, 1996.

Bibliography

- P Tiesinga, JM Fellous, and Terrence J Sejnowski. Regulation of spike timing in visual cortical circuits. *Nature Reviews Neuroscience*, 9(2):97–109, 2008.
- P H E Tiesinga. Chaos-induced modulation of reliability boosts output firing rate in downstream cortical areas. *Physical review E*, 69(3 Pt 1):031912, Mar 2004.
- A Toib, V Lyakhov, and Shimon Marom. Interaction between duration of activity and time course of recovery from slow inactivation in mammalian brain na⁺ channels. *J Neurosci*, 18(5):1893–903, Mar 1998.
- Maria Toledo-Rodriguez, Barak Blumenfeld, Caizhi Wu, Junyi Luo, Bernard Attali, Philip Goodman, and Henry Markram. Correlation maps allow neuronal electrical properties to be predicted from single-cell gene expression profiles in rat neocortex. *Cerebral Cortex*, 14(12):1310–1327, 2004.
- J Touboul and R Brette. Dynamics and bifurcations of the adaptive exponential integrate-and-fire model. *Biological Cybernetics*, 99:319–334, 2008.
- J Touboul and R Brette. Spiking dynamics of bidimensional integrate-and-fire neurons. *SIAM J Applied Dynamical Systems*, 8(4):1462–1506, 2009.
- T Toyozumi, J Pfister, K Aihara, and W Gerstner. Optimality model of unsupervised spike-timing-dependent plasticity: synaptic memory and . . . *Neural computation*, Jan 2007.
- T Toyozumi, KR Rad, and L Paninski. Mean-field approximations for coupled populations of generalized linear model spiking neurons with markov refractoriness. *Neural Computation*, 21(5):1203–1243, 2009.
- A Treves and Stefano Panzeri. The upward bias in measures of information derived from limited data samples. *Neural computation*, 7:399–407, 1995.
- Wilson Truccolo, U Eden, M Fellows, John P Donoghue, and Emery N Brown. A point process framework for relating neural spiking activity to spiking history, neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *Journal of Neurophysiology*, 93:1074–1089, 2005.
- Wilson Truccolo, Leigh R Hochberg, and John P Donoghue. Collective dynamics in human and monkey sensorimotor cortex: predicting single neuron spikes. *Nature Neuroscience*, 13(1):105–111, Jun 2010.
- M. Tsodyks and H. Markram. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proc. Natl. Academy of Sci., USA*, 94:719–723, 1997.
- Alan M. Turing. Computing machinery and intelligence. *Mind*, LIX(236):433–460, 1950.
- N Ulanovsky, L Las, D Farkas, and I Nelken. Multiple time scales of adaptation in auditory cortex neurons. *Journal of Neuroscience*, 24(46):10440, 2004.

- N. G. van Kampen. *Stochastic processes in physics and chemistry*. North-Holland, Amsterdam, 2nd edition, 1992.
- M. C. W. van Rossum. A novel spike distance. *Neural Computation*, 13:751–763, 2001.
- C. van Vreeswijk and H. Sompolinsky. Chaotic balanced state in a model of cortical circuits. *Neural Computation*, 10:1321–1371, 1998.
- M C Vanier and J M Bower. A comparative survey of automated parameter-search methods for compartmental neural models. *J Comput Neurosci*, 7(2):149–171, 1999.
- J.D. Victor and K.P. Purpura. Metric-space analysis of spike trains: theory, algorithms, and application. *Network: Comput. Neural Syst*, 8:127–164, 1997.
- Jonathan D Victor. Spike train metrics. *Current Opinion in Neurobiology*, 15:585–592, 2005.
- Jonathan D Victor and K Purpura. Nature and precision of temporal coding in visual cortex: a metric-space analysis. *Journal of Neurophysiology*, 76(2):1310–1326, 1996.
- VG Voinov. Unbiased estimation of powers of the inverse of mean and related problems. *Sankhyā: The Indian Journal of Statistics, Series B*, 47(3):354–364, 1985.
- VG Voinov and M. S Nikulin. *Unbiased estimators and their applications*, volume v. 263, 362 of *Mathematics and its applications*. Kluwer Academic Publishers, Dordrecht, 1996.
- L Wang, R Narayan, G Grana, M Shamir, and K Sen. Cortical discrimination of complex natural stimuli: Can single neurons match behavior? *Journal of Neuroscience*, 27(3):582–589, 2007.
- Xiao-Jing Wang, Yinghui Liu, Maria V Sanchez-Vives, and David A McCormick. Adaptation and temporal decorrelation by single neurons in the primary visual cortex. *J Neurophysiol*, 89(6):3279–93, Jun 2003.
- Y. Wang, A. Gupta, M. Toledo-Rodriguez, C.Z. Wu, and H. Markram. Anatomical, physiological, molecular and circuit properties of nest basket cells in the developing somatosensory cortex. *Cerebral Cortex*, 12:395–410, 2002.
- S Weisberg. *Applied Linear Regression*. Wiley/Interscience, 2005.
- John A White, Jay T Rubinstein, and Alan R Kay. Channel noise in neurons. *Trends in Neuroscience*, 23(3):1–7, Feb 2000.
- H. R. Wilson and J. D. Cowan. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, 12:1–24, 1972.
- H Wolf and KG Pearson. Comparison of motor patterns in the intact and deafferented flight system of the locust. iii: Patterns of interneuronal activity. *Journal of comparative physiology. A, Sensory, neural, and behavioral physiology*, 165(1):61–74, 1989.

Bibliography

Nanping Wu, Akifumi Enomoto, Susumu Tanaka, Chie-Fang Hsiao, Duane Q Nykamp, Eugene Izhikevich, and Scott H Chandler. Persistent sodium currents in mesencephalic v neurons participate in burst generation and control of membrane excitability. *J Neurophysiol*, 93(5): 2710–22, May 2005.

Z Xu, J R Payne, and M E Nelson. Logarithmic time course of sensory adaptation in electrosensory afferent nerve fibers in a weakly electric fish. *Journal of Neurophysiology*, 76(3): 2020–2032, 1996.

Richard Naud - Curriculum Vitae

Personal Information

Nationality: Canadian
Marital Status: Married
Address: Vers la Doey, Corbeyrier, CH-1856
Phone: +41 (0) 78 849 3147
E-mail: richard.naud@epfl.ch
Web: lcn.epfl.ch/~naud/

Education

2007-present: **Ph. D. Candidate**
Neuroscience,
École Polytechnique Fédérale
de Lausanne (EPFL).
2004-2006: **Master of Science (M. Sc.)**
Physics,
McGill University.
2000-2003: **Bachelor (B. Sc.)**
Honors Physics,
McGill University.

Research Experience

2007-present: **Theoretical Neuroscience**
Prof. Wulfram Gerstner,
EPFL.
2003-2007: **Biophysics and Microscopy**
Prof. Paul Wiseman,
McGill University and
Prof. Paul De Koninck,
Université Laval.
Winter 2005: **Quantum Optics and
Microscopy**
Prof. Oscar Martinez,
Universidad de Buenos Aires.
Summer 2002: **Oceanography**
Prof. Brad deYoung,
Memorial University of
Newfoundland and Labrador.
Summer 2001: **Atmospheric sciences**
Prof. Henry Leighton,
McGill University.

Scientific Publications

Mensi, S., **Naud, R.**, Pozzorini, C., Avermann, M., Petersen, C.H., Gerstner, W., **Parameter Extraction and Classification of Three Neuron Types Reveals two Different Adaptation Mechanisms**, *under review*.
Naud, R. and Gerstner, W., **The Performance (and limits) of Simple Neuron Models: Generalizations of the Leaky Integrate-and-Fire Model**. Chapter to be published in textbook: *Computational Systems Neurobiology*.
Naud, R., Gerstner, W., **Can We Predict Every Spike?**, chapter in *Spike Timing: Mechanisms and Function*, Patricia M. DiIorenzo and Jonathan D. Victor (Editors), (2012).
Naud, R., Gerhard, R., Mensi, S., Gerstner, W., **Improved Similarity Measures for Small Sets of Spike Trains**, to appear in *Neural Computation*.
Mensi, S., **Naud, R.**, Gerstner, W., **From Non-linear Integrate-and-Fire to Generalized Linear Model**, NIPS (2011).
Gerstner W. and **Naud, R.**, **How Good Are Neuron Models?**, *Science* (viewpoint), (2009).
Naud, R., Marcille, N., Clopath, C., Gerstner, W., **Firing patterns in the adaptive exponential integrate-and-fire model**, *Biological Cybernetics*, (2008).
Jolivet, R., Schuermann, F., Berger, T. K., **Naud, R.**, Roth, A., Gerstner, W., **The quantitative single-neuron modeling competition**, *Biological Cybernetics*, (2008).
Jolivet, R., Kobayashi, R., Rauch, A., **Naud, R.**, Shinomoto, S., Gerstner, W., **A bench-mark test for a quantitative assessment of simple neuron models**, *Journal of Neurosciences Methods*, (2008).

Richard Naud - Curriculum Vitae

Publications in Preparation

Naud, R., Gerstner, W., Coding and Decoding with Adapting Neurons: A Population Approach.

Naud, R., Bathellier, B., Gerstner, W., Spike-timing prediction in a neuron model with active dendrites.

Pozzorini, C., **Naud, R.**, Gerstner, W., **Power-law Spike-triggered Adaptation.**

Prizes and Scholarships

- Ph. D. scholarship from FQRNT
Value: 60 000\$, (2007-2010).
- Master's scholarship from FQRNT
Value: 30 000\$, (2004-2006).
- Summer research scholarship NSERC
Value: 8 000\$ (2003)
- Summer research scholarship NSERC
Value: 8 000\$ (2002)
- Summer research scholarship NSERC
Value: 8 000\$ (2001)
- W. McConnell Prize, McGill University,
(2001).
- Emily Ross Crawford scholarship, McGill
University, (2001).
- Sir William McDonald scholarship
McGill University, (2001).
- E. R. Pounder prize in Physics
McGill University, (2001).

Invited Presentations

Improving the Discriminability of Similarity Measures for Small Sets of Spike Trains, *External Speaker Seminar of Norwegian University*, Oslo, Norway, (2011).

Improving the Discriminability of Similarity Measures for Small Sets of Spike Trains, *CNS Workshop on Spike Train Analysis*, Stockholm, Sweden, (2011).

Extracting multi-scale adaptation parameters of spiking neuron models from data, *Spike-Frequency Adaptation in Neural Systems*, Dresden, Germany, (2010).

Extracting multi-scale adaptation parameters of spiking neuron models from data, *Neural Coding*, Limassol, Cyprus, (2010).

Comparing Spike-Time Predictions, *Spike-Time Metrics Workshop*, Plymouth, United Kingdom, (2010).

The Spike-Timing Prediction Competition, FACETS Meeting, Leysin, Switzerland, (2009).

The Spike-Timing Prediction Competition *Lemanic Circle*, Morges, Switzerland, (2009).

Firing Patterns in Simple Neuron Models, FACETS Meeting, Dresden, Germany, (2009).

Firing Patterns in Simple Neuron Models, *FACETS Meeting*, Debrecen, Hungary, (2008).

Poster Presentations

Mensi, S., **Naud, R.**, Pozzorini, C., Avermann, M., Petersen, C. H., Gerstner, W., **Automatic Characterization of Three Cortical Neuron Types Reveals Two Distinct Adaptation Mechanisms**, CNS, Stockholm, (2011). **Poster Award**

Pozzorini, C., **Naud, R.**, Gerstner, W., **Multiple timescales of adaptation in Single Neuron Models**, Cosyne, Salt lake city, (2011).

Naud, R., Gerhard, F., Mensi, S., Gerstner, W., **Improving the Discriminability of Similarity Measures for Small Sets of Spike Trains**, Neural Coding, Limassol, (2010).

Naud, R., Gerhard, F., Mensi, S., Gerstner, W., **Improving the Discriminability of Similarity Measures for Small Sets of Spike Trains**, BCCN, Berlin (2010).

Richard Naud - Curriculum Vitae

Naud, R., Bathellier, B., Larkum, M., Gerstner, W., **Spike-timing prediction in a neuron model with active dendrites**, SAND, Pittsburgh (2010).

Mensi, S., **Naud, R.**, Berger, T. K., Gerstner, W. **Complexity and performance in simple neuron models**, INCF, Pilsen, (2009).

Naud, R., Berger, T. K., Bathellier, B., Carandini, M., Gerstner, W., **Quantitative Single-Neuron Modeling: Competition 2009**, INCF, Pilsen, (2009).

Naud, R., Bathellier, B., Gerstner, W., **Spike-timing prediction in a neuron model with active dendrites**, CNS, Berlin, (2009).

Naud, R., Berger, T. K., Badel, L., Roth, A., Gerstner, W., **Quantitative Single-Neuron Modeling: Competition 2008**, INCF, Stockholm, (2008). *Poster Highlight *

Naud, R., Gerstner, W., **Spike-timing prediction in a neuron model with active dendrites**, FENS, Geneva, (2008).

Thesis

Naud, R. Particle tracking with application to the trafficking of synaptic receptors, *McGill University*, (2007).

External Collaborations

- ▶ Bathellier, B., Universität Wien, experiments for active dendrites project.
- ▶ Carandini, M. from University College London, Spike-timing prediction competition.
- ▶ Gregor, R. from Université de Fribourg (Switzerland), cortical coding and decoding with adaptation.
- ▶ Larkum, M. from Universität Bern, active dendrites project.
- ▶ Roth, A., from University College London, Spike-timing prediction competition.

Other Academic Accomplishments

- ▶ Reviewer for *Journal of Computational Neuroscience* and *Journal of Neuroscience Methods*.
- ▶ Supervision of 2 Master students, 3 internship students, 1 Bachelor student.

Teaching Experience

- 2007-2011: **Teaching assistant**, Neural Networks and Biological Modeling for Physics and Life Sciences, EPFL.
- 2004-2006: **Teaching assistant**, Introduction to Physics Lab, McGill University.
- 2003: Co-founder of the **Mathematics Help Desk** at McGill University, with John Labute and Frédéric Laliberté.

Languages

- ▶ French: Mother tongue.
- ▶ English: Fluent written and spoken.
- ▶ Spanish: Working proficiency.
- ▶ German: Elementary proficiency.