

Computational principles of single neuron adaptation

THÈSE N° 6461 (2014)

PRÉSENTÉE LE 17 DÉCEMBRE 2014

À LA FACULTÉ INFORMATIQUE ET COMMUNICATIONS
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

Christian Antonio POZZORINI

acceptée sur proposition du jury:

Prof. M. D. C. Sandi Perez, présidente du jury
Prof. W. Gerstner, directeur de thèse
Prof. R. Brette, rapporteur
Prof. A. Longtin, rapporteur
Prof. R. Schneggenburger, rapporteur



ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Suisse
2014

Acknowledgements

First and foremost I wish to thank my advisor, Prof. Wulfram Gerstner, for the freedom, the trust and the support he offered me during my entire time at the LCN. I am equally indebted to Richard Naud, who, during the first part of my thesis, has been a mentor for me. A special thank goes to my friend and colleague Skander Mensi. Most of the ideas that led to this thesis originated during the countless coffee breaks I have spent with him at the cafeteria Giacometti. I also would like to thank Olivier Hagens and Shovan Naskar for trusting my ideas, keeping calm and carrying on patch-clamping.

I thank the former and current members of the LCN for their help and discussions. In particular, I am grateful to Guillaume Hennequin, Laureline Lagiaco, David Kastner, Carlos Stein Naves de Brito, Friedeman Zenke, Moritz Deger, Tilo Schwalger for their advices, criticisms and patience and to Moritz Deger, Marco Lehmann, Laureline Lagiaco, Olivia Gozel and Samuel Muscinelli for proofreading parts of this thesis. A special acknowledgement goes to Friedeman Zenke who always did his best to kindly answer my technical questions. I am also grateful to my external collaborators Christof Koch and Carl Petersen for their time and enthusiasm.

For great discussions, *footshocks* and *spikes* loosely related to neuroscience, as well as for letting me sometimes forget that the LCN is a working place, I warmly thank Lorric Ziegler, Felipe Gerhard, Carlos Stein Naves de Brito, Skander Mensi, Alex Seeholzer, Jonathan Bieler, Colin Molter, Danilo Jimenez Rezende and Tim Vogels. Coming to work has always been a pleasure for me.

Finally, I thank my mother, my father and the rest of my family for supporting me since the beginning and for giving me the opportunity to study at EFPL. Without them, none of this would have been possible. I am incredibly indebted to the future mother of my child, Josephine, who has always been of great support, especially in the worst times. I look forward to our new life.

Lausanne, 16 September 2014

C. P.

Abstract

Cortical neurons continuously transform sets of incoming spike trains into output spike trains. This input-output transformation is referred to as single-neuron computation and constitutes one of the most fundamental process in the brain. A deep understanding of single-neuron dynamics is therefore required to study how neural circuits support complex behaviors such as sensory perception, learning and memory.

The results presented in this thesis focus on single-neuron computation. In particular, I address the question of *how* and *why* cortical neurons adapt their coding strategies to the statistical properties of their inputs. A new spiking model and a new fitting procedure are introduced that enable reliable nonparametric feature extraction from *in vitro* intracellular recordings. By applying this method to a new set of data from L5 pyramidal neurons, I found that cortical neurons adapt their firing rate over multiple timescales, ranging from tens of milliseconds to tens of second. This behavior results from two cellular processes, which are triggered by the emission of individual action potentials and decay according to a power-law. An analysis performed on *in vivo* intracellular recordings further indicates that power-law adaptation is near-optimally tuned to efficiently encode natural inputs received by single neurons in biologically relevant situations. These results shade light on the functional role of spike-frequency adaptation in the cortex.

The second part of this thesis focuses on the long-standing question of whether cortical neurons act as temporal integrators or coincidence detectors. According to standard theories relying on simplified spiking models, cortical neurons are expected to feature both coding strategies, depending on the statistical properties of their inputs. A model-based analysis performed on a second set of *in vitro* recordings demonstrates that the spike initiation dynamics implements a complex form of adaptation to make cortical neurons act as coincidence detectors, regardless of the input statistics. This result indicates that cortical neurons are well-suited to support a temporal code in which the relevant information is carried by the precise timing of spikes.

The spiking model introduced in this thesis was not designed to study a particular aspect of single-neuron computation and achieves good performances in predicting the spiking activity of different neuronal types. The proposed method for parameter estimation is efficient

Abstract

and only requires a limited amount of data. If applied on large datasets, the mathematical framework presented in this thesis could therefore lead to automated high-throughput single-neuron characterization.

Keywords: Single-neuron computation – Spike-frequency adaptation – Adaptive coding – Spiking-neuron model – Patch-clamp recordings – Parameter estimation

Riassunto

Il sistema nervoso centrale è composto da un elevato numero di neuroni che, connessi tra loro, processano, codificano e trasmettono informazioni sotto forma di segnali elettrici chiamati impulsi nervosi. Generalmente un neurone è costituito da tre parti funzionali distinte: i dendriti, il soma e l'assone. I dendriti ricevono gli impulsi nervosi in entrata e li dirigono verso il soma dove sono trasformati in una serie di impulsi di uscita che, a loro volta, vengono trasmessi lungo l'assone a migliaia di altri neuroni. La trasformazione di impulsi nervosi effettuata da parte di singoli neuroni, comunemente chiamata *calcolo neurale*, è uno dei processi cerebrali più importanti e costituisce l'oggetto principale di questo studio. In particolare, questa tesi verte sui principi computazionali, i meccanismi biologici e le ragioni per cui il calcolo e la codifica di informazioni neurali si adattano in funzione delle proprietà statistiche degli impulsi nervosi di entrata.

Nella prima parte di questa tesi, viene presentato un nuovo modello matematico capace di predire con grande precisione temporale l'emissione di impulsi nervosi da parte di singoli neuroni. I risultati ottenuti applicando questo modello matematico a dati ottenuti tramite esperimenti elettrofisiologici eseguiti *in vitro* indicano che i neuroni corticali piramidali adattano il loro comportamento su più scale temporali, comprese tra il millisecondo e la decina di secondi. Questo fenomeno è originato da due processi cellulari attivati dall'emissione di singoli impulsi nervosi e il cui decadimento segue una legge di potenza. Tramite un'analisi spettrale di dati sperimentali ottenuti *in vivo*, viene inoltre dimostrato che l'adattamento neurale basato su una legge di potenza permette la codifica ottimale di stimoli ricevuti da singoli neuroni corticali in condizioni naturali. Globalmente, questi risultati dimostrano il ruolo dell'adattamento neurale nella corteccia cerebrale.

La seconda parte di questa tesi ha come scopo quello di capire se i neuroni corticali operino come integratori di impulsi nervosi asincroni o come rilevatori di impulsi nervosi temporalmente coincidenti. Secondo i modelli standard di dinamica neurale, i neuroni corticali adottano l'una o l'altra modalità in funzione delle proprietà statistiche degli impulsi nervosi afferenti. Un'analisi basata su risultati ottenuti usando un nuovo modello matematico e un nuovo set di dati sperimentali rivela che la complessa dinamica sottostante la generazione di impulsi neurali implementa una forma di adattamento che, indipendentemente dai segnali in entrata, forza i neuroni ad agire come rilevatori di coincidenza temporale. Indirettamente,

Abstract

questo risultato indica che, nella corteccia cerebrale, l'informazione neurale è rappresentata dalla precisa sequenza temporale con la quale gli impulsi nervosi vengono emessi.

Le proprietà elettrofisiologiche e computazionali dei neuroni possono variare notevolmente a seconda del tipo cellulare e dell'area cerebrale in questione. Essendo robusti e flessibili, i modelli e i metodi matematici presentati in questa tesi potranno in futuro essere impiegati al fine di caratterizzare e classificare automaticamente i dati ottenuti mediante esperimenti elettrofisiologici.

Parole chiave: Modello matematico di dinamica neurale – Calcolo e adattamento neurale – Informazione neurale – Registrazioni intracellulari da singoli neuroni – Estrazione di parametri da dati sperimentali

Contents

Acknowledgements	i
Abstract	iii
List of figures	xi
1 Introduction	1
1.1 Basic notions	2
1.1.1 The brain's fundamental building blocks	2
1.1.2 The Hodgkin-Huxley model	3
1.1.3 Simplified threshold models	6
1.1.4 Synaptic connections	7
1.1.5 Studying single-neuron computation <i>in vitro</i>	9
1.2 Sensory coding and sensory adaptation	13
1.2.1 The efficient coding hypothesis	13
1.2.2 Adaptive coding	15
1.2.3 The timescale of spike-frequency adaptation	16
1.2.4 Efficient coding of time-dependent signals by temporal whitening	18
1.2.5 The timescale of contrast adaptation	20
1.3 The adaptive properties of single neurons	22
1.3.1 Spike-frequency adaptation in single neurons	22
1.3.2 Single neurons as fractional differentiators	24
1.3.3 Spiking models of spike-frequency adaptation	26
1.3.4 Stochastic spike emission and Generalized Linear Model	28
1.3.5 Complex forms of single-neuron adaptation	30
1.4 Thesis contribution	33
1.5 Personal contribution	35
2 Temporal whitening by power-law adaptation in neocortical neurons	39
2.1 Introduction	39
2.2 Results	40
2.2.1 SFA is mediated by two power-law spike-triggered processes	40
2.2.2 Power-law SFA explains neural activity on short timescales	44
2.2.3 Power-law SFA explains neural activity on long timescales	46

Contents

2.2.4	Power-law SFA is optimally tuned for <i>temporal whitening</i>	48
2.3	Discussion	51
2.3.1	Extent of spike-triggered effects	51
2.3.2	Biophysical implementation of power-law adaptation	51
2.3.3	How general is power-law adaptation?	52
2.3.4	Functional implications	52
2.3.5	<i>Temporal Whitening</i> vs. <i>Noise-Shaping</i>	53
2.4	Materials and Methods	54
2.4.1	<i>In-vitro</i> electrophysiological recordings	54
2.4.2	Current Injections	55
2.4.3	Linear analysis	55
2.4.4	Generalized Leaky Integrate-and-Fire model (GLIF- ξ)	56
2.4.5	Data preprocessing: Active Electrode Compensation	57
2.4.6	Fitting the GLIF- ξ model on <i>in vitro</i> recordings	61
2.4.7	Performance evaluation	65
2.4.8	Estimating the statistical properties of the input current received <i>in vivo</i> by neocortical pyramidal neurons	65
2.4.9	Simulating the population response to <i>in vivo</i> like inputs	66
2.4.10	Statistics	66
2.5	Author contributions	67
3	Automated high-throughput characterization of single neurons by means of simplified spiking models	69
3.1	Introduction	69
3.1.1	Simplified threshold models	71
3.1.2	Extracting model parameters from experimental recordings	74
3.1.3	Model validation	77
3.2	Results	78
3.2.1	A protocol for automated high-throughput single-neuron characterization	78
3.2.2	Applying the protocol on <i>in silico</i> recordings	81
3.2.3	Applying the protocol on <i>in vitro</i> patch-clamp recordings	83
3.3	Discussion	85
3.3.1	GIF model limitations and comparison with other spiking models	85
3.3.2	Limitations of point-neuron models	88
3.3.3	Conclusion	89
3.4	Materials and Methods	90
3.4.1	GIF model parameter extraction	90
3.4.2	Determining the amount of data required to perform GIF model parameter extraction and validation	91
3.4.3	Active electrode compensation	94
3.4.4	Multi-compartmental model simulations	95
3.4.5	Electrophysiological recordings	95

3.4.6	Generalized Linear Model	96
3.5	Author contributions	97
4	Enhanced temporal coding by nonlinear threshold dynamics in neocortical neurons	99
4.1	Introduction	99
4.2	Results	101
4.2.1	Enhanced sensitivity to rapid input fluctuations	101
4.2.2	The voltage threshold for spike initiation depends on the input statistics	104
4.2.3	Modeling the firing threshold dynamics	106
4.2.4	iGIF model parameter extracted from intracellular recordings reveals a nonlinear coupling between membrane potential and firing threshold	109
4.2.5	The iGIF model captures enhanced sensitivity to rapid input fluctuations and predicts spikes with millisecond precision	111
4.2.6	Enhanced sensitivity to input fluctuations results from a nonlinear interaction between spike-dependent and voltage-dependent threshold adaptation	114
4.2.7	The iGIF model captures and explains the complex forms of adaptation revealed by the GLM-based analysis	117
4.2.8	L5 Pyr neurons feature two distinct forms of adaptation	119
4.3	Discussion	121
4.3.1	Biophysical implementation of nonlinear threshold dynamics	122
4.3.2	Are simplified spiking models getting complicated?	123
4.3.3	Connection to sensory adaptation	124
4.3.4	Conclusion	124
4.4	Materials and Methods	125
4.4.1	Electrophysiological recordings	125
4.4.2	Current injections	125
4.4.3	Data preprocessing	127
4.4.4	Extracting voltage threshold for spike initiation from <i>in vitro</i> recordings	127
4.4.5	Generalized Linear Model (GLM)	128
4.4.6	Inactivating Generalized Integrate-and-Fire model (iGIF)	128
4.4.7	iGIF model parameter extraction	129
4.4.8	Extracting the effective membrane timescale from intracellular recordings	132
4.4.9	iGIF model linearization	132
4.4.10	Performance evaluation	134
4.5	Author contributions	134
	Bibliography	139
	Curriculum Vitae	157

List of Figures

1.1	Schematic representation of an excitable membrane and its equivalent electric circuit	3
1.2	Single-neuron computation	10
1.3	Comparison between a LIF model and a L5 pyramidal neuron responding to a set of stationary fluctuating currents with different statistics	11
1.4	Efficient coding of non-stationary signals requires sensory adaptation	15
1.5	Sensory adaptation by motion-sensitive neurons in the fly visual system	16
1.6	Fractional differentiation by cortical pyramidal neurons	23
1.7	LIF models extended to account for spike-frequency adaptation	27
1.8	Escape-rate mechanism for stochastic spike generation and Generalized Linear Model	30
2.1	Experimental protocol and spiking neuron model GLIF- ξ	41
2.2	Adaptation filter of the GLIF- ξ model extracted from <i>in vitro</i> recordings	42
2.3	Effective adaptation filters of individual L5 pyramidal neurons	44
2.4	The GLIF- ξ model predicts the occurrence of single spikes with millisecond precision	45
2.5	The GLIF- ξ model accurately predicts the firing rate response on multiple timescales	47
2.6	GLIF- ξ_L and GLIF- ξ_{PL} have similar predictive power	49
2.7	Power-law adaptation is near-optimally tuned to perform <i>temporal whitening</i>	50
2.8	The functional shape of the effective adaptation filter does not shape the power spectrum of the effective noise	54
2.9	Active Electrode Compensation (AEC) removes artifacts on short timescales	58
2.10	Active Electrode Compensation removes artifacts on long timescales	60
2.11	Electrode properties are stable during single injections	61
2.12	Influence of extending the fitting procedure with the hidden variable I_{drift}	63
3.1	Schematic representation of the protocol for high-throughput single-neuron characterization	78
3.2	Testing the protocol for high-throughput single-neuron characterization on <i>in silico</i> recordings	82

List of Figures

3.3	Testing the protocol for high-throughput single-neuron characterization on <i>in vitro</i> patch-clamp recordings	86
3.4	Estimating the amount of data required to perform accurate GIF model parameter extraction and validation	93
4.1	Pyr neurons maintain sensitivity to rapid input fluctuations over a wide range of depolarizing offsets	102
4.2	A comparison between GLM filters extracted from the spiking response to different input statistics reveals complex forms of adaptation	103
4.3	Standard analysis of intracellular recordings suggests an intricate dependence of the firing threshold on the input statistics	105
4.4	Simplified integrate-and-fire model illustrating firing threshold modulation by fast Na ⁺ -channel inactivation	108
4.5	Inactivating Generalized Integrate-and-Fire model (iGIF) with parameters extracted from intracellular recordings in Pyr neurons	110
4.6	The iGIF model captures enhanced sensitivity to input fluctuations and predicts spikes with millisecond precision	112
4.7	Nonlinear interaction between spike-dependent and voltage-dependent threshold adaptation	116
4.8	The iGIF model captures and explains complex forms of adaptation	120

1 Introduction

The brain is composed of a large number of neurons that communicate between each others by means of short electrical pulses called action potentials or *spikes*. Since neurons are the fundamental building blocks of the central nervous system, it is of crucial importance to understand their electrical properties and the way in which they process information. Single-neurons continuously transform incoming spike trains into output spike trains, which are then transmitted to neighboring neurons. This input-output transformation is often referred to as *single-neuron computation* and is one of the most fundamental process occurring in the brain.

The main goal of this thesis is to investigate single-neuron computation by means of accurate mathematical models capable of describing the single-neuron response as it is observed experimentally. Rather than modeling in great detail all of the biophysical processes known to occur in single neurons, the applied strategy is to keep the mathematical models as simple as possible in order to identify and understand the fundamental principles underlying single-neuron computation. In particular, a central aim of this thesis is to understand how single neurons adapt their coding strategy to the statistical properties of the incoming signals.

This *Introduction* is divided in three sections and reviews some important results about the electrical properties of single-neurons, sensory adaptation and single-neuron adaptation.

1.1 Basic notions

1.1.1 The brain's fundamental building blocks

Neurons are highly specialized cells and constitute the fundamental building blocks of the central nervous system. Like all the other cells, neurons are enclosed by a lipid bilayer, which forms a membrane impermeable to ions. The electrical properties of neurons result from specific protein complexes called *ion-pumps* and *ion-channels*, which are embedded into the cell membrane and through which ions flow generating membrane currents (Fig. 1.1a). The most important ions underlying these currents are: sodium (Na^+), potassium (K^+), calcium (Ca^{2+}) and chloride (Cl^-) (Dayan and Abbott, 2001).

Ion-pumps actively move specific ions inside and outside the cell (Fig. 1.1a). As a result, intracellular and extracellular ion concentrations differ creating both a chemical and an electrical gradient across the cell membrane. The sum of all the electrical gradients generated by different ions is called *membrane potential* and its temporal evolution constitutes one of the most relevant signals in the brain. Typically, there is an excess of positive charges outside the cell that results in a negative membrane potential of around -70 mV. During an action potential, the membrane potential transiently rises to around $+50$ mV and then, after one or two milliseconds, returns to its normal values (Dayan and Abbott, 2001). In contrast to small membrane potential fluctuations, action potentials are actively transmitted to neighboring neurons.

Chemical and electrical gradients counteract each others by moving ions in opposite directions. A dynamic equilibrium between inwards and outwards ion-fluxes is achieved at a specific membrane potential E , known as *reversal potential*. According to the Nernst equation, the specific value of E is proportional to the logarithm of the ratio between the extracellular and the intracellular concentration of a positively charged particle, as well as to the temperature of the system (Kandel et al., 2000). Since different ions have different concentrations, reversal potentials are ion-dependent. For example, Na^+ and Ca^{2+} are much more concentrated outside than inside the cell. Consequently, their reversal potentials $E_{\text{Na}} \approx +70$ mV and $E_{\text{Ca}} \approx +140$ mV are positive. Conversely, K^+ is more concentrated inside the cell and its reversal potential $E_{\text{K}} \approx -90$ mV is negative (Izhikevich, 2007). Ion-pumps are constantly at work to maintain these reversal potentials over time.

A number of ion-conducting channels are embedded into the cell membrane and confer to neurons their specific electrical properties (Kandel et al., 2000) (Fig. 1.1a). The net conductance mediated by an ion-channel typically varies over time and depends on one or more of the following factors: i) the membrane potential (voltage-gated channels), ii) the concentration of a neurotransmitter such as glutamate or GABA (ligand-gated channels) and iii) other signals such as the intracellular Ca^{2+} or Na^+ concentration (Ca^{2+} - or Na^+ -dependent channels) (Kandel et al., 2000). Some ion-channels are permeable to several particles, others are highly selective to particular ions. In the latter case, the resulting current drives the mem-




Figure 1.1: Schematic representation of an excitable membrane and its equivalent electric circuit. **a)** Schematic illustration of a patch of membrane (black) embedding (from left to right): a sodium-potassium pump that actively moves Na^+ in the extracellular space and K^+ inside the cell; a non-selective conductance implementing a leakage through the membrane, a voltage-activated Na^+ -channel and a voltage-activated K^+ -channel. As a result of the constant action of different ion-pumps, Na^+ and Ca^{2+} are more concentrated in the extracellular space and K^+ and Cl^- are more concentrated in the intracellular space. While pumps actively move ions against chemical gradients, ion-channels mediate passive currents that follow these gradients. In the schematic representation, both Na^+ - and K^+ -channels are open. In reality, these channels open and close depending on the electrical potential V across the membrane. **b)** Equivalent circuit representing the patch of membrane shown in panel **b**. As indicated in the scheme, the leak conductance g_L is constant and does not depend on V . Panel **b** was adapted from Gerstner and Kistler (2002).

brane potential to the reversal potential associated with the ion underlying the current. For example, during the emission of an action potential, the transient opening of Na^+ -channels initially rises the membrane potential to $E_{\text{Na}} \approx +70$ mV (i.e., to the reversal potential of Na^+). A few milliseconds later, the delayed activation of rectifying K^+ -channels hyperpolarizes the membrane potential towards $E_{\text{K}} \approx -90$ mV (Dayan and Abbott, 2001).

A large variety of ion-channels exists with different selectivities, kinematics and activation properties (Hille, 1992; Kandel et al., 2000). A mathematical formalism introduced more than sixty years ago by Alan Hodgkin and Andrew Huxley enables to study how the overall electrical properties of different neurons emerge from specific sets of ion-channels (Koch, 1999).

1.1.2 The Hodgkin-Huxley model

The electrical behavior of a neuron is accurately described by an analogue electrical circuit composed of a capacitor, which describes how opposite charges accumulate on both sides of the cell membrane, and a set of time-dependent conductances $g_i(t)$, for the different ion channels (Fig. 1.1b). The chemical gradients generated by ion-pumps are assumed to be constant over time and, in the analogue electrical circuit, act as batteries providing the driving force for ions to flow through selective channels (Dayan and Abbott, 2001).

The Hodgkin-Huxley model of membrane potential dynamics is obtained by applying the Kirchhoff's law of current to the electrical circuit shown in Figure 1.1b (Gerstner and Kistler,

2002):

$$C\dot{V}(t) = -\sum_i I_i(t) + I_{\text{ext}}(t) + I_{\text{syn}}(t), \quad (1.1)$$

where the *dot* denotes the first-order temporal derivative, $I_C(t) = C\dot{V}(t)$ is the capacitive current (see Fig. 1.1b), C is the membrane capacitance, $\{I_i(t)\}$ is the set of ion-currents mediated by different channels expressed by the neuron, $I_{\text{ext}}(t)$ is an external input that can be experimentally controlled and $I_{\text{syn}}(t)$ is the net synaptic input received from neighboring neurons.

Membrane currents $I_i(t)$ mediated by persistent channels¹ are typically described by the following equation:

$$I_i(t) = \bar{g}_i \cdot m(t) \cdot (V(t) - E_i), \quad (1.2)$$

where \bar{g}_i is the maximal conductance (i.e., the total conductance obtained when all of the expressed channels of type i are simultaneously open), E_i is the reversal potential of the ion to which the channel is permeable and $m(t)$ is a gating variable describing the probability of the channel being active at a given moment in time² (Dayan and Abbott, 2001). In the case of voltage-dependent channels, the gating variable m depends on the membrane potential and its temporal dynamics is modeled by the following equation:

$$\tau(V)\dot{m} = -m + m_\infty(V), \quad (1.3)$$

where $\tau(V)$ is the voltage-dependent timescale over which m tends to its steady-state value $m_\infty(V)$. Voltage-dependent channels are typically activated at high potentials and $m_\infty(V)$ is approximated by a sigmoidal function of V ranging from $m_\infty = 0$ at low voltages to $m_\infty = 1$ at high voltages. However, there are ion channels that activate when the membrane potential is hyperpolarized. These channels are known as hyperpolarization-activated channels and their steady-state functions are described by inverse sigmoidal functions of V (Dayan and Abbott, 2001).

Another class of voltage-dependent channels, known as transient channels, only mediate transient currents. These channels are characterized by the fact that activation is followed by inactivation. The sodium channels responsible for spike emission belong to this category. To accurately model the membrane currents mediated by transient conductances, Eq. 1.2 has to be extended with an inactivation gating variable h modeled with an equation similar to the one used to describe the dynamics of m (i.e., Eq. 1.3). Activation and inactivation are generally assumed to be independent processes. Consequently, the net membrane current

¹In contrast to transient channels, persistent channels do not undergo inactivation. Consequently, their dynamics can be modeled using a single gating variable (i.e., the activation gating variable), which describes the activation probability.

²Since single-neurons express multiple ion-channels of the same type, the gating variable $m(t)$ can also be interpreted as the fraction of channels that are simultaneously open at a given moment in time t .

$I_i(t)$ mediated by a transient conductance is modeled as

$$I_i(t) = \bar{g}_i \cdot m(t) \cdot h(t) \cdot (V(t) - E_i), \quad (1.4)$$

where, compared to $m_\infty(V)$, the steady-state function $h_\infty(V)$ of the inactivation variable features an opposite voltage dependence (Dayan and Abbott, 2001). In the particular case of Na^+ -channels responsible for spike emission, the interplay between activation and fast inactivation results in a strong but transient current which is responsible for the voltage rise observed during the initial phase of an action potential. As discussed in Chapter 4, fast Na^+ -channel inactivation also affects single-neuron computation by enhancing sensitivity to rapid input fluctuations (Platkiewicz and Brette, 2011).

The mathematical framework presented in the previous pages was introduced in the 1950s in order to explain how the combined action of different voltage-dependent channels confer to neurons the ability of generating action potentials (Hodgkin and Huxley, 1952). The original Hodgkin-Huxley (HH) model is defined as:

$$C\dot{V}(t) = -g_L(V - E_L) - \bar{g}_{\text{Na}}m^3h(V - E_{\text{Na}}) - \bar{g}_{\text{K}}n^4h(V - E_{\text{K}}) + I_{\text{ext}}(t), \quad (1.5)$$

where the first three terms on the right hand side of the equation respectively describes: a leakage current (I_L) mediated by different ions types and characterized by a constant conductance g_L , a transient Na^+ -current (I_{Na}) responsible for the action potential upswing and a delayed-rectifying K^+ -current (I_{K}) responsible for the action potential downswing (Fig. 1.1b). According to the HH model, when a neuron is stimulated with a positive current $I_{\text{ext}}(t) > 0$, the membrane potential rises. If the input is sufficiently strong, the membrane potential reaches -50 mV, where Na^+ -channels rapidly activate (i.e., where the activation variable m rapidly grows to one) mediating a positive feedback current which further increases the membrane potential up to +50 mV. Immediately after the action potential upswing, Na^+ -channels inactivates (i.e., the inactivation variable h drops to zero) and rectifying K^+ -channels activates (i.e., the activation variable n goes to one) mediating a current that restores a negative membrane potential³(Hodgkin and Huxley, 1952).

After the seminal work of Hodgkin and Huxley, conductance-based models became a standard mathematical framework used to build biophysical models of single neurons (Catterall et al., 2012). In particular, the original HH model has been extended with different active conductances in order to understand how different ion-channels interact to shape the electrical behavior of single-neurons (see, e.g., Connor and Stevens (1971); Morris and Lecar (1981)). The HH model also provides the starting point to develop multi-compartmental models that take into account the spatial structure of single-neurons (Rall, 1977; De Schutter and Bower, 1994; Mainen and Sejnowski, 1996; Stuart and Spruston, 1998; Howell et al., 2000). In contrast to single-compartmental models, these models are generally based on anatomical reconstruc-

³A detailed description of the dynamics of the HH model during the emission of an action potential can be found in Dayan and Abbott (2001).

tions of single-neuron morphologies and explicitly describe the temporal evolution of the membrane potential $V(t, x)$ at different spatial locations x . When active conductances are accounted for in different parts of a neuron (e.g., when nonlinear dendritic conductances are considered), the single-neuron morphology is generally discretized in a finite set of compartments $V_x(t)$, and each of them is modeled using the standard HH formalism⁴ (Dayan and Abbott, 2001). For this reason, modern conductance-based models, and in particular multi-compartmental models (see, e.g., Hay et al. (2011)), consist of a complicated set of coupled nonlinear differential equations, whose understanding mostly relies on intensive numerical simulations (Herz et al., 2006). Developing multi-compartmental models and increasing the level of detail included in conductance-based models are nowadays active fields of research (Markram, 2006; Waldrop, 2012).

Mainly because of their extreme complexity, such detailed biophysical (DB) models suffer from several drawbacks. In particular, DB models are analytically intractable, require high-computing power and have an extremely large number of parameters, which are difficult to extract from experimental data⁵ (Herz et al., 2006). A second class of spiking neuron models, known as simplified single-compartment threshold models (i.e., integrate-and-fire models), exists in which the spatial structure of the neuron is neglected and where the intricate dynamics of different ion-channels is accounted for by simpler phenomenological descriptions (Lapicque, 1907; Gerstner and Kistler, 2002; Izhikevich, 2007). In contrast to DB models, simplified threshold models are analytically tractable, are suitable for large-scale network simulations and have a smaller number of parameters that can be extracted from intracellular recordings (Herz et al., 2006). For all these reasons, integrate-and-fire (IF) models are often preferred by computational neuroscientist aiming at understanding the principles underlying single-neuron computation and the emergent properties of neural networks.

1.1.3 Simplified threshold models

The leaky integrate-and-fire (LIF) model is probably the most popular model amongst simplified threshold models and is obtained by replacing the complex dynamics of the Na^+ - and K^+ -channels in the standard HH model with a simple threshold process. In the LIF model, the subthreshold dynamics of the membrane potential is governed by the following linear differential equation (Gerstner and Kistler, 2002):

$$C\dot{V}(t) = -g_L(V - E_L) + I_{\text{ext}}(t), \quad (1.6)$$

where, in comparison to Eq. 1.5, only the passive properties of the membrane are retained. The LIF model does not explicitly model the voltage dynamics during an action potential, but

⁴In multi-compartmental models, the Hodgkin-Huxley equations are extended in order to account for the passive currents flowing across neighboring compartments.

⁵Inferring DB model parameters from simple intracellular recordings is an ill-posed problem. Indeed, as it has been experimentally demonstrated, different combinations of ion-channel densities can give rise to similar firing behaviors (Goaillard et al., 2009).

simply states that, each time V crosses a certain threshold V_T from below, an action potential is fired and the membrane potential is reset to a low value $V_{\text{reset}} < V_T$. The LIF model can equivalently be formulated in its integral version given by:

$$V(t) = E_L + \int_{\hat{t}_{\text{last}}}^t K_m(t-s) I_{\text{ext}}(s) ds + \eta_R(t - \hat{t}_{\text{last}}) \quad (1.7)$$

where \hat{t}_{last} denotes the time of the last spike before time t , $K_m(t) = \Theta(t) \cdot C^{-1} \exp\left(-\frac{t}{\tau_m}\right)$ is the membrane filter of timescale $\tau_m = C/g_L$ and $\eta_R(t) = C(V_{\text{reset}} - E_L)K_m(t)$ describes the after-spike reset $V \rightarrow V_{\text{reset}}$, with $\Theta(t)$ being the Heaviside step function (Gerstner and Kistler, 2002).

The first IF model was introduced more than a century ago⁶ (Lapicque, 1907; Hill, 1936; Brunel and Van Rossum, 2007) and is still nowadays widely adopted to investigate the emergent properties of neural networks. While capturing the most essential feature of single-neurons, namely the fact that the neural information is encoded by means of action potentials, the LIF model is oversimplified and poorly captures the spiking activity of real neurons (Jolivet et al., 2008b). During the last century, the LIF model has been continuously revisited and extended⁷ in order to account for different features such as smooth spike-initiation (Abbott and van Vreeswijk, 1993; Latham et al., 2000; Feng, 2001; Fourcaud-Trocmé et al., 2003; Badel et al., 2008), subthreshold resonance (Izhikevich, 2001; Richardson et al., 2003), spike-frequency adaptation (Baldissera et al., 1976; Benda and Herz, 2003; Paninski et al., 2005; La Camera et al., 2006), firing threshold dynamics (Hill, 1936; Fuortes and Mantegazzini, 1962; Geisler and Goldberg, 1966; Jolivet et al., 2006b; Badel et al., 2008; Kobayashi et al., 2009; Platkiewicz and Brette, 2011) and stochastic spike emission (Stein, 1965; Gluss, 1967; Plesser and Gerstner, 2000; Pillow et al., 2005; Truccolo et al., 2005). Modern integrate-and-fire models (Izhikevich et al., 2003; Brette and Gerstner, 2005; Truccolo et al., 2005; Mihalas and Niebur, 2009), sometimes called generalized integrate-and-fire (GIF) models, provide good descriptions of single-neurons responding to somatically injected currents (Jolivet et al., 2008b,a; Gerstner and Naud, 2009).

1.1.4 Synaptic connections

The most important connection points between neurons are called chemical synapses⁸. At these places, an intricate biological machinery actively transforms the action potentials emitted by a presynaptic neuron into transient postsynaptic currents $I_{\text{syn}}(t)$. Briefly, when an

⁶Note that the first IF model was proposed before the biophysical mechanisms underlying action potential generation were discovered.

⁷Exhaustive reviews about existing simplified threshold models are provided by Gerstner and Kistler (2002); Izhikevich (2007); Naud and Gerstner (2013).

⁸Neurons also communicate with each other by means of electrical synapses (gap junctions), consisting of conductive links between the membranes of two neurons. Compared to chemical synapses, gap junctions are fast and often bidirectional. Electrical synapses cannot amplify the strength of the transmitted signal and mainly act as passive couplings between neurons (Dayan and Abbott, 2001).

Chapter 1. Introduction

action potential reaches a presynaptic terminal, the release of a specific neurotransmitter (e.g., glutamate, GABA) is triggered⁹. Neurotransmitters then bind to and activate specific ion-channels (e.g., AMPA, NMDA, GABA_A and GABA_B receptors), which are embedded in the membrane of the postsynaptic neuron. Depending on the specific ions to which these channels are selective to, synaptic inputs can either depolarize (excitatory synapses) or hyperpolarize (inhibitory synapses) the postsynaptic neuron (Kandel et al., 2000).

In the HH formalism, the total postsynaptic current $I_{\text{syn}}(t)$ is modeled by the following equation (Dayan and Abbott, 2001):

$$I_{\text{syn}}(t) = -g_{\text{syn}}(t) \cdot (V - E_{\text{syn}}), \quad (1.8)$$

where E_{syn} is the reversal potential of the postsynaptic channel¹⁰ and $g_{\text{syn}}(t)$ is the total synaptic conductance induced by a presynaptic spike train $S(t) = \sum_{\hat{t}_j} \delta(t - \hat{t}_j)$, with $\delta(\cdot)$ denoting the δ -Dirac function and $\{\hat{t}_j\}$ being the presynaptic spike times. Different models exist that describe how incoming spike trains are transformed into time-dependent synaptic conductances. These models range from complex ones, accounting for nonlinear voltage dependences (Jahr and Stevens, 1990) and short-term plasticity (Abbott et al., 1997; Tsodyks and Markram, 1997), to simpler ones stating that $g_{\text{syn}}(t)$ equals a filtered version of the input spike train (Destexhe et al., 2001):

$$g_{\text{syn}}(t) = \int_0^{\infty} \epsilon_g(s) S(t-s) ds = \sum_{\hat{t}_j} \epsilon_g(t - \hat{t}_j). \quad (1.9)$$

According to Eq. 1.9, each individual spike triggers a postsynaptic conductance with stereotypical shape $\epsilon_g(t)$ and the effects induced by different spikes add up linearly. An even more simplistic but widely adopted model of synaptic transmission is obtained by neglecting the fact that $I_{\text{syn}}(t)$ depends on the postsynaptic voltage. In this model, informally referred to as *current-based* synaptic model, $I_{\text{syn}}(t)$ is simply given by (Brunel and Hakim, 1999; Richardson, 2004):

$$I_{\text{syn}}(t) = \int_0^{\infty} \epsilon(s) S(t-s) ds = \sum_{\hat{t}_j} \epsilon(t - \hat{t}_j), \quad (1.10)$$

where $\epsilon(t)$ describes the time course of a postsynaptic current (PSC) induced by a presynaptic spike. In computational and theoretical studies, PSCs are typically modeled using a single

⁹According to Dale's principle, a neuron releases the same transmitters at all of its synapses. For example, all of the presynaptic terminals of GABAergic inhibitory neurons release GABA and all of the presynaptic terminals of glutamatergic excitatory neurons release glutamate.

¹⁰While for glutamatergic excitatory synapses the reversal potential is $E_{\text{syn}} \approx 0$ mV, the reversal potential of GABAergic inhibitory synapses is defined by the reversal potential of Chloride ions $E_{\text{syn}} \approx -80$ mV. This fact explains why excitatory and inhibitory synapses generally depolarize and hyperpolarize the postsynaptic neuron, respectively (Dayan and Abbott, 2001).

exponential function (Vogels et al., 2005):

$$\epsilon(t) = w \cdot \tau_{\text{syn}}^{-1} \cdot \exp\left(-\frac{t}{\tau_{\text{syn}}}\right), \quad (1.11)$$

with w being the strength of the synaptic connection and $\tau_{\text{syn}} \approx 3 - 10$ ms being the timescale over which the PSCs decay (Vogels et al., 2005). In current-based synaptic models, a synaptic connection is defined as excitatory if $w > 0$ and as inhibitory otherwise.

Although neuronal plasticity is not the topic of this thesis, it is worth mentioning that the strength of a synapse is not static, but evolves over time. In particular, the amplitude w of a PSC is subject to long-lasting changes induced by the activity of both the presynaptic and the postsynaptic neurons, as well as by the presence of other factors such as different neuromodulators. This important phenomenon is known as synaptic plasticity and is currently thought to be the neural basis of memory formation and learning (Kandel et al., 2000). In particular, synaptic changes have been shown to depend on the precise timing of presynaptic and postsynaptic spikes (Markram et al., 1997; Bi and Poo, 1998). Amongst others, this result highlights the importance of spiking models capable of capturing the activity of single neurons with a high degree of accuracy.

1.1.5 Studying single-neuron computation *in vitro*

In vivo, cortical neurons continuously receive inputs from thousands of other neurons. The ideal goal of single-neuron computation studies is to understand how nerve cells transform arbitrary sets of incoming spike trains $\{S_i^{(\text{in})}(t)\}$ into output spike trains $S_{\text{out}}(t)$ (Fig. 1.2). In reality, given the complexity of this transformation, as well as the experimental difficulties in independently and simultaneously controlling the activity of multiple synapses (Boucsein et al., 2005; Branco et al., 2010), single-neuron computation is often studied *in vitro* using simplified paradigms based on somatic current-clamp injections¹¹. In these experimental paradigms, active dendritic processes are generally neglected (but see, e.g., Häusser et al. (2000); Larkum et al. (2009)) and single-neuron coding is investigated by studying how somatic input currents $I_{\text{syn}}(t)$ are transformed into output spike trains $S^{(\text{out})}(t)$ (Fig. 1.2). The injected currents $I_{\text{syn}}(t)$ are generally constructed in such a way as to mimic the net somatic currents resulting from dendritic integration (Fig. 1.2c).

According to Eqs. 1.10-1.11, in a realistic scenario where a neuron receives independent Poisson spike-trains from two homogeneous populations of N_I inhibitory and N_E excitatory neurons, $I_{\text{syn}}(t)$ can be approximated by a continuous Gaussian process called Ornstein-Uhlenbeck process (Van Kampen, 1992), which is defined by the following stochastic differen-

¹¹Single-neuron computation is also studied *in vitro* using a more sophisticated technique called dynamic-clamp (Destexhe and Bal, 2009). In contrast to current-clamp, dynamic-clamp makes it possible to introduce artificial synaptic conductances into biological neurons. With this technique, single neurons can be tested using *conductance-based* synaptic currents (e.g., Eq. 1.8) (Destexhe et al., 2003).

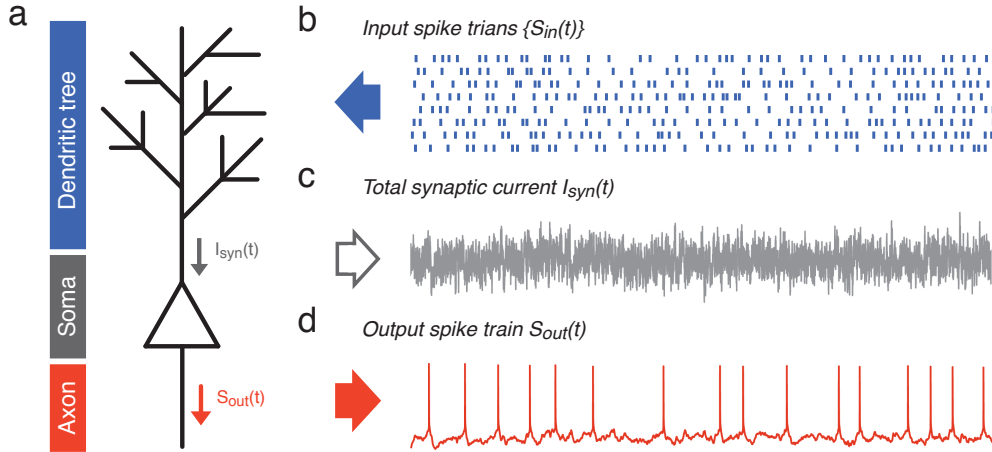


Figure 1.2: Single-neuron computation. **a)** Schematic representation of a cortical neuron divided in three functional parts: dendritic tree, soma and axon. **b)** Spiking input $\{S_{in}(t)\}$ sent by a number of presynaptic neurons and integrated at the dendritic tree of the postsynaptic neuron. Each row indicates the spike train sent by a single presynaptic neuron. **c)** Total current $I_{syn}(t)$ resulting from dendritic integration and sent to the soma. **d)** The total synaptic current $I_{syn}(t)$ is converted into an output spike train $S_{out}(t)$ and then sent through the axon to other neurons. The horizontal axes in panels *b-d* indicate time. Ideally, single-neuron computation studies aim at understanding how an arbitrary set of incoming spike trains $\{S_{in}(t)\}$ (panel *b*) is transformed into an output spike train $S_{out}(t)$ (panel *d*). In practice, single neuron computation is often investigated by studying how somatically injected currents (panel *c*) are transformed into output spike trains $S_{out}(t)$ (panel *d*).

tial equation (Tuckwell, 1988) (Fig. 1.3a):

$$\tau_{syn} \dot{I}_{syn}(t) = -I(t) + \mu_I + \sqrt{2\tau_{syn}\sigma_I} \xi(t), \quad (1.12)$$

where τ_{syn} is the characteristic timescale over which excitatory and inhibitory PSCs decay, $\xi(t)$ is a zero-mean white-noise process with unitary variance. The mean μ_I and the variance σ_I^2 of the total synaptic current $I_{syn}(t)$ are given by (Brunel, 2000):

$$\mu_I = N_E w_E \lambda_E - N_I w_I \lambda_I \quad (1.13)$$

$$\sigma_I^2 = \frac{1}{2} (N_E w_E^2 \lambda_E / \tau_E + N_I w_I^2 \lambda_I / \tau_I) \quad (1.14)$$

with w_E (and w_I) and τ_E (and τ_I) denoting the amplitude and the timescale of individual excitatory (and inhibitory) PSCs, and λ_E (and λ_I) being the average firing rates of individual excitatory (and inhibitory) neurons¹². Since the *in vivo* activity of neuronal populations is not constant, Eqs. 1.13 and 1.14 indicate that both the mean and the variance of the total synaptic input received at the soma of single neurons are likely to change over time. Moreover,

¹²Total synaptic currents described as Ornstein-Uhlenbeck processes play a crucial role in mean-field analysis of random networks (Brunel and Wang, 2003). Understanding how single neurons respond to this type of currents is therefore of great theoretical importance.




Figure 1.3: Comparison between a LIF model and a L5 pyramidal neuron responding to a set of stationary fluctuating currents with different statistics. **a)** Top: example of a fluctuating current $I_{\text{ext}}(t)$ (gray) generated according to an Ornstein-Uhlenbeck process (see Eq. 1.12). Solid and dashed black lines indicate the mean μ_I and the standard deviation σ_I of $I_{\text{ext}}(t)$, respectively. Bottom: LIF model activity evoked by the input current shown in gray. Black: membrane potential, V ; red: voltage threshold, V_T . **b)** The average firing rate of a LIF model responding to an Ornstein-Uhlenbeck current is plotted as a function of μ_I . Different colors indicate different levels of input fluctuations σ_I . At high μ_I , the LIF model loses sensitivity to input fluctuations and the output firing rate does not depend on σ_I . **c)** The average firing rate of a layer 5 pyramidal neuron responding to an Ornstein-Uhlenbeck current is plotted as a function of μ_I . Conventions are as in panel *b* (black: 50 pA; blue: 150 pA; red: 300 pA). At odds with the prediction of the LIF model, layer 5 pyramidal neurons maintain sensitivity to input fluctuations over a broad range of depolarizing offsets μ_I . Panel *c* was adapted from Arsiero et al. (2007).

single-neuron input statistics are also affected by changes in the input synchrony, changes in brain state and external stimuli. One of the main questions addressed in this thesis is to understand how and why single-neuron computation changes depending on input statistics.

Theoretical studies based on the LIF model predict that, in response to increased levels of input fluctuations (i.e., in response to increased values of σ_I), the output firing rate of a neuron rises, but only when the mean input μ_I is not sufficiently strong to evoke spikes by itself (Tuckwell, 1988) (Fig. 1.3b). In this regime, technically called *subthreshold regime*¹³, the voltage threshold for spike initiation is only reached in response to input fluctuations induced by the coincident arrival of several input spikes. Neurons operating in the subthreshold regime are therefore thought to act as coincidence detectors of synchronous spikes. In response to currents characterized by a strong mean μ_I , input fluctuations are expected to lose the ability to raise the average firing rate of a neuron, output spikes are mainly driven by μ_I and input fluctuations are simply expected to introduce little jitters in the spike train (Fig. 1.3b). In this regime, known as *suprathreshold regime*, the instantaneous firing rate of a neuron simply reflects the average strength of the input current (i.e., the net number of excitatory inputs received during an interspike interval) and neurons are thought to operate as temporal integrators. To summarize, according to the LIF model, a smooth transition from coincidence

¹³A neuron is said to operate in the subthreshold regime when the mean of $I_{\text{ext}}(t)$ is not sufficiently strong to drive the membrane potential above the firing threshold.

detection to temporal integration is expected to occur when the mean current μ_I crosses the rheobase¹⁴ (Fig. 1.3b). A long-standing debate persists over whether *in vivo* cortical neurons operate as coincidence detectors, thereby supporting a neural code in which the precise timing of individual spikes matters, or as temporal integrators, thereby supporting a rate code in which the information is simply carried by spike counts (see, e.g., Abeles (1982); Softky and Koch (1993); Shadlen and Newsome (1994); König et al. (1996); Rossant et al. (2011b); Bruno (2011); Stanley (2013); Ratté et al. (2013)).

The theoretical prediction of the LIF model (see Fig. 1.3b) has initially been confirmed by different studies performed in cortical pyramidal neurons (Chance et al., 2002; Rauch et al., 2003; Shu et al., 2003). However, more recent results from rat prefrontal cortex (Arsiero et al., 2007) and hippocampus (Fernandez et al., 2011), have demonstrated that pyramidal neurons maintain sensitivity to input fluctuations over a broad range of depolarizing currents μ_I (Fig. 1.3c). These results are of particular importance because they suggest that pyramidal neurons could act as coincidence detectors regardless of the strength of μ_I . The origins of enhanced sensitivity to input fluctuations remain unclear and are investigated in Chapter 4.

Understanding how single neurons adapt their coding strategy to different input statistics is one of the central questions of this thesis. In the past, a lot of effort has been put into understanding how early sensory systems adaptively encode external stimuli with changing statistics. The most important theoretical concepts and experimental findings obtained in the field of sensory coding are reviewed in the next section.

¹⁴That is, when the mean current μ_I becomes sufficiently strong to evoke spikes by itself.

1.2 Sensory coding and sensory adaptation

Making sense of the electrical activity generated by groups of neurons is one of the central topics of systems neuroscience. In particular, sensory neuroscience aims at understanding how sequences of action potentials represent relevant information about the external world. During the last decades, information theory (Shannon, 1949) has been successfully applied to understand several aspects of the neuronal activity of different sensory systems, such as the retina, the lateral geniculate nucleus (LGN), the motion sensitive (H1) neurons of the fly or the rodent barrel cortex.

This section reviews some fundamental concepts and the most relevant results obtained by previous studies of sensory adaptation. For readers interested in a more detailed introduction to this topic, I highly recommend the textbooks of Rieke et al. (1997) and Dayan and Abbott (2001). A good review about sensory adaptation is provided by Wark et al. (2007).

1.2.1 The efficient coding hypothesis

In 1961, Horace Barlow proposed that, given a limited capacity of transmitting information, sensory systems evolved to efficiently represent (i.e., encode) relevant stimuli of the external world. In particular, Barlow hypothesized that one of the main operations performed by early sensory pathways is to discard redundant information (Barlow, 1961). As it will be discussed, a central prediction of this theory is that the input-output transformation of a sensory system should be matched to the statistical properties of the signal that has to be encoded. This theory is nowadays known as *Barlow's efficient coding hypothesis* and has been very successful in explaining different aspects of the neural code such as sensory adaptation.

To introduce the concept of efficient coding, it is convenient to consider a simple situation in which an external stimulus s (e.g. the light intensity at a particular location of the visual field) has to be encoded by some neuronal activity r (e.g. the firing rate of a retinal ganglion cell) (Laughlin et al., 1981; Dayan and Abbott, 2001; Wark et al., 2007). In particular, it is assumed that both r and s are scalar quantities and that, at each moment in time, a new stimulus s is independently drawn from a probability distribution $P(s)$. For simplicity, a situation is considered in which the input-output transformation performed by a neural system of interest can be described by the following model:

$$r = f(s) + \eta, \tag{1.15}$$

where f is a nonlinear function of the stimulus and η is a noise term modeled as a random variable independently drawn from the noise distribution $P(\eta)$. The amount of information $I(s, r)$ that the response r carries about the stimulus s is quantified by the mutual information defined as:

$$I(s, r) = H[P(r)] - H[P(r|s)], \tag{1.16}$$

where $H[P(r)]$ is the entropy of the response (i.e. the entropy of the distribution $P(r)$) and $H[P(r|s)]$ is the noise entropy (i.e., the average entropy of the distribution $P(r|s)$, where the average is taken with respect to s), which describes the variability of the response once the stimulus is determined (Cover and Thomas, 2012).

In the absence of noise and under the assumption that f is a monotonic function of s , each stimulus corresponds to a unique response. Consequently the noise entropy vanishes and, according to Eq. 1.16, optimal coding maximizing the mutual information $I(s, r)$ is obtained by maximizing the entropy of the response $H[P(r)]$. Given a fixed dynamic range $\Delta r = r_{\max} - r_{\min}$ (i.e., assuming that r can only range between a minimal value r_{\min} and a maximal value r_{\max}), $H[P(r)]$ is maximized when $P(r)$ is uniform on the support $[r_{\min}, r_{\max}]$ (Dayan and Abbott, 2001):

$$P(r) = \begin{cases} \Delta r^{-1} & , \text{if } r \in [r_{\min}, r_{\max}] \\ 0 & , \text{otherwise.} \end{cases} \quad (1.17)$$

This theoretical result, known as *histogram equalization*, implies that the optimal way of encoding a stimulus s is to use an input-output function f with shape defined by the integral of the input distribution (Dayan and Abbott, 2001) (Fig. 1.4a):

$$f(s) = r_{\min} + \Delta r \cdot \int_{-\infty}^x P(s) ds. \quad (1.18)$$

In 1981, this theoretical principle was successfully applied by Simon Laughlin to show that information theory, and in particular the principle of maximum entropy encoding, provides a theoretical explanation for how the membrane potential fluctuations of large monopolar cells (LMC) in the visual system of the fly represent contrast fluctuations (Laughlin et al., 1981). In particular, it was shown that the input-output function of LMCs matched the theoretical prediction obtained by integrating the distribution of contrasts present in the visual scenes encountered by flies in their natural environments (Laughlin et al., 1981).

In the case of spiking neurons, constraining the maximal firing rate might not be an appropriate choice. A more reasonable constraint would for example be to impose the average of the output firing rate, which is related to the energy required to represent the stimulus¹⁵ (Attwell and Laughlin, 2001). Under this constraint, the output distribution $P(r)$ maximizing the entropy of the response is not the uniform distribution, but is the exponential distribution (Cover and Thomas, 2012).

Despite its simplicity, and regardless of the particular constraint used to determine the maximum entropy solution, the simplified scenario discussed in this section is important because it illustrates the theoretical result that the optimal coding strategy is not universal, but depends

¹⁵It has been proposed that sensory systems evolved to maximize the information transfer while minimizing energy consumption. While constraining the average firing rate goes in the right direction, the accurate study of the information-energy tradeoff probably requires more sophisticated models of energy consumption (see, e.g., Sengupta et al. (2014)).

on the input statistics (Wark et al., 2007).

1.2.2 Adaptive coding

The external world is complex and the statistics of relevant stimuli are not constant, but change over time. For example, the average intensity of visual scenes observed in different locations or at different moments in time can differ by up to nine orders of magnitude (Rushton, 1965). This means that local input statistics (i.e., the input distributions observed during restricted intervals of time or space) can strongly differ from the global distribution. As shown in the previous section, optimal coding requires a match between the input-output transformation performed by a neural system and the input distribution. Consequently, optimal coding predicts that, in response to changes in the input statistics, sensory systems should adapt their coding strategies (Wark et al., 2007). For example, if the mean of an external input increases, neurons should in principle modify their behavior in such a way as to shift their input-output function $f(s)$ horizontally (Fig. 1.4b). Similarly, a change in the input contrast (i.e., a change in the variance of the input distribution) should be accompanied by a rescaling of the input-output gain (Fig. 1.4c). The latter form of adaptation, also known as adaptive gain control, has been shown to occur in retinal ganglion cells (Smirnakis et al., 1997; Brenner et al., 2000), in motion sensitive neurons of the fly (Fairhall et al., 2001b) and in the mouse barrel cortex (Maravall et al., 2007).




Figure 1.4: Efficient coding of non-stationary signals requires sensory adaptation. **a)** Efficient coding of a stimulus s with stationary statistics by histogram equalization. Given a stimulus distribution $P(s)$ (red) and a limited dynamic range $r \in [r_{\min}, r_{\max}]$, the optimal solution maximizing the entropy of the output distribution $P(r)$ (blue) is achieved by means of a nonlinear transformation $f(s)$, whose shape is defined by the cumulative distribution of the stimulus (see Eq. 1.18). With this coding strategy, the local gain of f is large in regions where $P(s)$ is large and low where $P(s)$ is low. This way, all possible responses r between r_{\min} and r_{\max} occur with equal probability (blue). **b)** In response to an increase in the mean of the input distribution (bottom), efficient coding by histogram equalization requires an optimal system to adapt its coding strategy by shifting $f(s)$ to the right (top). **c)** In response to a decrease in the variance of the input distribution (bottom), an optimal encoder should increase its gain (top). Panel *a* was adapted from Laughlin et al. (1981). Panels *b* and *c* were adapted from Wark et al. (2007).

1.2.3 The timescale of spike-frequency adaptation

Over the last decade, inspired by the dynamic aspects of natural input statistics, sensory adaptation has been often investigated using an experimental paradigm known as the *switching experiment* (Smirnakis et al., 1997) (Fig. 1.5). In a switching experiment, the spiking activity $r(t)$ of a single neuron is recorded *in vivo* in response to a controlled sensory stimulus $s(t)$ that rapidly fluctuates over time. To assess contrast adaptation to local input statistics, the stimulus is generated according to a random process whose standard deviation is periodically switched between two values, with a cycle period T (Fig. 1.5a).

In response to a sudden increase in the input variance, sensory neurons generally increase their output firing rate, which then slowly decays until a steady-state is reached. Similarly, when the input variance is suddenly decreased, the firing rate response initially drops and then slowly recover to a new steady-state value (Smirnakis et al., 1997; Fairhall et al., 2001b) (Fig. 1.5b). This form of adaptation, known as *spike-frequency adaptation*, is ubiquitous in the central nervous system and was observed for the first time in 1926 by Adrian and Zotterman (1926)¹⁶.




Figure 1.5: Sensory adaptation by motion-sensitive neurons in the fly visual system. The spiking activity of motion sensitive neurons (H1 neurons) was recorded in response to a controlled visual stimulus consisting of a vertical bar pattern moving horizontally. H1 neurons encode the instantaneous velocity $s(t)$ at which the visual input moves across the visual field (Fairhall et al., 2001b). **a)** The external signal $s(t)$ was generated according to a white-noise process whose standard deviation periodically changed between two values, σ_{low} and σ_{high} , with cycle period T (only one cycle is shown). The same experiment was repeated four times using different cycle periods T , ranging from 5 to 40 seconds. **b)** The average firing rate of an H1 neuron is plotted as a function of time. Different gray levels show the responses to different periods T (see legend). The light gray line indicates the firing rate response evoked by a stimulus with cycle period $T = 40$ s (see panel a). **c)** Effective timescale τ_A of spike-frequency adaptation as a function of the cycle period T . The timescale of adaptation τ_A was

¹⁶In the experiments conducted by Adrian and Zotterman (1926), spike-frequency adaptation was discovered by studying the electrical activity of a stretch receptor evoked by the application of a static load.

1.2. Sensory coding and sensory adaptation

extracted from the data shown in panel *b* by fitting a single exponential function to the firing rate decay after an upward switch in σ . The linear relationship between τ_A and T is a signature of scale-free (i.e. power-law) adaptation. **d**) Same data as in panel *b*, but plotted with the time axis normalized by the cycle period T (i.e., time is shown in units of T). The responses to different periods T were also normalized to obtain a mean of zero and a standard deviation of one. The normalized responses are nearly invariant with respect to the cycle period T , indicating that spike-frequency adaptation is scale-free. **e**) The firing rate response of an H1 neuron was fitted with a LN model (see Eq. 1.23). The LN model was independently fitted four times using data observed at different moments in time relative to the cycle period (see panel *f*). In response to an external input of low variance, the input-output transformation is characterized by a steep nonlinearity $f(\bar{s})$. While responding to a signal of high variance, the gain of $f(\bar{s})$ is reduced, indicating optimal gain control. While spike-frequency adaptation lasts for several seconds, the adjustment of $f(\bar{s})$ is extremely rapid (compare the late response to σ_{high} with the early response to σ_{low}). This result suggests that spike-frequency adaptation and contrast adaptation are two different forms of adaptation. **f**) Firing rate response to a cycle period $T = 20$ s. Different gray areas indicate the data used to fit the LN models shown in panel *e*. All panels were adapted from Fairhall et al. (2001b).

In motion sensitive neurons of the fly, the effective timescale τ_A of spike-frequency adaptation was shown to scale linearly with the period T of the switching experiment (Fairhall et al., 2001a,b) (Fig. 1.5c). This result demonstrates that adaptation occurs on multiple timescales ranging from hundreds of milliseconds to tens of seconds. Moreover, the linear relationship between τ_A and T indicates that spike-frequency adaptation does not have a preferred timescale, but is characterized by a scale-invariant (i.e., power-law) process (Thorson and Biederman-Thorson, 1974; Drew and Abbott, 2006). When the data shown in Figure 1.5b are replotted by normalizing the time axis by T , recordings acquired with different cycle periods tend indeed to superimpose on each other (Fig. 1.5d). Since the seminal studies of Landgren (1952) on mammalian baroreceptors, power-law adaptation has been observed in a variety of systems such as the retina (Wark et al., 2009), the auditory cortex (Ulanovsky et al., 2003, 2004), the barrel cortex (Lundstrom et al., 2010), the mechanoreceptors of insects (French and Torkkeli, 2008) and the electrosensory system of the weakly electric fish (Xu et al., 1996). From a biophysical perspective, it remains unclear whether power-law adaptation is implemented by a single mechanism (Toib et al., 1998) or if it emerges from the interaction of multiple processes operating on different timescales (Drew and Abbott, 2006; Lundstrom et al., 2008). From a functional perspective, since many natural stimuli change over multiple timescales (Simoncelli and Olshausen, 2001), it has been proposed that power-law adaptation might be advantageous because it could confer to sensory systems the ability of matching the timescales of adaptation to the timescales of external stimuli (Fairhall et al., 2001b; Lundstrom et al., 2008).

What is the optimal timescale of adaptation? If a sensory system adapts its coding strategy too rapidly, slow input fluctuations would be incorrectly interpreted as changes in the local input statistics and the system would erroneously modify its input-output function. On the other hand, if the timescale of adaptation is too slow, the input-output transformation would not be

able to rapidly track real changes in the local input statistics and the encoded signal would not be optimal (Dunn and Rieke, 2006; Wark et al., 2007). Overall, sensory systems face the problem of finding the correct balance between adapting too rapidly and adapting too slowly (Wark et al., 2007). In a recent study, Wark et al. (2009) showed that, in retinal ganglion cells, the timescale of adaptation to changes in both mean luminescence and contrast is correctly predicted by the time that an optimal Bayesian observer would require to detect a change in the input statistics. In particular, adaptation was slow when the discriminability of the change was low and fast when the discriminability was high. Overall, this result provides a powerful theory to explain why sensory adaptation occurs on multiple timescales.

An alternative theory exists according to which spike-frequency adaptation implements a form of high-pass filtering to remove temporal correlations from the encoded signals (Dong and Atick, 1995a; Rieke et al., 1997; Fairhall and Bialek, 2002). This theory, known as temporal decorrelation (or temporal whitening), plays a central role in the results presented in Chapter 2 and is further discussed in the next section.

1.2.4 Efficient coding of time-dependent signals by temporal whitening

In the simple encoding problem discussed in Section 1.2.1, the external stimulus s was modeled as a random variable that is independently drawn at each time step. Consequently, the optimal solution was obtained by considering both s and r as scalar values and not as continuous functions of time. However, in many relevant situations, the stimuli that sensory systems have to encode are characterized by strong temporal correlations. In these situations, efficient coding theory has to be extended to time-dependent signals (Rieke et al., 1997).

In this section a scenario is considered in which a neural system of interest has to represent a time dependent signal $s(t)$ by its electrical activity $r(t)$. For simplicity, it will be assumed that the encoding process can be modeled by a linear filter $H(t)$ and by the presence of two independent sources of noise denoted $\eta_1(t)$ and $\eta_2(t)$. Mathematically, the encoding model reads (Rieke et al., 1997):

$$r(t) = \int_{-\infty}^t H(\tau) \cdot [s(t - \tau) + \eta_1(t - \tau)] d\tau + \eta_2(t). \quad (1.19)$$

Assuming that $s(t)$, $\eta_1(t)$ and $\eta_2(t)$ are well approximated by random Gaussian processes, their statistical properties are completely characterized by their power spectral densities $S(f)$, $N_1(f)$ and $N_2(f)$. Under these assumptions, the amount of information per unit of time I_{rate} (i.e., the information rate) that the response carries about the stimulus is given by (Shannon, 1949)

$$I_{\text{rate}}(s, r) = \frac{1}{2} \int_{-\infty}^{+\infty} \log_2(1 + \text{SNR}(f)) df, \quad (1.20)$$

where the signal-to-noise ratio $\text{SNR}(f)$ depends on the power spectral densities $S(f)$, $N_1(f)$

and $N_2(f)$ and on the squared amplitude response $|\hat{H}(f)|^2$ of the encoder (Rieke et al., 1997):

$$\text{SNR}(f) = \frac{|\hat{H}(f)|^2 \cdot S(f)}{|\hat{H}(f)|^2 \cdot N_1(f) + N_2(f)}. \quad (1.21)$$

At high signal-to-noise ratio, and under the constraint that the total variance of the output fluctuations is fixed (i.e., that the total energy of $r(t)$ is fixed), the optimal encoding strategy maximizing $I_{\text{rate}}(s, r)$ is such that¹⁷ (Rieke et al., 1997):

$$|\hat{H}(f)|^2 \cdot S(f) = \text{Const.} \quad (1.22)$$

This theoretical result tells us that, at high SNR, the information transfer is maximized when the filtered signal has a flat power spectrum. Consequently, optimal encoding is achieved by removing all temporal correlations from the signal that has to be transmitted. For this reason, this encoding strategy is often referred to as *temporal decorrelation* or, due to the analogy with the flat spectrum of white light, as *temporal whitening*. This result can be intuitively understood by remembering that, in the absence of noise, the mutual information between stimulus and response is the highest when the response entropy is maximal (see Eq. 1.16). Since correlations reduce the total entropy of a signal, maximizing the information transfer requires temporal decorrelation¹⁸.

The theoretical principle of efficient coding by decorrelation was originally developed in the spatial domain to explain why retinal ganglion cells encode the light intensity at the center of their receptive fields relative to their surroundings (Srinivasan et al., 1982; Atick and Redlich, 1990, 1992). Center-surround receptive fields, as well as nonlinear processing (Pitkow and Meister, 2012), have been shown to decorrelate the spiking activity of neighboring retinal cells and can therefore be understood as optimal filters *designed*¹⁹ to remove the strong spatial correlations that characterize natural visual scenes (Field, 1987). The theory of efficient coding by decorrelation was later extended to the temporal domain in order to provide a theoretical explanation for the high-pass filtering properties of neurons in the lateral geniculate nucleus²⁰ (Dong and Atick, 1995a).

A key aspect of efficient coding theory is that, given the statistical properties of natural stimuli, it makes clear predictions about how sensory systems should operate (Dong and Atick, 1995a). Natural stimuli are often characterized by scale-invariant power-spectra of the type $S(f) \propto 1/f^\alpha$ (Field, 1987; Dong and Atick, 1995a; Simoncelli and Olshausen, 2001). Consequently, at low frequencies, the signal power is extremely strong and the SNR is likely high. According

¹⁷The full derivation of this analytical result can be found in Appendix 19 of Rieke et al. (1997).

¹⁸Note, however, that in the presence of significant levels of noise, removing all temporal correlations from the encoded signal is not the optimal strategy. Redundancy can indeed improve the quality of a code by conferring robustness against noise (Shannon, 1949; Van Hateren, 1992; Van Steveninck and Laughlin, 1996; Cover and Thomas, 2012).

¹⁹I am neither a creationist nor a defender of intelligent design (Pozzorini et al., 2008).

²⁰Natural visual inputs are characterized by strong temporal correlations (Dong and Atick, 1995b; Van Hateren and Van Der Schaaf, 1996).

to efficient coding theory, sensory systems are thus expected to operate as high-pass filters $H(t)$ with squared aptitude response $|\hat{H}(f)|^2 = f^\alpha$ (see Eq. 1.22) (Rieke et al., 1997). The results presented in Chapter 2 provide evidence that this occurs in adapting neurons.

1.2.5 The timescale of contrast adaptation

In a switching experiment (see Section 1.2.3), contrast adaptation is typically assessed by fitting linear-nonlinear (LN) models (Schwartz et al., 2006) to the responses evoked by inputs with different contrast levels. In the LN model, the stimulus $s(t)$ is first passed through a linear filter $K_{\text{LN}}(t)$ and then transformed into an instantaneous firing rate $r(t)$ by a pointwise nonlinearity. Mathematically, the LN model is defined as

$$r(t) = f\left(\int_0^\infty K_{\text{LN}}(\tau)s(t-\tau)d\tau\right), \quad (1.23)$$

where the linear filter $K_{\text{LN}}(t)$ describes the feature of the stimulus to which the neuron is most sensitive²¹ and $f(\bar{s})$ maps the filtered input $\bar{s} = \int_0^\infty K_{\text{LN}}(\tau)s(t-\tau)d\tau$ into an instantaneous firing rate. Despite its simplicity, the LN model provides a fairly good description of experimental data collected in many early sensory systems²². Moreover, its parameters (i.e., the linear filter K_{LN} and the nonlinearity f) can be easily extracted from experimental data using a technique called spike-triggered analysis (Schwartz et al., 2006). Mainly for these reasons, the LN model is nowadays one of the most popular tools used to study how the activity of individual neurons is related to external stimuli. The LN model does not account for adaptation, but simply describes how rapid input fluctuations are transformed into output firing rates. However, by independently fitting the data acquired with different input statistics, the LN model allows to characterize sensory adaptation.

By estimating the steady-state nonlinear function $f(\bar{s})$ at different stimulus contrast levels σ_s , it was shown that many sensory systems can perform optimal gain control (Brenner et al., 2000; Fairhall et al., 2001b; Maravall et al., 2007, 2013). Indeed, when the nonlinear functions $f(\bar{s})$ observed in response to different σ_s are plotted as a function of the normalized filtered stimulus $\bar{s}_{\text{norm}} = \bar{s}/\sigma_s$, they tend to superimpose (see Fairhall et al. (2001b)). In a switching experiment, the temporal evolution of adaptive gain rescaling can be tracked over time by extracting the LN model parameters from different chunks of data recorded at different times relative to the stimulus switch. While Smirnakis et al. (1997) reported that the timescale τ_C of contrast gain control was similar to the timescale of spike-frequency adaptation, other studies showed that the readjustment of the input-output nonlinearity $f(\bar{s})$ requires less than 100 ms and is at least one order of magnitude faster than spike-frequency adaptation (Fairhall et al., 2001b; Baccus and Meister, 2002) (Fig. 1.5e,f). Moreover, Baccus and Meister (2002) showed that in the retina, changes in the input contrast are also accompanied by fast changes in the

²¹The filter $K_{\text{LN}}(t)$ can be interpreted as the receptive field of the neuron.

²²However, the accuracy of the model is significantly improved if spike-history effects are taken into account (Pillow et al., 2005).

1.2. Sensory coding and sensory adaptation

single-neuron receptive field $K_{LN}(t)$, which converge much earlier than rate adaptation²³. This separation of timescales suggests that spike-frequency adaptation, adaptive gain control and receptive field changes are distinct forms of adaptation, possibly controlled by independent mechanisms²⁴.

In most sensory coding studies, the spiking activity of a neuron (e.g., a retinal ganglion cell) is directly related to an external input of interest (e.g. light intensity). With this experimental paradigm, it is therefore difficult, if not impossible, to understand whether sensory adaptation requires network effects and identifying its underlying biological mechanisms is even harder. The forms of sensory adaptation discussed in the previous sections are widespread across sensory systems, which raises the appealing hypothesis that adaptation processes could result from the intrinsic properties of single neurons. Several experimental studies have shown that single neurons responding to *in vitro* current injections adapt their coding strategy to the input statistics. It remains however unclear whether and how spike-frequency adaptation, contrast gain control and receptive field changes can be simultaneously supported by the intrinsic properties of single neurons.

²³The same result was reported by Fairhall et al. (2001b), but data were not shown.

²⁴Contrast gain adaptation has been proposed to simply result from the static nonlinearity associated with the threshold process for spike generation (Sakai et al., 1995; Paninski et al., 2003), rather than from a *real* adaptation process. But see, e.g., Yu et al. (2005).

1.3 The adaptive properties of single neurons

Single neurons face a problem which closely resembles the one that sensory systems have to solve, namely to efficiently represent a continuous signal using a train of action potentials. It is indeed possible to interpret the total synaptic current $I_{\text{syn}}(t)$ as the result of single-neuron computation and the soma²⁵ as an analog-digital converter that transforms $I_{\text{syn}}(t)$ into a spike train for transmission to neighboring neurons (see Fig. 1.2). Consequently, it is reasonable to hypothesize that some of the computational principles underlying sensory adaptation also apply to single neurons.

The following sections discuss some important experimental results about single-neuron adaptation, as well as some popular spiking models that describe this phenomenon.

1.3.1 Spike-frequency adaptation in single neurons

When stimulated *in vitro* with a depolarizing current step, single neurons initially respond by emitting action potentials at a high firing rate. However, if the stimulus persists, the firing rate of many types of neurons declines until a steady-state is reached. This phenomenon has been experimentally observed in a large number of different preparations, demonstrating that spike-frequency adaptation is a common feature of spiking neurons.

Biophysically, spike-frequency adaptation can be induced by different cellular mechanisms such as (Benda and Herz, 2003): M-currents mediated by voltage-dependent high-threshold K^+ -channels (Brown and Adams, 1980) and AHP-currents mediated by Ca^{2+} -activated (Madsen and Nicoll, 1984) or Na^+ -activated (Bhattacharjee and Kaczmarek, 2005) K^+ -channels. All of these mechanisms induce spike-frequency adaptation by means of a common working-principle: the emission of an action potential rapidly triggers a cascade of biophysical events that ultimately activate a hyperpolarizing current that decays slowly. Compared to processes responsible for spike generation, the biophysical mechanisms underlying spike-frequency adaptation generally operate on much slower timescales. Consequently, the effects induced by consecutive spikes accumulate, progressively strengthening the total adaptation current and therefore decreasing the output firing rate.

Spike-frequency adaptation can also result from slow inactivation of Na^+ -channels (Fleiderovich et al., 1996; Mickus et al., 1999; Henze and Buzsaki, 2001; Benda and Herz, 2003; Jolivet et al., 2006b). After the emission of an action potential, Na^+ -channels inactivate, effectively reducing the number of channels available for the initiation of future action potentials and thus increasing the voltage threshold for spike initiation (Hodgkin and Huxley, 1952). Similar to the dynamics of channels mediating adaptation currents, recovery from inactivation can occur on relatively slow timescales such that the effect induced by consecutive spikes accumulate. As a result, in response to a step of current evoking repetitive firing, the voltage threshold for spike initiation progressively increases, producing a negative feedback process that gradually

²⁵More precisely the axon hillock, where action potentials are initiated.

1.3. The adaptive properties of single neurons




Figure 1.6: Fractional differentiation by cortical pyramidal neurons. Layer 2/3 and layer 5 pyramidal neurons from rat somatosensory cortex were tested *in vitro* with a set of six square-wave currents of different cycle periods $T \in [1, 32]$ s (panels *a-c*), as well as with a set of sinusoidal-wave currents of different cycle periods $T \in [1, 32]$ s (panels *d-f*). **a**) Typical response (black) of a pyramidal neuron to a square-wave current (gray) with $T = 2$ s. **b**) The average firing rate evoked by a set of square-wave currents of different periods $T = \{4, 8, 16\}$ s is plotted as a function of time, relative to the cycle onset. For each cycle period, upward adaptation (i.e., the firing rate decay after a sudden increase of the input current) and downward adaptation (i.e., the firing rate recovery after a sudden decrease of the input current) were independently fitted with single-exponential functions (blue). **c**) The effective timescale τ_A of both upward (blue) and downward (black) adaptation scales linearly with the cycle period T . This result indicates that spike-frequency adaptation does not have a preferred timescale, but is nearly scale-free (i.e., power-law). The effective timescale τ_A of adaptation was estimated by single-exponential fits (see panel *b*). Error bars indicate one standard deviation across 8 different neurons. **d**) Typical firing rate response (black) to a sinusoidal current of period $T = 16$ s (gray). In order to characterize the transfer function $\hat{H}(f)$ of single neurons, the same experiment was repeated multiple times by varying the period T of the sinusoidal current. **e-f**) Average amplitude response \hat{H}_A and average phase response \hat{H}_Φ of 8 neurons estimated by measuring the amplitude and the phase of the firing rate relative to the sinusoidal input. Dashed red lines indicate the best fits of a fractional differentiation model (see Eq. 1.26) of order α . Black lines with squares indicate the result obtained from 11 neurons using a different protocol in which the input current was generated according to a Gaussian process of constant mean and sinusoidally modulated standard deviation. In this case, the transfer function $\hat{H}(f)$ links the standard deviation of the input to the output firing rate. **e**) In agreement with the fractional differentiation model, single neurons act as high-pass filters $\hat{H}_A(f) \propto f^\alpha = T^{-\alpha}$. **f**) In agreement with the fractional differentiation model, the phase response $\hat{H}_\Phi(T)$ of single neurons is nearly constant, reflecting scale-invariance. All panels were adapted from Lundstrom et al. (2008).

reduces the output firing rate²⁶ (Fleiderovich et al., 1996).

²⁶In the HH model, slow Na^+ -channel inactivation is generally modeled by extending the Na^+ -current (see Eq.

To characterize the effective timescale over which spike-frequency adaptation occurs, a classical protocol consists of fitting the firing rate response induced *in vitro* by a depolarizing step of current with a single exponential function (see, e.g., Sanchez-Vives et al. (2000b)). Using this approach, spike-frequency adaptation has been shown by different studies to occur on a large variety of temporal scales, ranging from tens of millisecond up to several seconds. Other experimental studies demonstrated that to accurately capture the firing rate response of individual neurons, adaptation had to be described with two or more independent processes operating on different timescales (La Camera et al., 2006). An even more surprising result was later obtained by measuring the effective timescale of spike-frequency adaptation τ_A in cortical pyramidal neurons responding *in vitro* to a set of square-wave currents with different cycle periods T (Fig. 1.6a,b). Similar to the activity recorded *in vivo* from H1 neurons of the fly visual system (see Fig. 1.5), Lundstrom et al. (2008) demonstrated that the effective timescale over which cortical pyramidal neurons adapt their firing rate scales linearly with the characteristic timescale of the input T used in the experiment (Fig. 1.6c), indicating that spike-frequency adaptation is a scale-free process (i.e., power law).

1.3.2 Single neurons as fractional differentiators

The linear relationship between τ_A and T suggests that spike-frequency adaptation is mediated by a scale invariant process with power-law memory (Drew and Abbott, 2006). Lundstrom et al. (2008) demonstrated that this particular behavior is consistent with a model of fractional differentiation (Miller, 1995), according to which the firing rate response $r(t)$ to a current whose average $\mu_I(t)$ is slowly modulated in time²⁷ is given by:

$$r(t) = r_0 + k \cdot \frac{d^\alpha}{dt^\alpha} \mu_I(t) \quad (1.24)$$

where r_0 and k are constants and the fractional differentiation $\frac{d^\alpha}{dt^\alpha}$ of order $\alpha \in [0, 1]$ is a linear operator defined in the frequency domain as (Lundstrom et al., 2008):

$$\frac{d^\alpha}{dt^\alpha} x(t) \xrightarrow{\text{FT}} (i2\pi f)^\alpha \cdot X(f), \quad (1.25)$$

with $x(t)$ being a function of time, $X(f)$ being its Fourier transform (FT) and $\hat{H}(f) = (i2\pi f)^\alpha$ being the characteristic transfer function of a fractional differentiator²⁸. By rewriting the

1.4) with a third gating variable which is similar to h , but operates on much slower timescales (Fleiderovich et al., 1996).

²⁷The fractional differentiation model is also valid in cases where the standard deviation $\sigma_I(t)$ of the input, and not the mean, is slowly modulated in time (see. Fig. 1.6e,f). In this case, the output firing rate of the neuron is correctly described as a fractional derivative of $\sigma_I(t)$ (Lundstrom et al., 2008).

²⁸Fractional derivatives extend the concept of integer-order derivatives (e.g., $\frac{d}{dt}$, $\frac{d^2}{dt^2}$, ..., $\frac{d^n}{dt^n}$, with n being an integer number) to non-integer orders α (e.g., $\frac{d^{0.1}}{dt^{0.1}}$, $\frac{d^{0.2}}{dt^{0.2}}$, ..., $\frac{d^\alpha}{dt^\alpha}$, with α being a real number). This can be understood by noting that, in the Fourier domain, the transfer function of a differentiator of integer order n is given by $\hat{H}(f) = (i2\pi f)^n$.

transfer function $\hat{H}(f)$ in polar coordinates one obtains:

$$\hat{H}(f) = (2\pi f)^\alpha \cdot \exp\left(i\frac{\alpha\pi}{2}\right). \quad (1.26)$$

Eq. 1.26 tells us that fractional differentiators act as high-pass filters with amplitude response $\hat{H}_A(f) = (2\pi f)^\alpha$ and are characterized by a phase response $\hat{H}_\Phi(f) = \frac{\alpha\pi}{2}$ that does not depend on the frequency f of the input. By testing cortical neurons with sinusoidal currents of different frequencies²⁹ (Fig. 1.6d), Lundstrom et al. (2008) demonstrated the validity of this model and found the fractional order of differentiation that best captures the data was around $\alpha \approx 0.15$ (Fig. 1.6e-f). In response to a step of current, the fractional differentiation model predicts that the firing rate decay induced by spike-frequency adaptation follows a power-law $r(t) \propto t^{-\alpha}$ (Lundstrom et al., 2008). Since power-laws are scale-invariant functions, this result explains why the effective timescale of adaptation estimated by forcing an exponential fit to the data depends on the duration of the experiment (Drew and Abbott, 2006).

The fractional differentiation model provides a compact and elegant mathematical formulation that nicely captures multiple timescales of spike-frequency adaptation with a single parameter α . However, despite its beauty, the model *only* provides a description of the experimental data without explaining the origin of power-law adaptation. In principle, a power-law decay can be approximated with a sum of single exponential processes covering a large range of timescales (Anderson, 2001), indicating that power-law adaptation could emerge in single neurons from the combined action of different ion-channels mediating adaptation currents on different timescales (Drew and Abbott, 2006). In agreement with this hypothesis, Lundstrom et al. (2008) demonstrated that an HH model extended with several active conductances was sufficient to qualitatively reproduce the experimental data shown in Figure 1.6e-f (Lundstrom et al., 2008). Alternatively, power-law adaptation could be implemented with a single channel type that intrinsically features a scale-free dynamics (Toib et al., 1998; Gilboa et al., 2005). Also, the fractional differentiation model only accounts for firing rate modulations induced by slow changes in the input statistics and is therefore not suitable to describe the neural activity evoked by rapidly fluctuating inputs. Ideally, one would like a spiking model that, while being consistent with a fractional differentiator, also captures the occurrence of individual spikes with millisecond precision.

The functional role of power-law adaptation remains poorly understood. A previous study demonstrated that due to its high-pass filtering properties, spike-frequency adaptation over multiple timescales might improve the information transfer by removing temporal correlations from relevant signals that single-neurons have to encode (Wang et al., 2003). However, to carefully test this hypothesis, one should first characterize the statistics, and in particular the temporal correlations, of the total synaptic currents $I_{\text{syn}}(t)$ received at the soma of single neurons in biologically relevant stimulations (see Section 1.2.4). The results presented in Chapter 2 address, and partially solve, these issues.

²⁹The fractional differential model also captures the average firing rate response evoked by a fluctuating current whose standard deviation $\sigma_1(t)$ is modulated by a sum of sinusoidal functions.

1.3.3 Spiking models of spike-frequency adaptation

To accurately predict the spiking activity of single neurons, accounting for spike-frequency adaptation is of crucial importance (Jolivet et al., 2008b,a; Gerstner and Naud, 2009). In order to capture spike-frequency adaptation, the LIF model (i.e., Eq. 1.7) can be extended by including a set of $N \geq 1$ adaptation currents $A_i(t)$ that phenomenologically describe the activity of multiple channels (Izhikevich et al., 2003; Brette and Gerstner, 2005; Paninski et al., 2005; Pillow et al., 2005; Drew and Abbott, 2006):

$$C\dot{V} = -g_L(V - E_L) + I_{\text{ext}}(t) - \sum_{i=1}^N A_i(t). \quad (1.27)$$

A common approach followed in previous studies consists of modeling the dynamics of each of individual adaptation current $A_i(t)$ with a linear differential equation (Izhikevich et al., 2003; Brette and Gerstner, 2005):

$$\tau_i \dot{A}_i = -A_i + b_i \tau_i \cdot \sum_{\hat{t}_j < t} \delta(t - \hat{t}_j), \quad (1.28)$$

where $\{\hat{t}_j\}$ denotes the timing of previous spikes. According to Eq. 1.28, each time an action potential is emitted, the adaptation current $A_i(t)$ is instantaneously increased by a certain amount defined by b_i and then decays to zero with a single timescale τ_i . Since the contributions of multiple action potentials accumulate, the more the neuron spikes the more the total adaptation current $I_A(t) = \sum_i A_i(t)$ rises, delaying the emission of future spikes (Fig. 1.7a,b). By integrating Eq. 1.28 over time, it is easy to show that the total adaptation current is given by a sum of spike-triggered currents $I_A(t) = \sum_{\hat{t}_j < t} \eta(t - \hat{t}_j)$, defined as $\eta(t) = \Theta(t) \cdot \sum_{i=1}^N b_i \exp\left(-\frac{t}{\tau_i}\right)$, where $\Theta(t)$ denotes the Heaviside step function. Consequently, in a spiking neuron model defined by Eqs. 1.27-1.28, the number of timescales over which spike-frequency adaptation is modeled is directly controlled by the parameter N .

By relaxing the assumption that spike-triggered currents are generated by sums of single exponential processes, it is possible to define a more general model, according to which the membrane potential dynamics is governed by the following differential equation (Fig. 1.7a):

$$C\dot{V} = -g_L(V - E_L) + I_{\text{ext}}(t) - \sum_{\hat{t}_j < t} \eta(t - \hat{t}_j) \quad (1.29)$$

where $\eta(t)$ is now an arbitrary function of time (Paninski et al., 2005; Pillow et al., 2005). By defining the spike-triggered current as a power-law function $\eta(t) \propto (t + b)^{-a}$, one obtains the LIF model with power-law adaptation that was successfully used by Drew and Abbott (Drew and Abbott, 2006) to describe sensory adaptation, as it is observed in the weak electric fish (Xu et al., 1996). This model is particularly interesting because, similar to the fractional differentiation model of Lundstrom et al. (2008), it accounts for multiple timescales of adaptation by means of only two parameters a and b .

1.3. The adaptive properties of single neurons




Figure 1.7: LIF models extended to account for spike-frequency adaptation. **a)** Schematic representation of a LIF model extended with a spike-triggered current $\eta(t)$ (Eq. 1.29). The input current $I_{ext}(t)$ is transformed into a membrane potential V by means of a convolution with an exponential filter (leaky integration). Each time the voltage reaches the firing threshold V_T^* , an action potential is fired and an adaptation current $\eta(t)$ is triggered. **b)** Response of the model shown in panel *a* to a step of current (gray). Each spike triggers an adaptation current $\eta(t)$. Adaptation currents triggered by consecutive spikes accumulate. As a result, the total adaptation current I_A slowly increases over time (blue), progressively reducing the output firing rate. Membrane potential and constant firing threshold are shown in black and red, respectively. **c)** Schematic representation of a LIF model extended with a dynamic threshold (Eq. 1.30). This model closely resembles to the one shown in panel *a*. Each time an action potential is fired, a movement of the firing threshold V_T is triggered. The model does not feature a spike-triggered current. **d)** In response to a step of current (gray), the model shown in panel *d* features spike-frequency adaptation. Each spike triggers a movement of the firing threshold $\gamma(t)$. Similar to the dynamics of I_A in panel *b*, threshold movements induced by consecutive spikes accumulate progressively increasing the firing threshold (red). The schematic representations in panels *a* and *c* omit the voltage reset occurring after a spike.

Alternatively, spike-frequency adaptation can be accounted for by a dynamic threshold (Fig. 1.7c,d). Soon after its introduction, the original LIF model was extended by making V_T spike-dependent. In the resulting model, the voltage dynamics is exactly as in the standard LIF model (see Eq. 1.7), but the firing threshold V_T is defined as (Hill, 1936; Kobayashi et al., 2009; Jolivet et al., 2006b; Mihalaş and Niebur, 2009):

$$V_T(t) = V_T^* + \sum_{\hat{t}_j < t} \gamma(t - \hat{t}_j) \quad (1.30)$$

where V_T^* is a baseline and, similar to $\eta(t)$, $\gamma(t)$ describes the stereotypical trajectory followed by V_T after each spike. Compared to a LIF in which spike-frequency adaptation is modeled with a spike-triggered current, Eq. 1.30 is more appropriate in situations where adaptation

results from Na^+ -channel inactivation and modify the spiking activity without affecting the subthreshold dynamics of the membrane potential (Fleidervish et al., 1996).

Since in the LIF model action potentials are emitted each time the condition $V = V_T$ is verified, decreasing V by means of a spike-triggered current $\eta(t)$ or of increasing V_T with a spike-triggered movement of the firing threshold have a similar impact on the spiking activity of the model. However, to accurately capture the subthreshold dynamics of the membrane potential, the dynamics of these two forms of adaptation have to be independently modeled. In Chapters 2 and 3, a new fitting procedure is introduced that solves this problem.

1.3.4 Stochastic spike emission and Generalized Linear Model

When stimulated *in vitro* with multiple injections of the same fluctuating current $I_{\text{ext}}(t)$, cortical neurons emit quasi-reliable trains of action potentials in which the precise timing of individual spikes shows trial-to-trial variability (Mainen and Sejnowski, 1995; Jolivet et al., 2006b). This important result indicates that the spiking response of a single neuron can not be predicted with certainty and that single-neuron computation should in principle be characterized with a stochastic model determining the conditional probability of observing an output spike train $S_{\text{out}}(t)$ given an input current $I_{\text{ext}}(t)$ (Pillow, 2007). Estimating $P(S_{\text{out}}|I_{\text{ext}})$ for any possible input $I_{\text{ext}}(t)$ is obviously impossible. A better approach to tackle this problem is to first define a model $P_{\theta}(S_{\text{out}}|I_{\text{ext}})$ that could provide a good approximation of the real distribution $P(S_{\text{out}}|I_{\text{ext}})$ and then extract the model parameters θ from an available set of experimental data (Pillow, 2007).

Any integrate-and-fire model can be transformed into a probabilistic spiking model by defining that action potentials are generated according to a point process with firing intensity $\lambda(t)$ given by (Plesser and Gerstner, 2000; Gerstner and Kistler, 2002):

$$\lambda(t) = \lambda_0 \cdot \exp\left(\frac{V(t) - V_T(t)}{\Delta V}\right), \quad (1.31)$$

where the dynamics of $V(t)$ and $V_T(t)$ is deterministic, λ_0 is a constant and ΔV defines the level of stochasticity³⁰ (Fig. 1.8a). According to point process theory (Gerstner and Kistler, 2002), the probability of observing an action potential at time $\hat{t} \in [t, t + \Delta t]$ is:

$$P(\hat{t} \in [t, t + \Delta t]) = 1 - \exp\left(-\int_t^{t+\Delta t} \lambda(s) ds\right) \approx \lambda(t)\Delta t, \quad (1.32)$$

and, in the limit of $\Delta V \rightarrow 0$, spike generation becomes again deterministic and action potentials are fired whenever $V(t) = V_T(t)$. This model is commonly referred to as the *escape-rate model* (Gerstner and Kistler, 2002) and has been shown to accurately describe stochastic spike generation in cortical neurons responding to fluctuating currents (Jolivet et al., 2008a).

³⁰Alternatively, deterministic IF models can be made stochastic by corrupting the external input $I_{\text{ext}}(t)$ with a noise current $I_{\text{noise}}(t)$ typically modeled as a white-noise process (see, e.g., Pillow et al. (2005)). The resulting models are referred to as *IF models with diffusive noise* (Gerstner and Kistler, 2002).

The Generalized Linear Model (GLM) (Truccolo et al., 2005; Pillow et al., 2008) provides a simple but powerful model that simultaneously accounts for stochastic spike emission and spike-frequency adaptation (Fig. 1.8b). Similarly to an IF model equipped with the escape-rate mechanism, GLMs generate action potentials according to a point process with conditional firing intensity $\lambda_{\text{GLM}}(t)$ defined as:

$$\lambda_{\text{GLM}}(t|I_{\text{ext}}, \theta) = \lambda_0 \cdot \exp \left(\int_0^\infty K_{\text{GLM}}(s) I_{\text{ext}}(t-s) ds - \sum_{\hat{t}_j} h_{\text{GLM}}(t - \hat{t}_j) \right), \quad (1.33)$$

where λ_0 is a constant, $K_{\text{GLM}}(t)$ is a linear filter acting on the external current and defining the features of the input to which the neuron is most sensitive, $h_{\text{GLM}}(t)$ is a filter accounting for spike-history effects and $\theta = \{\lambda_0, K_{\text{GLM}}, h_{\text{GLM}}\}$ denotes the model parameters³¹. Importantly, the log-probability of an output spike-train $S_{\text{out}}(t) = \sum_{\hat{t}_j} \delta(t - \hat{t}_j)$ being generated by a GLM in response to an input current $I_{\text{ext}}(t)$ can be computed and reads (Brillinger, 1988; Pillow, 2007):

$$\log P_\theta(S_{\text{out}}|I_{\text{ext}}) = Z + \sum_{\hat{t}_j} \log \lambda_{\text{GLM}}(\hat{t}_j|I_{\text{ext}}, \theta) - \int_0^T \lambda_{\text{GLM}}(t|I_{\text{ext}}, \theta) dt \quad (1.34)$$

where T denotes the total duration of the observations and Z is a normalization constant. Moreover, since the nonlinearity used in Eq. 1.33 is at the same time convex and log-concave in its argument, the log-likelihood $\log P_\theta(S_{\text{out}}|I_{\text{ext}})$ is guaranteed to be a convex function of θ (Paninski, 2004), making Eq. 1.34 an extremely powerful tool to extract model parameters from experimental data (see Chapter 3).

The GLM does not simply provide a good model for spiking neurons responding *in vitro* to fluctuating currents (Mease et al., 2014), but has also been successfully employed to characterize the activity of individual neurons responding *in vivo* to sensory stimuli $s(t)$ (Paninski et al., 2007; Calabrese et al., 2011). In this case, the input current $I_{\text{ext}}(t)$ in Eq. 1.33 is simply replaced by $s(t)$ and $K_{\text{GLM}}(t)$ is interpreted as a receptive field³². GLMs can also be extended to model the spiking activity $\{S_i^{\text{(out)}}(t)\}_{i=1}^N$ recorded *in vivo* from a population of $N > 1$ connected neurons (Truccolo et al., 2005; Pillow et al., 2008). Indeed, if the interaction between neurons is described with an additional set of linear filters $\{\epsilon_{ji}(t)\}_{j,i}$ describing how a spike emitted by neuron i affects the spiking probability of neuron j , the likelihood function $\log P_\theta(\{S_i^{\text{(out)}}(t)\}_{i=1}^N | s(t))$ closely resembles Eq. 1.34 and remains convex in the model parameters $\theta = \{\lambda_{0,i}, K_{\text{GLM},i}, h_{\text{GLM},i}, \epsilon_{ij}\}$ (Pillow, 2007). GLMs can thus be used to analyze multi-electrode recordings and infer the functional connectivity between neurons (Okatan et al., 2005; Pillow et al., 2008; Gerhard et al., 2013).

³¹In contrast to IF models extended with the escape-rate mechanism (see Eq. 1.32), GLMs do not explicitly model the membrane potential and the firing threshold, but directly relate the input current to the spiking response. Consequently, the parameter ΔV does not appear explicitly in Eq. 1.33, but is absorbed as a scaling factor in the K_{GLM} and h_{GLM} . Also, the GLM does not feature an after-spike reset.

³²Note that the GLM can be interpreted as an LN model (see Eq. 1.23) extended with a mechanism for stochastic spike generation (i.e., a Linear-Nonlinear Poisson model) as well as with a term accounting for spike-history effects.




Figure 1.8: Escape-rate mechanism for stochastic spike generation and Generalized Linear Model.

a) Schematic representation of the escape rate model. The probability $\lambda(t)$ of emitting an action potential at time t depends on the distance between membrane potential $V(t)$ (black) and firing threshold $V_T(t)$ (red). Spikes can occur even if $V(t) < V_T(t)$. It is also possible that the membrane potential crosses the threshold without spiking. **b)** Schematic representation of the GLM. The input current $I_{\text{ext}}(t)$ is first passed through the linear filter $K_{\text{GLM}}(t)$. The resulting signal is then transformed by an exponential nonlinearity into a firing intensity $\lambda_{\text{GLM}}(t)$, according to which spikes are stochastically emitted. Each time an action potential is fired, a feedback process $h_{\text{GLM}}(t)$ is triggered that accounts for all spike-history effects (e.g., adaptation currents, threshold movements and after-spike reset). Importantly, the functional shapes of $K_{\text{GLM}}(t)$ and $h_{\text{GLM}}(t)$ are not defined *a priori* and can be extracted from experimental data using a non-parametric method. Panel *a* and panel *b* were adapted from Gerstner and Kistler (2002) and Pillow (2007), respectively.

1.3.5 Complex forms of single-neuron adaptation

In a recent study, Mease et al. (2014) used a GLM to characterize the spiking activity of cortical neurons responding to a set of fluctuating currents $I_{\text{ext}}(t)$ with different DC components μ_I . When fitted independently on different experiments, GLMs were able to capture the experimental data with great accuracy and outperformed the other models considered in that study. However, the GLM parameters extracted from different experiments were not identical, but changed systematically with μ_I . In particular, increasing the DC component of $I_{\text{ext}}(t)$ resulted in a progressive shortening of the linear filter $K_{\text{GLM}}(t)$, reflecting increased sensitivity to rapid input fluctuations. Moreover, the adaptation filter $h_{\text{GLM}}(t)$ also changed with μ_I , revealing the existence of *meta-adaptation*³³. Overall, these results demonstrate that single-neurons feature complex forms of adaptation that GLMs cannot explain.

Since the output firing rate of single neurons generally increases with μ_I , the shortening in $K_{\text{GLM}}(t)$ reported by Mease et al. (2014) could possibly be captured by replacing the spike-triggered adaptation current in Eq. 1.29 by a conductance-based current defined as (MacGre-

³³Since the spike-history filter $h_{\text{GLM}}(t)$ accounts for adaptation, changes in its functional shape can be interpreted as a form of meta-adaptation.

gor and Sharpless, 1973; Dayan and Abbott, 2001):

$$I_A(t) = \sum_{\hat{t}_j} \eta_g(t - \hat{t}_j)(V(t) - E_A), \quad (1.35)$$

where E_A is a reversal potential and $\eta_g(t)$ describes the conductance change triggered by the emission of a previous spike. Alternatively, since increasing μ_I also shifts the subthreshold voltage distribution towards more depolarized potentials, single neurons could in principle regulate their effective timescale of integration by means of a subthreshold adaptation current $I_V(t)$ governed by the following differential equation (Deemyad et al., 2012):

$$\tau_V \dot{I}_V(t) = -I_V(t) + f(V(t)), \quad (1.36)$$

where τ_V is a timescale and f is a nonlinear function of the subthreshold membrane potential.

Since GLMs do not explicitly model the subthreshold dynamics of the membrane potential, $K_{\text{GLM}}(t)$ cannot be directly interpreted as a membrane filter mapping $I_{\text{ext}}(t)$ to $V(t)$ ³⁴. It is thus possible that adaptive changes in $K_{\text{GLM}}(t)$ arise from dynamical properties of the firing threshold. In a recent theoretical study, Platkiewicz and Brette (2010) demonstrated that fast Na^+ -channel inactivation can implement a nonlinear coupling between subthreshold membrane potential $V(t)$ and firing threshold $V_T(t)$. This mechanism can be included in a LIF model by defining its firing threshold dynamics as (Platkiewicz and Brette, 2011):

$$\tau_T \dot{V}_T(t) = -(V_T(t) - V_T^*) + f(V(t)), \quad (1.37)$$

where τ_T is the timescale of fast Na^+ -channel inactivation and f is a nonlinear function of the subthreshold membrane potential $V(t)$. Similar to a nonlinear subthreshold adaptation current, the threshold equation of Platkiewicz and Brette (2011) adaptively shortens the effective timescale over which input currents are integrated. Moreover, Eq. 1.37 has recently been shown to provide a good description of the firing threshold dynamics, as it is observed *in vitro* (Higgs and Spain, 2011) and *in vivo* (Fontaine et al., 2014b).

As we have seen, during the last decades, the LIF model has been repeatedly extended with different mechanisms in order to account for single-neuron adaptation. When considered independently, these forms of adaptation are well understood. Whether and how multiple adaptation mechanisms interact remains however poorly understood. Another important problem that has not yet been solved is how to efficiently extract model parameters from experimental data when complex forms of adaptation are taken into account. In Chapter 4, a new generalized IF model is introduced that simultaneously features: the escape-rate mechanism for stochastic spike emission (Eq. 1.32), a spike-triggered conductance (Eq. 1.35), a spike-triggered movement of the firing threshold (Eq. 1.30) as well as a nonlinear coupling

³⁴The standard procedure used for GLM parameter extraction does not exploit the information provided by subthreshold voltage fluctuations, but only relies on spiking data (Pillow, 2007). Consequently, when extracted from experimental data, the linear filter $K_{\text{GLM}}(t)$ also accounts for mechanisms that change the spiking probability without affecting the membrane potential (see Chapter 3).

Chapter 1. Introduction

between membrane potential and firing threshold (Eq. 1.37). Despite its relative complexity, the model remains amenable to analytical treatment and its parameters can be efficiently extracted from intracellular recordings with a maximum-likelihood approach. As we will see, a non-trivial interaction between different adaptation mechanisms explains and unifies the phenomenon of enhanced sensitivity to input fluctuations (Arsiero et al. (2007), see Fig. 1.3) as well as the complex forms of single-neuron adaptation recently observed by Mease et al. (2014).

1.4 Thesis contribution

This thesis summarizes the main results I have obtained during my PhD at EPFL between 2010 and 2014 under the supervision of Prof. Wulfram Gerstner and in collaboration with Skander Mensi, Richard Naud, Olivier Hagens and other colleagues. The main goal of this thesis was to investigate how single neurons transform *in vivo*-like fluctuating currents into output spike trains. In particular, I focused on the question of *how* and *why* single-neurons adapt their coding strategies to the statistical properties of their inputs. To answer these questions, electrophysiological patch-clamp recordings of cortical neurons were combined with mathematical models and numerical simulations. The thesis is divided in three chapters.

In **Chapter 2**, a new Generalized Integrate-and-Fire (GIF) model is introduced to study spike-frequency adaptation over multiple timescales. The model features both a spike-triggered current $\eta(t)$ (see Eq. 1.29) and a spike-triggered movement of the firing threshold $\gamma(t)$ (see Eq. 1.30). By applying a non-parametric fitting procedure to intracellular recordings obtained from L5 pyramidal neurons, we found that both adaptation processes feature a power-law decay extending from tens of milliseconds to tens of seconds. This finding provides an explanation to a previous result of Lundstrom et al. (2008) that, in cortical pyramidal neurons, the effective timescale of spike-frequency adaptation is not fixed, but appears to scale linearly with the duration of the experiment (see Fig. 1.6). In contrast to the fractional differentiation model of Lundstrom et al. (2008) (see Section 1.3.2), the GIF model was additionally able to predict the occurrence of individual spikes with millisecond precision. In the second part of Chapter 2, the GIF model is combined with efficient coding theory and *in vivo* patch-clamp recordings to demonstrate that power-law spike-frequency adaptation is near-optimally tuned to remove temporal correlations from *natural inputs* received by single-neurons in biologically relevant situations. This result provides insights about the functional role of spike-frequency adaptation in cortical neurons.

The last years have seen a growing interest in automating electrophysiological patch-clamp recordings. To make sense of the large amount of data that emergent technologies will make available in the near future, adequate computational tools are required. In **Chapter 3**, the GIF model and the fitting procedure introduced in Chapter 2 are proposed as a powerful tool to perform high-throughput single-neuron characterization based on automated patch-clamp recordings. By extracting GIF model parameters from *in silico* data generated by the GIF model itself, from *in silico* data generated using the multi-compartmental model of Hay et al. (2011) and from *in vitro* patch-clamp recordings performed in L5 pyramidal neurons, we demonstrate that a short experimental protocol of around 5 minutes is sufficient to characterize the electrical properties of single neurons. This chapter also provides a concise review of simplified threshold models as well as of existing methods to extract model parameters from intracellular recordings.

Chapter 4 focuses on complex forms of single-neuron adaptation and on the question of whether cortical neurons operate as temporal integrators or coincidence detectors. A new set

of *in vitro* patch-clamp recordings is introduced, in which cortical neurons are tested with fluctuating currents generated by systematically varying the intensity of the mean current μ_I and the amplitude of rapid input fluctuations σ_I (see Eq. 1.12). In agreement with previous results (see Fig. 1.3), we found that the average firing rate of L5 pyramidal neurons always increased with σ_I , even in the case of strong depolarizing offsets μ_I . Explaining why pyramidal neurons maintain sensitivity to rapid input fluctuations, a GLM-based data analysis revealed that the effective timescale over which the input current is integrated progressively shortened with increasing μ_I . To identify the computational principles underlying this form of adaptation, the GIF model introduced in Chapter 2 was extended with a dynamic coupling between firing threshold and subthreshold membrane potential (see Eq. 1.37). In the following, we refer to this new model as the inactivating Generalized Integrate-and-Fire (iGIF) model. Results obtained by extracting iGIF model parameters from *in vitro* recordings indicate that, in L5 pyramidal neurons, the voltage threshold for spike initiation depends nonlinearly on the subthreshold membrane potential and linearly on previous spikes. This result is consistent with a previous biophysical model in which Na^+ -channel inactivation is accounted for with both a fast and a slow gating variable (Fleidervish et al., 1996; Platkiewicz and Brette, 2011). The analysis of the iGIF model's dynamics further revealed that a non-trivial interaction between spike-dependent and voltage-dependent threshold movements adaptively regulates the effective timescale of somatic integration. Overall, the results reported in this chapter demonstrate that the firing threshold dynamics implements a complex form of adaptation, which in turn makes single neurons act as coincidence detectors over a broad range of input statistics. These findings indicate that cortical pyramidal neurons are well-suited to support a temporal code in which the precise timing of individual spikes conveys relevant information.

My contribution to each of these findings is summarized in the next section and at the end of each chapter in a specific section called *Author contributions*.

1.5 Personal contribution

The work presented in **Chapter 2** was done in collaboration with Richard Naud, Skander Mensi, Shovan Naskar, Carl Petersen and Wulfram Gerstner. During my master thesis (at the Laboratory of Computational Neuroscience, EPFL), Richard Naud gave me the original idea of explaining scale-free spike-frequency adaptation (Lundstrom et al., 2008) by means of a power-law spike-triggered current that could possibly be extracted from experimental data using mathematical methods. In a parallel project, Skander Mensi was working on deterministic spiking models and in particular on a linear regression method to extract the time course of adaptation currents from intracellular voltage recordings. During the first year of my PhD, I extended the method of Skander Mensi by making the spiking model stochastic with the final goal of extracting from data the dynamics of the firing threshold (maximum likelihood method). The fitting procedure obtained by combining the method of Skander Mensi (linear regression on subthreshold voltage fluctuations) and mine (maximum likelihood on spiking response) was first published as a Methods paper in *Journal of Neurophysiology* in 2011 (see Mensi et al. (2011)). To obtain a reliable estimate of the adaptation current up to several seconds after spike generation, the accuracy of the linear regression method originally published in Mensi et al. (2011) was further improved by extending the model with an hidden current aimed at absorbing experiential drifts. This improvement was an original idea of Skander Mensi, which I personally implemented. All of the *in vitro* recordings reported in Chapter 2 were designed by me and performed by Shovan Naskar in the laboratory of Carl Petersen (EPFL) using a graphical user interface that I personally developed to execute and calibrate the whole-cell current-clamp injections in a semi-automated way. During the early phase of the *in vitro* experiments, I assisted Shovan Naskar with the goal of gaining a better understanding of the experimental setup and to improve the protocol. Carl Petersen supervised the experimental side of the project. The hypothesis that power-law adaptation could implement a form of efficient coding by temporal whitening of the natural inputs received by cortical neurons in biologically relevant situations originated from several discussions between Richard Naud and me. The *in vivo* recordings that have been used to verify this hypothesis were kindly provided by Sylvain Crochet and were acquired in the laboratory of Carl Petersen. This dataset was previously published by Crochet et al. (2011). The spectral analysis of the *in vivo* recordings has been originally performed by Skander Mensi and redone by me before final publication. All of the data preprocessing (i.e., Active Electrode Compensation), numerical simulations, model fitting and data analysis reported in Chapter 2 were implemented and performed by me. I produced all of the figures and I wrote most of the text. Wulfram Gerstner and Richard Naud contributed to the writing of the manuscript and supervised the project.

The work presented in **Chapter 3** was done in collaboration with Skander Mensi, Olivier Hagens, Richard Naud, Christof Koch and Wulfram Gerstner. The original idea of combining the mathematical methods presented in Mensi et al. (2011) and in Pozzorini et al. (2013) with high-throughput patch-clamp recordings originated from a discussion between Wulfram

Chapter 1. Introduction

Gerstner, Christof Koch and me. The study was designed by Skander Mensi, Christof Koch, Wulfram Gerstner and me. In particular Christof Koch proposed to assess the performance of our GIF model using artificial data generated by simulating the multi-compartmental model of Hay et al. (2011). Skander Mensi performed all of the numerical simulations and the data analysis reported in the manuscript (except for Active Electrode Compensation that was done by me). The structure of the manuscript was conceived by me with the help of Skander Mensi and Wulfram Gerstner. I wrote the initial draft and Skander Mensi produced all the figures (except for the schematic drawings appearing in Boxes 1-4, which were done by me). Christof Koch, Wulfram Gerstner, Skander Mensi and Richard Naud contributed to the writing of the manuscript by correcting and improving the original draft. A previous version of this manuscript appeared in the thesis of Skander Mensi, was later improved and is currently under review at PLOS Computational Biology. In Chapter 3 of this thesis, the original manuscript has been modified and extended by me with the main goal of making the article accessible to a broader audience, including experimentalists. The structure of the paper has been modified, two new sections have been written by me to review previous efforts in designing and fitting simplified spiking neuron models to *in vitro* electrophysiological recordings, a new section was written by Richard Naud to discuss the limitations of point-neuron models and four boxes have been written by me to make the technical aspects of the manuscript accessible to a broad audience. The improved manuscript was corrected and improved by Wulfram Gerstner and Christof Koch. Since Skander Mensi has been less involved in this second part of the project, Skander Mensi and I took the decision to modify the order of the authors list by putting my name as the first coauthor.

The work presented in **Chapter 4** was done in collaboration with Skander Mensi, Olivier Hagens and Wulfram Gerstner. I had the original idea of explaining enhanced sensitivity to input fluctuations (Arsiero et al., 2007) by extending the GIF model with a nonlinear coupling between firing threshold and membrane potential (Brette, 2011) that could possibly be extracted from *in vitro* intracellular recordings. I also had the original idea of linking our results to previous findings about sensory adaptation (Fairhall et al., 2001b; Baccus and Meister, 2002). During the exploratory phase of the project, Skander Mensi and I worked in parallel to test different models, fitting strategies and theories. More precisely, Skander Mensi mostly focused on the fitting procedure and model validation. I mostly worked on data analysis (i.e., f-I curves and model-free firing threshold estimation), protocol design and model reduction. Based on the results obtained during this exploratory phase, Skander Mensi and I designed the figures and the structure of the paper, Skander Mensi redid all of the fits, model validation, model reduction and produced Figures 4.2, 4.5-4.8. In parallel, I redid the data analysis, produced Figures 4.1, 4.3, 4.4 and entirely wrote the manuscript. All of the *in vitro* experiments included in the manuscript were performed by Olivier Hagens in the lab of Henry Markram. During the exploratory phase of the project, I assisted Olivier Hagens to improve the protocol design. Wulfram Gerstner supervised the project. A previous version of this manuscript appeared in the thesis of Skander Mensi. At this stage of the project, the main goal of the study was to

understand the differences between pyramidal and fast spiking neurons. The main results supporting our early conclusions were therefore provided by model fitting and validation, in which Skander Mensi was clearly more involved than me. For this reason, Skander Mensi appears in his thesis as the first coauthor of the manuscript (this decision was taken with my approval). In the late phase of the project, Skander Mensi and I decided to focus more in depth on pyramidal neurons and the results obtained from fast spiking neurons were removed from the manuscript. In the version of the manuscript reported in Chapter 4, the central aspect of the study is to understand and explain the differences between the membrane timescales and effective timescale of somatic integration (which turns out to be affected by the firing threshold dynamics). Since I was more involved in the theoretical aspect of the project and in understanding the functional implications of our results, Skander and I took the decision to modify the order of the authors list by putting my name as the first coauthor.

2 Temporal whitening by power-law adaptation in neocortical neurons

Christian Pozzorini, Richard Naud, Skander Mensi and Wulfram Gerstner (published in Nature Neuroscience)

Abstract

Spike-frequency adaptation is widespread in the central nervous system, but its functional role remains unclear. In neocortical pyramidal neurons, adaptation manifests itself by an increase in neuronal firing threshold and by adaptation currents triggered after each spike. Combining electrophysiological recordings with modeling, we found that in mice these adaptation processes last for more than 20 seconds and decay over multiple time scales according to a power-law. The power-law decay associated with adaptation mirrors and cancels the temporal correlations of input current received *in vivo* at the soma of L2/3 somatosensory pyramidal neurons. These findings suggest that, in the cortex, spike-frequency adaptation causes temporal decorrelation of output spikes (*temporal whitening*), an energy efficient coding procedure that, at high signal-to-noise ratio, improves the information transfer.

2.1 Introduction

Neural signaling requires a large amount of metabolic energy (Attwell and Laughlin, 2001). Consequently, neurons are thought to communicate using efficient codes in which redundant information is discarded (Laughlin, 2001). Theories of efficient coding (Barlow, 1961) successfully predict several features of sensory systems. At early stages of visual processing, inputs coming from the external world are *decorrelated* both in space and time (Srinivasan et al., 1982; Dong and Atick, 1995a; Dan et al., 1996; Pitkow and Meister, 2012); through sensory adaptation (Wark et al., 2007), codes are dynamically modified so as to maximize information transmission (Wainwright, 1999; Brenner et al., 2000; Fairhall et al., 2001b; Maravall et al., 2007); and sensory adaptation on multiple timescales (Fairhall et al., 2001b; Baccus and Meister, 2002; Ulanovsky et al., 2004) could possibly reflect the statistics of the external world (Simoncelli and Olshausen, 2001).

Sensory adaptation is at least partially due to intrinsic properties of individual neurons and, in particular, to spike-frequency adaptation (SFA). SFA is not only observed at the early stages of sensory processing, but is also widespread in cortical neurons embedded in highly recurrent networks. Often modeled by a single process with one specific timescale (Izhikevich et al., 2003; Brette and Gerstner, 2005), SFA also occurs on multiple timescales (Spain and Schwindt, 1991; Gilboa et al., 2005; La Camera et al., 2006). In pyramidal neurons of the rat somatosensory cortex, three or more processing steps away from the sensory receptors, SFA is scale-free (Lundstrom et al., 2008), meaning that the effective speed at which individual neurons adapt is not fixed but depends on the input. Scale-free adaptation can be captured by simple threshold models with a power-law decaying spike-triggered process (Drew and Abbott, 2006) that possibly describes the combined action of Na⁺-channel inactivation (Fleidervish et al., 1996; Mickus et al., 1999; Melnick et al., 2004) and ionic channels mediating adaptation currents (Madison and Nicoll, 1984; Schwindt et al., 1989; Sanchez-Vives et al., 2000a).

Three questions therefore arise: First, can the temporal features of spike-triggered currents and spike-triggered changes in the firing threshold, possibly spanning multiple timescales, directly be extracted from experimental data? Second, can SFA be explained by these spike-triggered effects? And finally, do the timescales of SFA match the temporal statistics of the inputs received by individual neurons? If temporal characteristics of inputs and SFA were matched, SFA could lead to a perfect decorrelation of the information contained in one spike with that of the previous one of the same neuron, a phenomenon called *temporal whitening* (Wang et al., 2003). Temporal whitening in turn implies that, at high signal-to-noise ratio, information transmission is enhanced (Rieke et al., 1997).

2.2 Results

The question of whether SFA is optimally designed for efficient coding can only be addressed if both the dynamics of SFA and the statistical properties of the inputs generated in biologically relevant situations are known. Therefore, the Results section is organized as follow. We start with a combined theoretical and experimental approach so as to extract the dynamics of spike-triggered processes and SFA directly from *in vitro* recordings of cortical neurons. Then, we analyze the synaptically driven membrane potential dynamics recorded *in vivo* from somatosensory neurons during active whisker sensation (data from Crochet et al. (2011)). Our overall goal is to study whether adaptation optimally removes the temporal correlations in the input to single neurons embedded in the highly recurrent network of the cortex.

2.2.1 SFA is mediated by two power-law spike-triggered processes

To reveal adaptation on multiple timescales, we stimulated L5 somatosensory pyramidal neurons with sinusoidal noisy currents of period T (see Materials and Methods) chosen between 500 ms and 16 s (Fig. 2.1). Single neurons responded with a firing rate $r(t)$ characterized by

fast fluctuations around a sinusoidal mean $r_{\text{mean}}(t)$ given by the first-order approximation

$$r_{\text{mean}}(t) = r_0 + \hat{H}_A(w) \cdot \Delta I_{\text{mean}} \sin(\omega t + \hat{H}_\Phi(w)), \quad (2.1)$$

where $w = 2\pi/T$ is the angular frequency of the input modulation, $r_0 \approx 4$ Hz is the average firing rate, $\hat{H}_A(w)$ is the amplitude response and $\hat{H}_\Phi(w)$ is the phase response. In the Fourier domain, the transfer function $\hat{H}(w) = \hat{H}_A(w)e^{i\hat{H}_\Phi(w)}$ constitutes a linear model for the modulation of the output firing rate (Fig. 2.1).




Figure 2.1: Experimental protocol and spiking neuron model GLIF- ξ . To reveal SFA on multiple timescales, we repeatedly stimulated synaptically isolated L5 pyramidal neurons (*PYR neuron*) with fluctuating currents (*Input current*) generated by adding filtered Gaussian noise to sinusoidal waves with different angular frequencies $w = 2\pi/T$ (*Mean modulation*). The horizontal bars (bottom left and right) indicate the period T of modulation. The single neuron response (*Spiking response*, black) was recorded intracellularly and the firing rate $r(t)$ was estimated by counting the number of spikes in every time bin (*Firing rate modulation*, gray). The periodic oscillations of the firing rate $r_{\text{mean}}(t)$ (*Firing rate modulation*, black) was related to the mean input (*Firing rate modulation*, light gray) with a linear rate model defined in the Fourier domain by the the transfer function $\hat{H}(w)$. We then used the intracellular recordings to fit the Generalized Leaky Integrate & Fire model GLIF- ξ (top). In this model, the input current is first low-pass filtered by the membrane filter $K_m(t)$ and then transformed into a firing intensity by an exponential nonlinearity. Spikes are emitted stochastically (*Spiking response*, red) and trigger an adaptation process described by the effective adaptation kernel $\xi(s)$.

Since SFA is at least partly due to spike-triggered effects, the simple firing rate picture of Equation 2.1 must be complemented by a spike-based description. We therefore used intracellular recordings to fit a generalized leaky integrate-and-fire model (GLIF- ξ) with escape-rate noise (Gerstner and Kistler, 2002) for stochastic spike emission (Fig. 2.1). To capture spike-triggered

Chapter 2. Temporal whitening by power-law adaptation in neocortical neurons

adaptation, the model features an *effective* dynamic threshold, described by the function $\xi(s)$. This function (also called effective adaptation *filter* or *kernel*) summarizes the stereotypical sequence of biophysical events triggered by the emission of an action potential and accounts for both adaptation currents and physiological changes of the firing threshold. Since the effects induced by consecutive spikes accumulate, the effective dynamic threshold produces SFA. Importantly, the functional shape of $\xi(s)$, like all the other parameters of the model, were extracted from the data (see Materials and Methods).

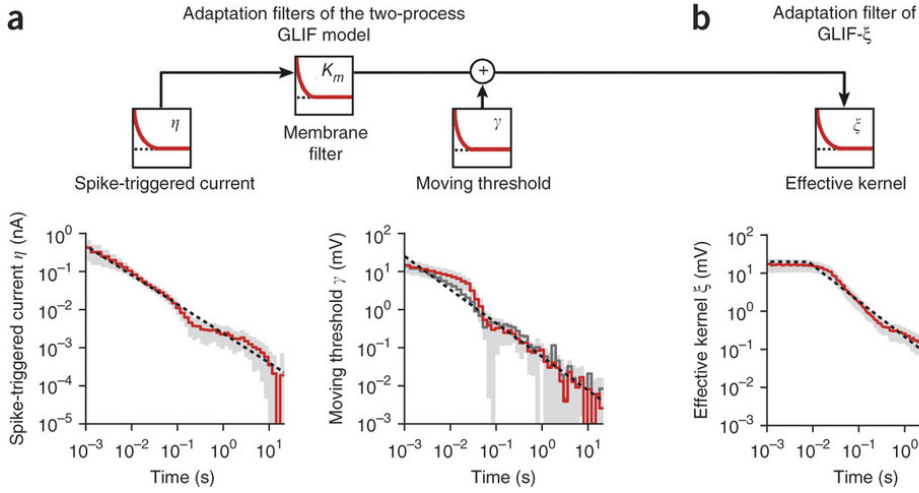


Figure 2.2: Adaptation filter of the GLIF- ξ model extracted from *in vitro* recordings. **a)** Adaptation filters of a two-process GLIF model that accounts for SFA with both a spike-triggered current $\eta(s)$ and a spike-triggered movement of the firing threshold $\gamma(s)$. Left: Mean spike-triggered current $\eta(s)$ (red) obtained by averaging the results of different cells ($n = 14$). The dashed black line shows the fit of a power-law function $\eta_{PL}(s) = \alpha_\eta s^{-\beta_\eta}$ with parameters $\alpha_\eta = 0.44$ nA, $\beta_\eta = 0.76$ and s is in milliseconds. Right: Mean moving threshold $\gamma(s)$ (red) obtained by averaging the results of different cells ($n = 14$). The dashed black line shows the fit of a power-law function $\gamma_{PL}(s) = \alpha_\gamma s^{-\beta_\gamma}$ with parameters $\alpha_\gamma = 24.4$ mV, $\beta_\gamma = 0.87$ and s in milliseconds. The dark gray line is a control showing an independent estimation of the average moving threshold $\gamma(t)$ obtained with an alternative fitting procedure (see Methods). **b)** The spike-triggered current $\eta(s)$ and the moving threshold $\gamma(s)$ were combined (block diagram) to obtain the effective adaptation filter $\xi(s)$ of the GLIF- ξ model. The mean adaptation filter $\xi_L(s)$ (red, GLIF- ξ_L) obtained by averaging the effective spike-triggered adaptation measured in individual cells ($n = 14$, see Fig. 2.3) is shown in red. The optimal fit of a truncated power-law $\xi_{PL}(s)$ (dashed black, GLIF- ξ_{PL}) yields an exponent $\beta_\xi = 0.93$ (c.f. Eq. 2.2). In all panels, the gray area indicates one standard deviation for the distribution of filters across different cells (asymmetric errors are due to log-scales).

As previously reported by Mensi et al. (2011), neocortical pyramidal neurons adapt their firing rates by means of two distinct biophysical mechanisms that respectively increase the firing threshold and lower the membrane potential after each spike. To get an accurate estimation of the effective adaptation filter $\xi(s)$, we first fitted a two-process GLIF model that explicitly features both a dynamic threshold and an adaptation current, described by the filters $\gamma(s)$

Term	Description	Mean	S.d.	Units
R	Cell resistance	93.2	39.2	$M\Omega$
τ_m	Membrane timescale	15.3	7.1	ms
E_L	Reversal potential	-69.4	6.2	mV
V_T^*	Firing threshold baseline	-51.9	5.4	mV
ΔV	Firing threshold sharpness	0.75	0.15	mV
α_ξ	Magnitude of the effective adaptation filter ξ_{PL}^\dagger	19.42	5.72	mV
β_ξ	Scaling exponent of the effective adaptation filter ξ_{PL}^\dagger	0.90	0.17	-
T_ξ	Cutoff of the effective adaptation filter ξ_{PL}^\dagger	8.05	4.12	ms
T_{ref}	Absolute refractory period	2.0	-	ms
V_r	Reset Potential	-38.8	9.0	mV

Table 2.1: GLIF- ξ_{PL} model parameters extracted from of (n=14) L5 pyramidal neurons.

[†] The parameters obtained by fitting the average kernel shown in Fig. 2.2b are: $\alpha_\xi=19.3$ mV, $\beta_\xi=0.93$ and $T_\xi=8.3$.

and $\eta(s)$, respectively (see Mensi et al. (2011) and Materials and Methods). Since in the model the emission of action potentials only depends on the difference between the membrane potential and the firing threshold, spike-triggered currents $\eta(s)$ and movements of the firing threshold $\gamma(s)$ were then combined to obtain the effective adaptation filter $\xi(s)$ of the more parsimonious model GLIF- ξ (see Materials and Methods and Fig. 2.2).

We found that 22 seconds after the emission of an action potential a small but significant deflection remained in both the spike-triggered current $\eta(t)$ and the moving threshold $\gamma(t)$. Moreover, when displayed on log-log scales, the decay of both adaptation kernels was approximately linear over four orders of magnitude, meaning that both the adaptation current and the moving threshold are characterized by scale-free spike-triggered dynamics (Fig. 2.2a). Fitting $\eta(t)$ and $\gamma(t)$ with a power-law function $f_{\text{PL}}(t) = \alpha_f t^{-\beta_f}$, revealed that both spike-triggered processes have similar scaling exponents ($\beta_\eta = 0.76$, $\beta_\gamma = 0.87$). Consequently, the effective adaptation filter $\xi(t)$ is well described by a truncated power-law

$$\xi_{\text{PL}}(t) = \begin{cases} \alpha_\xi \cdot \left(\frac{t}{T_\xi}\right)^{-\beta_\xi} & \text{if } t > T_\xi \\ \alpha_\xi & \text{if } 0 < t \leq T_\xi \end{cases} \quad (2.2)$$

with parameters $\alpha_\xi = 19.2$ mV, $\beta_\xi = 0.93$ and $T_\xi = 8.3$ ms for the average kernel (Fig. 2.2b) and slightly different values for each individual cell (Fig. 2.3), indicating that scale-free SFA is an intrinsic property of individual neurons and not only of the average over several cells.

In the following, we will refer to a model with a single spike-triggered adaptation filter as GLIF- ξ_L , where GLIF stands for *Generalized Leaky Integrate-and-Fire* and ξ_L indicates that SFA is implemented by a 22-second *Long* filter obtained by combining the moving threshold and the spike-triggered current extracted from the experimental data. With the same logic, we denote GLIF- ξ_{PL} a model in which the effective filter $\xi(s)$ is described by the truncated power law ξ_{PL} defined by Equation 2.2. All the GLIF- ξ_{PL} model parameters are listed in Table 2.1.




Figure 2.3: Effective adaptation filters of individual L5 pyramidal neurons. The 14 effective filters $\xi(t)$ measured in individual cells were fitted with the truncated power-law function $\xi_{PL}(t)$ (Eq. 2.2). **a)** Distribution of magnitudes α_ξ measured in different cells. **b)** Distribution of scaling exponents β_ξ measured in different cells. **c)** Distribution of cutoff values T_ξ measured in different cells. **d)** Effective filters $\xi(t)$ extracted from different cells (red) with optimal truncated power-law fit (dashed black). Each subpanel corresponds to a different cell.

2.2.2 Power-law SFA explains neural activity on short timescales

Valid single neuron models should predict the occurrence of individual spikes with millisecond precision (Jolivet et al., 2008b). In response to a single injection of a fluctuating current (Fig. 2.4a) the neuron emitted spikes that the GLIF- ξ_L model was able to predict with a high degree of accuracy (Fig. 2.4b, red). When the same current was injected repetitively, the spiking responses revealed the stochastic nature of single neurons: certain action potentials were emitted reliably with a high temporal precision, while others did not occur at each repetition or were characterized by larger temporal jitters. The GLIF- ξ_L model also captured this aspect (Fig. 2.4c). To validate our model, we quantified its predictive power using a similarity measure denoted M_d^* (see Materials and Methods and Naud et al. (2011)). On average, GLIF- ξ_L was able to predict more than 80% of the spikes ($M_d^* = 0.807$, s.d. = 0.04) with a precision of ± 4 ms. Very

similar performances, statistically not different ($n=12$ cells, paired t -test, $t_{11} = 0.30$, $p=0.77$), were achieved by GLIF- ξ_{PL} ($M_d^* = 0.804$, s.d. = 0.05), indicating that spike-triggered processes are well described by a truncated power-law.

As expected, the subthreshold response observed *in vitro* was systematically overestimated by GLIF- ξ_L (Fig. 2.4b, red). This is explained by the fact that GLIF- ξ_L artificially translates spike-triggered currents into effective threshold movements. In a two-process GLIF model, where adaptation currents and threshold movements are described as two distinct features (i.e., when $\eta(t)$ and $\gamma(t)$ are not combined in a single effective kernel), model prediction of the membrane voltage and experimental data were indeed in good agreement (Fig. 2.4b, gray), confirming the validity of our fitting procedure. In terms of mere spike-timing prediction, the two-process GLIF model and the more parsimonious GLIF- ξ model are equivalent (Fig. 2.4c). For this reason, we work in the following with single-process model GLIF- ξ .

Overall, the spike time prediction paradigm demonstrates the ability of both GLIF- ξ_L and GLIF- ξ_{PL} to capture the spiking activity of single neurons on the timescale of milliseconds.




Figure 2.4: The GLIF- ξ model predicts the occurrence of single spikes with millisecond precision.

a) Typical 2.5-second segment of injected current. The same fluctuating current is presented several times (*frozen-noise*). The dashed black line represents 0 nA. **b)** The spiking response, but not the subthreshold membrane potential predicted by the GLIF- ξ_L model (red) is in close agreement with the experimental data (black). In the two-process GLIF model (gray), where spike-triggered currents and threshold movements are modeled by two distinct processes (i.e., $\gamma(s)$ and $\eta(s)$), the dynamics of the subthreshold membrane potential predicted by the model is in excellent agreement with the experimental data. Inset: comparison of subthreshold membrane potential (scale bars: 40 ms, 5 mV). **c)** The raster plots show the spiking response of both the neuron (black) and the GLIF- ξ_L model (red) to repetitive presentation of the same current.

2.2.3 Power-law SFA explains neural activity on long timescales

We wondered whether the 22-second long adaptation filter $\xi(s)$ could also predict the firing rate modulation on the much slower timescale of seconds. To this end, we used the GLIF- ξ_L model fitted on responses to different frequencies of modulation ($0.5 \leq T \leq 16$ s) to predict the firing rates recorded in the second part of the experiment, where one of the two slowest modulations ($T = 8$ or 16 s) was chosen and repetitively presented to the cell. Comparison of the raster plots obtained by injecting the same current in both the neuron and the GLIF- ξ_L model shows that the spiking activity of the real neuron closely resembles the one predicted by the model (Fig. 2.5a-c). Furthermore, the match between the running-mean PSTHs constructed for the model and the experimental data revealed that both responses share a similar phase advance (Fig. 2.5d), indicating that our GLIF- ξ_L model is sufficient to capture the characteristic signature of SFA under slow sinusoidal stimulation (Lundstrom et al., 2008).

To study the role of the 22-second adaptation filter of GLIF- ξ_L , we then fitted the same single-process model under the assumption that the adaptation filter $\xi(t)$ has a duration of only 1 second (GLIF- ξ_S , where S stands for *short* adaptation filter). Compared to GLIF- ξ_L , the firing rate predicted by GLIF- ξ_S (Fig. 2.5e, orange) was in phase with the input (Fig. 2.5e, dark gray) and not with the spike output of the cells, indicating that GLIF- ξ_S was unable to capture the slow components of SFA (i.e., the model with a short adaptation filter predicted a wrong phase advance). To provide even stronger evidence, we systematically quantified the ability to predict both the mean firing rate r_0 (Fig. 2.5f) and the phase lead \hat{H}_Φ (Fig. 2.5g). Whereas the GLIF- ξ_L model was capable of very good predictions, which are in statistical agreement with the experimental data (errors $\Delta r_0 = -0.01$ Hz, s.d. = 0.67; $n=12$ cells, Student t -test, $t_{11} = -0.04$, $p = 0.97$ and $\Delta \hat{H}_\Phi = -0.17$ deg, s.d. = 5.7; $n=12$ cells, Student t -test, $t_{11} = -0.10$, $p = 0.92$), GLIF- ξ_S had the tendency to both overestimate the average firing rate and underestimate the phase advance (errors $\Delta r_0 = 0.47$ Hz, s.d. = 0.72; $n=12$ cells, Student t -test, $t_{11} = 2.16$, $p = 0.05$ and $\Delta \hat{H}_\Phi = -17.9$ deg, s.d. = 6.5 deg; $n=12$ cells, Student t -test, $t_{11} = -9.16$, $p < 10^{-6}$), demonstrating that an adaptation filter of 1 second is not sufficient.

Finally, we measured the transfer function $\hat{H}(w)$ for both real neurons and spiking models by fitting Equation 2.1 to the firing rates observed in response to six frequencies of modulation (Fig. 2.5h-j). For both real neurons and GLIF- ξ_L , the amplitude response $\hat{H}_A(w)$ was stronger at higher frequencies compared to lower ones revealing high-pass filtering, a characteristic feature of SFA (Fig. 2.5h). Consistent with observations in L2/3 pyramidal neurons (Lundstrom et al., 2008), plotting on log-log scales the amplitude response \hat{H}_A as a function of the input frequency $f = T^{-1}$, revealed that the gain of L5 pyramidal neurons was approximatively power-law (Fig. 2.5i). Moreover, the phase response $\hat{H}_\Phi(w)$ was always positive meaning that, for all the frequencies tested in this study, the output firing rate led the input modulation (Fig. 2.5j). Overall, GLIF- ξ_L was able to capture the features of the transfer function observed in L5 pyramidal neurons. Similar results were obtained with GLIF- ξ_{PL} (Fig. 2.6), confirming that the spike-triggered processes observed *in vitro* were correctly modeled by a truncated power-law lasting 22 seconds. The experimental results reported in Figure 2.5h-j are very similar to those

obtained in L2/3 pyramidal neurons (Lundstrom et al., 2008) and provide an independent evidence for multiple timescales of adaptation.




Figure 2.5: The GLIF- ξ model accurately predicts the firing rate response on multiple timescales.

a) Input current (gray) with slow mean modulation (dark gray). **b)** Membrane potential recorded in a single trial. **c)** The firing activity (black) obtained by repetitive presentation of the same input current is compared with predictions of GLIF- ξ_L (red) and GLIF- ξ_S (orange). **d)** Data from *c* were used to build two PSTHs (black: data, red: GLIF- ξ_L). The two sinusoidal functions represent the input modulation

(dark gray) and the best fit of the experimental data (light gray). **e**) Same as in panel *d* but with the prediction of GLIF- ξ_S (orange). **f**) Performance in predicting the average firing rate r_0 of new stimuli. Left: Model predictions are plotted against experimental data. Each dot represents a different cell. Right: Each couple of open circles shows the prediction errors on the same cell. GLIF- ξ_L (red) is slightly more accurate than GLIF- ξ_S (orange) ($n = 12$, paired t -test, $t_{11} = -4.09$, $p = 0.002$). **g**) Performance in predicting the phase response \hat{H}_Φ to inputs at $T = 8$ or 16 s. GLIF- ξ_L (red) outperforms GLIF- ξ_S (orange) ($n = 12$ cells, paired t -test, $t_{11} = 6.31$, $p = 6.0 \cdot 10^{-5}$). Conventions as in panel *f*. **h**) Gain $\hat{H}_A(T)$ as a function of the period $T = 2\pi/w$. **i**) Log-log plot of the gain $\hat{H}_A(f)$ as a function of the input frequency $f = T^{-1}$. Experimental data were fitted by a power-law with scaling exponent $\beta_H = 0.12$ (dashed gray). **j**) Phase response $\hat{H}_\Phi(T)$ as a function of the period $T = 2\pi/w$. In panels *h-j*, data from individual cells ($n = 14$, gray lines) are averaged (black) and compared with the predictions of GLIF- ξ_L (red) and GLIF- ξ_S (orange). In all panels, error bars indicate one standard deviation and horizontal dashed lines indicate zero.

Overall, these results show that accounting for long-lasting spike-triggered effects with an appropriate adaptation filter is crucial to capture the response of L5 pyramidal neurons on multiple timescales.

2.2.4 Power-law SFA is optimally tuned for *temporal whitening*

Our model describes how the net current resulting from dendritic integration is encoded into a spike train at the soma of neocortical pyramidal neurons. To investigate the implications of power-law adaptation, we considered a situation in which a population of N uncoupled GLIF- ξ_{PL} neurons had to encode a common input $I(t) = I_0 + \Delta I(t)$ in the instantaneous firing rate $A(t)$, also called population activity. Note that, since the neurons in our population were all identical and received the same input, the population activity $A(t)$ is identical to the PSTH measured by repetitively injecting the same current into one single cell. For relatively small fluctuations around a mean activity A_0 , we can assume that the population operates in a linear regime and responds to an external input fluctuation $\Delta I(t)$ according to the first-order approximation

$$A(t) = A_0 + \int_0^t \Delta I(t-s)H(s)ds + n(t), \quad (2.3)$$

where the impulse response $H(s)$ is the inverse Fourier transform of $\hat{H}(w)$, the noise $n(t)$ is due to stochastic firing in a finite population and both terms depend on the intrinsic properties of the individual neurons and in particular on the precise shape of the adaptation filter $\xi(t)$.

For large populations, the noise term in Equation 2.3 becomes negligible and optimal coding is achieved by the removal of temporal correlations potentially present in the input (Atick, 1992; Rieke et al., 1997). This encoding strategy is known as *temporal whitening* and requires the population activity to have a flat power spectrum $A(f) = \text{Const}$.




Figure 2.6: GLIF- ξ_L and GLIF- ξ_{PL} have similar predictive power. **a)** The amplitude response \hat{H}_A predicted by GLIF- ξ_L (red) is compared with that of GLIF- ξ_{PL} (light blue). **b)** The phase response \hat{H}_Φ predicted by GLIF- ξ_L is compared with that of GLIF- ξ_{PL} . In panels *a* and *b*, experimental data (black and gray) and GLIF- ξ_L predictions (red) are as in Fig. 2.5h and 2.5j, respectively. **c)** Control showing that the average firing rate r_0 does not depend on the period of modulation T (c.f. Eq. 2.1). Colors are as in panel *a* and *b*. **d)** The performance of GLIF- ξ_L (red) and GLIF- ξ_{PL} (light blue) in predicting the occurrence of individual spikes with a precision of ± 4 ms was quantified using the similarity measure M_d^* . Each couple of open circles shows the performance of the two models on one cell. Performance of GLIF- ξ_{PL} ($M_d^* = 0.804$, s.d. = 0.05) were not significantly different ($n = 12$, paired t -test, $t_{11} = 0.30$, $p = 0.77$) from the ones obtained with GLIF- ξ_L ($M_d^* = 0.807$, s.d. = 0.04). **e)** Fig. 2.5f is completed with the performance of GLIF- ξ_{PL} (average error $\Delta r_0 = -0.15$ Hz, s.d. = 0.57). Predictions of GLIF- ξ_L and GLIF- ξ_{PL} were not significantly different ($n = 12$ cells, paired t -test, $t_{11} = 1.80$, $p = 0.10$). **f)** Fig. 2.5g is completed with the performance of GLIF- ξ_{PL} (average error $\Delta \hat{H}_\Phi = -4.4$ deg, s.d.=3.57). Predictions of GLIF- ξ_L and GLIF- ξ_{PL} were significantly different ($n = 12$ cells, paired t -test, $t_{11} = 2.73$, $p = 0.02$), however the difference was small. In all panels, error bars indicate one standard deviation.

SFA is known to implement high-pass filtering of the input current (Benda and Herz, 2003; Köndgen et al., 2008). In the particular case of power-law adaptation, the population response is characterized by a power-law gain (see Fig. 2.5h,i and Lundstrom et al. (2008)) suggesting that, in neocortical pyramidal neurons, spike-triggered processes might be optimally tuned to efficiently encode scale-free signals (i.e., signals that are temporally correlated across multiple timescales). However, the issue of whether the functional role of power-law adaptation is to implement *temporal whitening* can only be solved if the statistical properties of the input received *in vivo* by neocortical pyramidal neurons are known.

To this end, we analyzed the synaptically driven membrane potential dynamics recorded

Chapter 2. Temporal whitening by power-law adaptation in neocortical neurons

from somatosensory pyramidal neurons during active whisker sensation (see Materials and Methods). A spectral analysis performed on the data of Crochet et al. (2011) revealed that, at low frequencies, the power spectrum of the subthreshold membrane potential fluctuations was characterized by a power-law decay (Fig. 2.7a, red), indicating that *natural stimuli* received by somatosensory pyramidal neurons are indeed scale-free.

To provide further evidence, we simulated the activity of a population of GLIF- ξ_{PL} neurons in response to an *in vivo* like input characterized by a scale-free spectrum (Fig. 2.7a, black). The statistics of the subthreshold responses obtained in individual GLIF- ξ_{PL} neurons were consistent with the ones observed *in vivo* (Fig. 2.7b, gray). Moreover, we found that the power spectrum of the population activity $A(f)$ (Fig. 2.7b, blue) was much closer to a horizontal line than that of the input, indicating that a population of GLIF- ξ_{PL} neurons efficiently encodes *in vivo* like signals by removing temporal correlations present in the input. Similar results were obtained with a population of GLIF- ξ_L neurons, where the adaptation filter $\xi(t)$ was not an idealized power-law, but the average kernel extracted from intracellular recordings (Fig. 2.7b, gray).




Figure 2.7: Power-law adaptation is near-optimally tuned to perform *temporal whitening*. **a)** Power spectral density of the intracellular membrane potential fluctuations recorded *in vivo* from L2/3 pyramidal neurons (Voltage PSD, red). The power spectrum was computed using 20-second long recordings ($n = 57$) obtained from 7 different cells (data from Crochet et al. (2011)). Fitting a power law (not shown) on the frequency band $0.05 < f < 2$ Hz yields a scaling exponent $\beta_I = 0.67$. The power spectrum of the scale-free input used to stimulate a population of GLIF- ξ_{PL} neurons ($N = 100$) is shown in black (Input Current PSD). The power spectrum of the subthreshold response of individual GLIF- ξ_{PL} neurons (Voltage PSD, gray) is in good agreement with the *in vivo* recordings. **b)** The population activity of a group of GLIF- ξ_{PL} neurons in response to an *in vivo* like input (black, copied from panel *a*) has a nearly flat spectrum $A(f)$ (blue). Similar results were obtained with GLIF- ξ_L neurons (gray). To allow a direct comparison between input and output powers, all the spectra shown in panel *b* were normalized to have the same total power.

Overall, our results suggest that in neocortical pyramidal neurons, power-law spike-triggered adaptation mirrors and approximately cancels the temporal correlations of signals generated in a biologically relevant situation. This result provides evidence for efficient coding at the level of single neurons embedded in the highly recurrent network of the cortex.

2.3 Discussion

Neocortical pyramidal neurons are known to adapt their firing rate on multiple timescales (La Camera et al., 2006; Lundstrom et al., 2008). Here we found that SFA is due to two separable spike-triggered mechanisms: each time an action potential is fired, both an adaptation current and a movement of the firing threshold are induced. Our results show that these spike-triggered effects are surprisingly long (more than 20 s) and decay with a power-law (Fig. 2.2), highlighting the fact that SFA does not have a specific timescale. A GLIF model with an effective power-law spike-triggered process simultaneously captured both the fast dynamics critical for the prediction of individual spikes (Fig. 2.4) and the slow processes that modulate the firing rate (Fig. 2.5 and Fig. 2.6). Most importantly, we found that, in behaving mice, the currents resulting from dendritic integration and received as input at the soma of pyramidal neurons are characterized by long-range temporal correlations that are partially removed by power-law spike-triggered adaptation (Fig. 2.7). This final observation indicates that, in cortex, power-law SFA is near-optimally tuned for efficient coding.

2.3.1 Extent of spike-triggered effects

According to our results, 20 seconds after its emission, an individual spike can still affect the firing activity of a neuron. Possibly, spike-triggered effects have an even longer duration. However after 22 seconds the magnitudes of both the moving threshold and the spike-triggered current were too small to be measured by our method (for $t > 20$ s, $\eta(t) < 0.1$ pA and $\gamma(t) < 0.01$ mV, see Fig. 2.2a). Since the effects of consecutive spikes accumulate, these small contributions shaped the single neuron response in a significant way (Fig. 2.5).

Whereas power-law adaptation was necessary to capture the firing rate fluctuations, a model with spike-triggered processes that only last for 1 second (GLIF- ξ_s) achieved very high performances ($M_d^* = 0.80$, s.d. = 0.03) in predicting the occurrence of individual spikes. This fact probably explains why power-law adaptation has not been observed in previous studies in which model validation was only based on spike timing prediction.

2.3.2 Biophysical implementation of power-law adaptation

Our fitting procedure enabled us to discriminate between adaptation processes implemented by spike-triggered currents and physiological changes of the firing threshold. However, the biophysical details concerning the implementation of power-law dynamics are not part of our model. In principle, power-law relaxations can be approximated by a sum of exponentials covering a wide range of timescales (La Camera et al., 2006; Drew and Abbott, 2006). It is therefore likely that the spike-triggered current $\eta(s)$ we found results from the combined action of multiple ion-channels operating on different timescales such as Ca^{2+} -dependent, Na^+ -dependent and high-voltage activated K^+ -channels. Note, however, that a match of the relative strength of different currents implies a fine-tuned regulation of gene expression levels.

In line with this hypothesis, multiple timescales of SFA have been previously modeled by biophysical models with several channels mediating adaptation currents (La Camera et al., 2006; Wang et al., 2003; Lundstrom et al., 2008). Alternatively, scale-free dynamics could also be an intrinsic property of single channels. In particular, the power-law decay we found in the moving threshold $\gamma(s)$ might reflect the scale-free dynamics observed during Na^+ -channel deinactivation (Toib et al., 1998). In this alternative view, scale-free dynamics is likely to emerge from the presence of multiple inactivated states of ion-channels (Lowen et al., 1999; Gilboa et al., 2005).

2.3.3 How general is power-law adaptation?

All the *in vitro* results presented in this paper are from mouse layer 5. We also investigated SFA in L2/3 and obtained very similar results. In particular, we found that L2/3 pyramidal neurons adapt by means of power-law filters that closely resemble the ones observed in L5 and cause positive phase lead of the firing rate response to slow sinusoidal currents. These preliminary results suggest that L2/3 and L5 somatosensory pyramidal neurons share similar adaptation mechanisms. We also fitted GLIF models to the data of Lundstrom et al. (2008) and found that both L2/3 and L5 pyramidal neurons of the rat somatosensory cortex adapt by means of spike-triggered power-law processes indicating that this mechanism is conserved across species and could be a common feature of cortical pyramidal neurons.

2.3.4 Functional implications

Both the moving threshold and the spike-triggered current are characterized by power-law decays with very similar scaling exponents. This suggests that the particular shape of the adaptation filters plays an important role. Neural signaling consumes a large amount of metabolic energy (Attwell and Laughlin, 2001; Laughlin, 2001). The brain should therefore represent information using codes in which redundant information is discarded. According to efficient coding theory, optimality is achieved by adapting to the stimulus statistics and, at high signal-to-noise ratio (SNR), by completely removing correlations that are potentially present in the signals to be encoded (Barlow, 1961). Efficient coding theory has been used to explain neural processing at early stages of the visual system. In the retina, center-surround receptive fields coupled with nonlinear processing strongly attenuate spatial correlations of natural images (Srinivasan et al., 1982; Pitkow and Meister, 2012). Similarly, in primary visual cortex (V1), spatial decorrelation of features has been found (Simoncelli and Olshausen, 2001). In the temporal domain, neural firing was found to be decorrelated in the lateral geniculate nucleus of the cat (Dan et al., 1996) and pyramidal neurons of V1 adapt on multiple timescales, providing further temporal decorrelation (Wang et al., 2003). However, it remained unclear whether SFA serves for temporal redundancy reduction in the cortex.

To solve this issue, we estimated the statistics of input currents received *in vivo* at the soma of L2/3 somatosensory pyramidal neurons during active whisker sensation (data from Crochet

et al. (2011)). This current, which reflects spatio-temporal statistics of spike arrivals at the synapses as well as subsequent filtering in the dendritic tree, can be seen as the driving current for spike generation. We found that input currents of pyramidal neurons did not have a preferred timescale but were characterized by scale-free dynamics. Moreover, our numerical simulations indicated that power-law spike-triggered processes are near-optimally tuned to completely remove the temporal correlations revealed by the power-law decay of the input spectrum (Fig. 2.7). Overall, these results provide evidence for efficient coding in single neocortical neurons stimulated with behaviorally relevant signals.

The GLIF- ξ model implements a form of predictive coding. Indeed, the sum of adaptation processes $\xi(s)$ triggered by past spikes can be interpreted as a linear predictor of the future input. Consistent with predictive coding, further spiking only occurs when the real input exceeds the prediction. In line with our results, it has been shown that predictive coding of scale-free inputs by means of power-law spike-triggered kernels reduces the number of action potentials required to achieve a certain signal-to-noise ratio (Bohte and Rombouts, 2010).

2.3.5 Temporal Whitening vs. Noise-Shaping

For deterministic signals encoded in the absence of noise, efficient coding theory states that redundancy reduction is the optimal solution. However, in presence of noise, complete decorrelation can be detrimental. Redundancy can indeed improve the robustness of a code (Rieke et al., 1997). To assess optimal coding in small populations of neurons, the noise term $n(t)$ associated with stochastic firing (c.f. Eq. 2.3) has therefore to be considered.

Previous studies have shown that non-renewal firing activity with negatively correlated inter-spike intervals can achieve higher information rates by *noise-shaping* (Mar et al., 1999; Shin, 2001; Chacron et al., 2004). In this coding strategy, the SNR is increased in the frequency band of the input signal by transferring the effective noise power to other frequencies. As already hypothesized by Avila-Akerberg and Chacron (2011), we found that at low-frequencies, spike-triggered adaptation resulted in a reduction of noise which was completely counterbalanced by a similar modification of the gain that controls the amplitude of the signal, so that the SNR remained unchanged. Consequently, modifying the adaptation filter $\xi(s)$ did not affect the power spectrum of the effective noise (Fig. 2.8), indicating that *noise-shaping* is probably not the functional role of power-law adaptation. The question of how this result generalizes to different stimulation paradigms is out of the scope of this study.

In computational studies of memory and learning in neural networks, SFA is often neglected and, when considered, it is usually assumed to operate on short timescales. From our perspective, the power-law of spike-triggered adaptation could be helpful in bridging the gap between the millisecond timescale of spike timing and behavioral timescales. Moreover, our results suggest that power-law adaptation causes temporal decorrelation of output spikes, a procedure that, at high signal-to-noise ratio, improves information transfer.




Figure 2.8: The functional shape of the effective adaptation filter does not shape the power spectrum of the effective noise. The squared amplitude response $|\hat{H}(f)|^2$ and the noise spectrum $N(f)$ of a small population of $M = 10$ uncoupled GLIF- ξ_{PL} neurons were numerically calculated for 4 different scaling exponents $\beta_\xi = \{0.8, 1.0, 1.2, 1.4\}$ (see Eq. 2.2). The results were then used to compute the power spectrum of the effective noise defined as $N_{\text{eff}}(f) = N(f)/|\hat{H}(f)|^2$. To do so, the population activities $A_i(t)$ in response to repetitive injections of the same white-noise stimulus were simulated. Each neuron in the population received the same current. Furthermore, in all the simulations, the input current was tuned to evoke small fluctuations around a mean activity of $A_0 = 5$ Hz. Firing rates were computed by counting spikes in bins of 25 ms. **a)** The noise spectra $N(f)$ obtained for 4 different scaling exponents β_ξ are shown with 4 different colors. In the temporal domain, the noise is defined as $n_i(t) = A_i(t) - \langle A_i(t) \rangle_i$, where $A_i(t)$ denotes the population response to a single stimulation and $\langle \cdot \rangle_i$ denotes an average across repetitions i of the same injection. **b)** Squared amplitude response $|\hat{H}(f)|^2$ computed by dividing the power spectrum of the average population response $\langle A_i(t) \rangle_i$ by the constant defining the power spectrum of the white-noise input. **c)** The power spectrum of the effective noise $N_{\text{eff}}(f)$ is not affected by the scaling exponent β_ξ of the adaptation filter $\xi_{PL}(t)$. In the frequency band $0.05 < f \leq 2$ Hz (gray area) the effective noise spectrum is approximately flat, regardless of the value of the scaling exponent β_ξ . Indeed, changes in the noise spectra $N(f)$ are counterbalanced by similar changes in the amplitude response. Colors are the same in all panels and correspond to different scaling exponents β_ξ (see legend in panel c).

2.4 Materials and Methods

2.4.1 *In-vitro* electrophysiological recordings

All animal experiments were performed using published procedures (Lefort et al., 2009; Avermann et al., 2012) in accordance with the rules of the Swiss Federal Veterinary Office. Briefly, somatosensory brain slices were obtained from P14-18 Wild Type mice (C57BL6/J) and whole-cell patch-clamp recordings were performed at 35° C from L5 pyramidal neurons. The pipette solution was comprised of (in mM): 135 K-gluconate, 4 KCl, 4 Mg-ATP, 10 Na2-phosphocreatine, 0.3 Na3-GTP and 10 HEPES (pH 7.3, 290 mOsm). During the experiments, we blocked all excitatory synaptic transmission by adding CNQX (20 μ M) and D-AP5 (50 μ M) to the bath solution. All electrophysiological data were low-pass Bessel filtered at 10 kHz and digitized at 20 kHz. Measurements were not corrected for the liquid junction potential. Recordings characterized by instabilities in the action potential shape and/or large drifts in the baseline

firing rate r_0 were excluded from the dataset upon visual inspection.

2.4.2 Current Injections

To characterize single neurons with the standard tools of linear system analysis, we performed 64-s-long experiments in which noisy currents modulated by sinusoidal means were delivered in current-clamp mode. The injected current, denoted I_{ext} , was generated according to the following equation

$$I_{\text{ext}}(t) = I_0 + \Delta I_{\text{mean}} \cdot \sin\left(\frac{2\pi}{T} t\right) + \Delta I_{\text{noise}} \cdot \mathcal{N}(t) \quad (2.4)$$

where I_0 is a constant offset, ΔI_{mean} controls the amplitude of the sinusoidal mean and ΔI_{noise} defines the standard deviation of the noise. The noise $\mathcal{N}(t)$ was generated with an Ornstein-Uhlenbeck process with zero mean, unitary variance and a temporal correlation of 3 ms.

Each experiment consisted of many injections of currents generated according to Equation 2.4. In the first half of the experiment (*training set*), we performed six injections using different periods of modulation $T \in \{0.5, 1, 2, 4, 8, 16\}$ in seconds. Stimuli were delivered in random order and, for each of the six injections, a new realization of the noise $\mathcal{N}(t)$ was used. In the second part of the experiment (*test set*), one of the two slowest modulations ($T = 8$ or 16 s) was chosen and more injections were performed. To assess the reliability of single neurons, the same realization of noise $\mathcal{N}(t)$ was used (*frozen-noise*). All the injections were performed with interstimulus intervals of 1 minute.

Before and after each injection, we stimulated the neuron with two additional inputs. The first was a 2.5-s-long current composed of a hyperpolarizing step followed, after 500 ms, by a suprathreshold step. We used the response to this stimulus to identify the neuronal type (L5 burst-generating cells were not included in the dataset). The second was a 4-s long subthreshold noisy current generated with an Ornstein-Uhlenbeck process with zero mean and temporal correlation of 3 ms. We used this second injection to characterize the electrode response and perform Active Electrode Compensation (see Materials and Methods).

At the beginning of each experiment, we tuned the input parameters I_0 , ΔI_{mean} and ΔI_{noise} to obtain a firing rate r_{mean} oscillating periodically between 2 and 6 Hz. Typical values obtained after calibration were comprised in the range 100-450 pA for I_0 , 15-30 pA for ΔI_{mean} and 50-150 pA for ΔI_{noise} .

2.4.3 Linear analysis

For each neuron, we estimated the transfer function $\hat{H}(w)$ (Fig. 2.5h-j) using standard methods already used in previous studies (Lundstrom et al., 2008; Köndgen et al., 2008). Briefly, the experimental spike train $\{\hat{t}_j\}$ was built by selecting the times at which the membrane potential $V(t)$ crossed 0 mV from below. We then obtained the firing rate $r(t)$ by building a histogram

of the spike times. The bin size was chosen such that each period of modulation T was divided in 30 bins. For each input frequency $\omega = 2\pi/T$, we finally obtained the transfer function by minimizing the sum of squared errors (SSE) between the sinusoidal function $r_{\text{linear}}(t) = C_0 + C_1 \cdot \sin(\omega t + \phi)$ and the experimental firing rate $r(t)$, with $\{C_0, C_1, \phi\}$ being the only free-parameters. The transfer functions of GLIF- ξ models (Fig. 2.5h-j) were obtained with the same method.

2.4.4 Generalized Leaky Integrate-and-Fire model (GLIF- ξ)

The spiking neuron models discussed in this study are generalized leaky integrate-and-fire models equipped with a spike-triggered mechanism for SFA and with escape rate noise for stochastic spike emission (Fig. 2.1). Spikes are produced according to a point process with conditional firing intensity $\lambda(t)$ which exponentially depends on the momentary distance between the membrane potential $V(t)$ and the *effective* firing threshold $V_T(t)$ (Jolivet et al., 2006a):

$$\lambda(t) = \lambda_0 \exp\left(\frac{V(t) - V_T(t)}{\Delta V}\right), \quad (2.5)$$

where λ_0 has units of s^{-1} so that $\lambda(t)$ is in Hz and ΔV defines the sharpness of the threshold. Consequently, the probability of a spike to occur at a time $\hat{t} \in [t; t + \Delta t]$ is given by

$$P(\hat{t} \in [t; t + \Delta t]) = 1 - \exp\left(-\int_t^{t+\Delta t} \lambda(s) ds\right) \approx \lambda(t)\Delta t. \quad (2.6)$$

In the limit of $\Delta V \rightarrow 0$, the model becomes deterministic and action potentials are emitted at the moment when the membrane potential crosses the firing threshold. For finite ΔV and a membrane potential at threshold (i.e., when $V = V_T$), λ_0^{-1} defines the mean latency until a spike is emitted.

The subthreshold dynamics is modeled as a standard leaky integrator defined by the following ordinary differential equation for the membrane potential V

$$C\dot{V} = -g_L(V - E_L) + I_{\text{ext}}, \quad (2.7)$$

where the three parameters C , g_L and E_L determine the passive properties of the membrane, the *dot* denotes the temporal derivative and I_{ext} is the injected current.

The dynamics of the *effective* firing threshold $V_T(t)$ in Equation 2.5 is given by

$$V_T(t) = V_T^* + \sum_{\hat{t}_j < t} \xi(t - \hat{t}_j - T_{\text{ref}}) \quad (2.8)$$

where V_T^* is a constant, $\{\hat{t}_1, \hat{t}_2, \hat{t}_3, \dots\}$ are the times at which action potentials have been fired and $\xi(s)$ is an effective adaptation filter that accounts for all the biophysical events triggered

by the emission of an action potential. According to Equation 2.8, each time a spike is emitted, a threshold movement with stereotypical shape $\xi(s)$ is triggered, after a delay of absolute refractoriness T_{ref} . Threshold movements induced by different spikes accumulate and therefore produce SFA, if $\xi > 0$. For $s < 0$, we fixed $\xi(s) = 0$ so that only spikes in the past can affect the momentary value of the firing threshold. Importantly, the adaptation filter $\xi(s)$ also accounts for adaptation processes mediated by spike-triggered currents. Consequently, $V_T(t)$ does not describe the *physiological* threshold (i.e., the membrane potential at which action potentials are initiated *in vitro*) but has to be interpreted as a *phenomenological* model of spike-triggered adaptation. Finally, the functional shape of $\xi(s)$ was not defined *a priori* but was obtained by combining the effects of both spike-triggered currents and spike-triggered movement of the *physiological* threshold which, in turn, were extracted from the experimental data.

In principle, an absolute refractory period can be included in the adaptation kernel $\xi(s)$. However, here we prefer to work with an explicit reset after a dead time. Each time a spike is emitted the membrane potential is reset to V_T and the numerical integration is restarted after a short period of absolute refractoriness T_{ref} . The GLIF- ξ model only differs from a Generalized Linear Model (Truccolo et al., 2005; Pillow et al., 2008) due to this explicit reset.

The three GLIF- ξ models discussed in the paper differ in the duration and shape of the adaptation filter $\xi(s)$. In GLIF- ξ_L and GLIF- ξ_S , the functional shape of $\xi(s)$ is the one directly extracted from intracellular recordings. In these two models the duration of the adaptation filter is of 22 s and 1 s, respectively. In GLIF- ξ_{PL} , the adaptation filter $\xi(s)$ is modeled as a truncated power law and lasts for 22 s.

2.4.5 Data preprocessing: Active Electrode Compensation

In-vitro recordings were preprocessed to remove the bias due to the voltage drop across the recording electrode. For that, we performed Active Electrode Compensation (AEC) (Brette et al., 2008) following the procedure described in Badel et al. (2008). The electrode response was estimated *before*, *during* and *after* each 64-s long injection. Consequently, we were able to remove experimental drifts due to slow changes in the electrode properties.

Due to the voltage drop across the electrode resistance (V_e), the potential recorded by a stimulating electrode (V_{rec}) is a biased version of the real membrane potential (V):

$$V(t) = V_{\text{rec}}(t) - V_e(t). \quad (2.9)$$

In AEC, the electrode is assumed to be an arbitrary linear system operating on the timescale of a few milliseconds so that $V_e(t)$ can be modeled as a filtered version of the injected current. If the input current I_{ext} has both a DC component I_0 and a time varying component $\delta I(t)$, then the electrode potential is described by the following equation

$$V_e(t) = I_0 R_e(t) + \int_0^t K_e(s, t) \delta I(t-s) ds, \quad (2.10)$$

Chapter 2. Temporal whitening by power-law adaptation in neocortical neurons

where $K_e(s, t)$ is the electrode kernel at time t and $R_e(t) = \int_0^\infty K_e(s, t) ds$ is the electrode resistance (i.e. the *access resistance*). The argument t enables us to incorporate a potential slow drift of electrode parameters. The timescales on which electrodes operate are much faster than the timescales on which the electrode properties change. Consequently, the two terms on the right hand side of Equation 2.10 are responsible for slow-frequency and high-frequency artifacts, respectively. As it has already been shown, AEC removes high-frequency artifacts (Fig. 2.9). Here, we show that this technique also enables to compensate slow-frequency artifacts due to changes in the access resistance (Fig. 2.10).




Figure 2.9: Active Electrode Compensation removes artifacts on short timescales. **a)** Typical optimal filter $K^{(i)}(s)$ (red) obtained by averaging across bootstrap repetitions. The gray area represents one standard deviation. The tail of $K^{(i)}(s)$ is well fitted by an exponential function (dashed black line). The inset shows a zoom on the y-axis. **b)** For each bootstrap-repetition, the exponential fit from **a** (i.e. the putative membrane filter $K_m^{(i)}(s)$) is subtracted from the optimal filter to obtain an estimate of the electrode filter. The electrode filters obtained in the 15 bootstrap-repetitions are then averaged to obtain the electrode filter $K_e^{(i)}(s)$ (red) used for AEC. The gray area represents one standard deviation. Each electrode filter $K_e^{(i)}(s)$ is characterized by its timescale (estimated by fitting an exponential function) and by the access resistance $R_e^{(i)} = \int_0^\infty K_e^{(i)}(s) ds$. The two distributions plotted in the inset show the electrode properties measured in all the recordings included in this paper. **c)** The access resistance is plotted as a function of the electrode timescale. This plot indicates that high access resistances are often associated with longer electrode timescales. **d)** For each injection i , the membrane potential V (red) is estimated by subtracting from the recorded signal V_{rec} (black) the potential drop across the electrode. Since the injected current has a baseline $I_0 > 0$, the membrane potential is, on average, lower than the recorded potential. This difference is given by $I_0 R_e^{(i)}$. Inset: zoom illustrating the fact that AEC acts as a low-pass filter to remove artifacts on the short timescales. The signals shown in the inset have been shifted to have the same mean. Scale bars: 5 mV and 30 ms. **e)** Average shapes of the action potentials obtained from V_{rec} (black) and V (red). The two traces have been shifted to have the same mean.

In practice, it is impossible to accurately estimate the electrode filter K_e at each moment in time. However, since changes in the electrode properties usually occur slowly, it was sufficient to estimate K_e within separable experimental blocks. In our protocol, each experiment consists of many 64-second long injections (i.e., experimental blocks). The assumption that the electrode properties are stable throughout each injection was verified by a systematic estimation of the electrode filter *before*, *during* and *after* each injection (Fig. 2.11).

For each experimental block i , the electrode kernel $K_e^{(i)}$ was estimated following the procedure already used in Badel et al. (2008). Briefly, far from spikes (i.e., in the subthreshold regime) we assume the neuron to act as a linear system described by the membrane filter $K_m^{(i)}$. Consequently, the recorded potential can be modeled as a filtered version of the input current

$$V_{\text{rec}}(t) = V_0 + \int_0^t K^{(i)}(s) I_{\text{ext}}(t-s) ds, \quad (2.11)$$

where V_0 is the resting potential and $K^{(i)} = K_e^{(i)} + K_m^{(i)}$ accounts for both the electrode and the passive membrane. The filter $K^{(i)}$ was extracted from segments of subthreshold data by calculating the Wiener-Hopf optimal filter that provides the best estimate of the derivative of the recorded potential:

$$\dot{V}_{\text{rec}}(t) = \int_0^\infty K^{(i)}(s) \dot{I}_{\text{ext}}(t-s) ds. \quad (2.12)$$

The electrode kernel $K_e^{(i)}$ was then obtained by subtracting from the optimal filter an exponential function fitted on the tail of $K^{(i)}$. As in Brette et al. (2008) and Badel et al. (2008), this tail is interpreted as the membrane filter $K_m^{(i)}$. To improve the accuracy, this procedure was repeated 15 times by resampling experimental data from the available subthreshold segments. The final estimate of $K_e^{(i)}(t)$ was obtained by averaging across bootstrap-repetitions. The maximal length of the two kernels $K^{(i)}(t)$ and $K_e^{(i)}(t)$ was set to 100 ms and 7 ms, respectively.

Figure 2.9a,b shows a typical Winener-Hopf filter $K^{(i)}(s)$ and a typical electrode filter $K_e^{(i)}(s)$, respectively. Each electrode filter was characterized by its timescale (estimated by fitting $K_e^{(i)}(s)$ with an exponential function) and by its access resistance $R_e^{(i)} = \int_0^\infty K_e^{(i)}(s) ds$. The distribution of electrode properties measured in all the recordings included in this study is shown in Figure 2.9b (inset) and Figure 2.9c. As shown by a comparison between the recorded potential $V_{\text{rec}}(t)$ and the membrane potential obtained after Active Electrode Compensation $V(t)$ (Fig. 2.9d,e), AEC acts as a low-pass filter by removing artifacts on short timescales.

Since the electrode properties were estimated *during* each individual experimental block i , we were also able to compensate artifacts that are due to changes in the *access resistance*. Figure 2.10a,b illustrates a typical example in which AEC successfully removed a drift of the recorded potential induced by slow changes in the electrode filter. Note however that, in most cases, the electrode properties were stable throughout the entire experiment (a typical example of a stable recording is shown in Fig. 2.10c,d and summary data are presented in Fig. 2.10e,f).




Figure 2.10: Active Electrode Compensation removes artifacts on long timescales. Active Electrode Compensation (AEC) removes artifacts caused by slow changes in the access resistance. According to our protocol, each experiment is divided in several 64-s-long injections. AEC was performed using electrode filters $K_e^{(i)}$ estimated independently at each repetition (labeled i). **a**) Electrode filters $K_e^{(i)}$ estimated in 12 consecutive stimulations (gray level increases from $K_e^{(1)}$ to $K_e^{(12)}$). In this specific experiment, the properties of the electrode filter clearly change over time. **b**) Top: the access resistance $R_e^{(i)}$ (computed by integrating the electrode filters in panel *a*) increases with time. This produces a drift in the recorded potential that we were able to correct with AEC. Bottom: black and red dots show the average subthreshold potential computed using the recorded signal V_{rec} and the membrane potential V estimated with AEC, respectively. **c**, **d**) Same plots as in *a* and *b* showing the data of a typical experiment in which the electrode properties are stable. **e**) For each experiment included in this study, the access resistance is plotted as a function of the average subthreshold recorded potential. Groups of dots having the same color represent injections into the same neuron. Different colors represent different neurons. Most of the recordings are stable (in these cases the data points form a small cloud). Slow drifts in the recorded potential are always associated with changes in the access resistance. **f**) The access resistance is plotted as a function of the average membrane potential estimated with AEC. These results demonstrate the ability of AEC to compensate drifts due to changes in the access resistance. Arrows (1) and (2) indicate the two experiments shown in panels *a-b* and *c-d*, respectively.

Finally, the stability within individual injections was verified by estimating the electrode filter $K_e^{(i)}$ *before* and *after* each experiment block i . The results shown in Figure 2.11 confirm our assumption that the electrode properties were stable within individual injections.

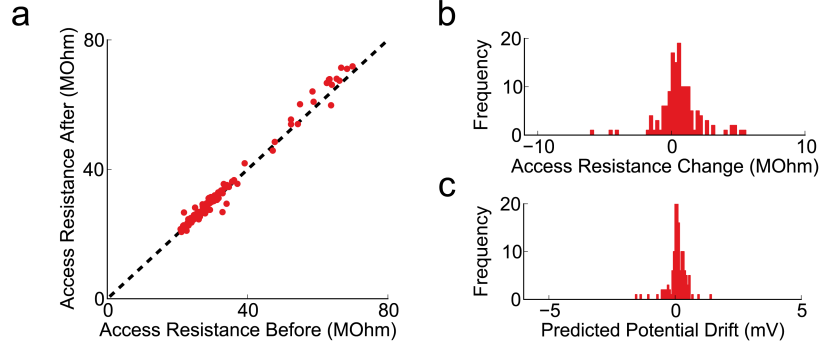


Figure 2.11: Electrode properties are stable during single injections. According to our protocol, each experiment consists of many injections of 64 seconds, labeled i . To remove artifacts due to changes in the electrode properties, a new electrode filter $K_e^{(i)}$ is estimated at each injection i . The assumption was made that the electrode filter does not change during single injections. To validate this hypothesis, the electrode filter was estimated before and after each injection using the response to 4-s-long subthreshold noise. **a)** Each dot shows the access resistance measured before and after each injection. **b)** Distribution of the changes $\Delta R_e^{(i)}$ in the access resistance observed in the experiments. $\Delta R_e^{(i)}$ is defined as the difference between the access resistance measured *after* and *before* the injection i . On average, the change in the access resistance was of $\Delta R_e = 0.62 \text{ M}\Omega$ ($s.d. = 1.54$). **c)** Histogram of membrane potential drifts expected to occur during single injections. The expected drift $\Delta V^{(i)}$ is computed by multiplying $\Delta R_e^{(i)}$ with the baseline current I_0 . On average, the expected drift was of $\Delta V = 0.09 \text{ mV}$ ($s.d. = 0.34$) confirming the hypothesis that electrode properties are sufficiently stable during individual injections.

2.4.6 Fitting the GLIF- ξ model on *in vitro* recordings

To fit GLIF- ξ models, the method introduced in Mensi et al. (2011) was extended to get a more accurate estimate of $\xi(s)$. This was done with an additional hidden variable $I_{\text{drift}}(t)$ able to absorb small drifts that are likely to occur in long recordings.

To get an accurate estimation of the effective adaptation filter $\xi(s)$, *in vitro* recordings were first fitted with a two-process GLIF model that explicitly features both a spike-triggered current $\eta(s)$ and a spike-triggered movement of the firing threshold $\gamma(s)$ (Fig. 2.2). We then obtained the effective adaptation filter $\xi(s)$ by combining $\eta(s)$ and $\gamma(s)$ according to the following formula

$$\xi(t) = \int_0^\infty K_m(t-s)\eta(s)ds + \gamma(t), \quad (2.13)$$

where $K_m(s) = \Theta(s) \frac{R}{\tau_m} e^{-\frac{s}{\tau_m}}$ is the membrane filter, $\Theta(s)$ is the Heaviside step function, $R = g_L^{-1}$ and $\tau_m = RC$. Importantly, the functional shapes of $\eta(s)$ and $\gamma(s)$ were not assumed *a priori*

but were directly extracted from the experimental data by the following two-step procedure.

In the first step, we extracted the functional shape of $\eta(s)$, together with all the parameters that determine the subthreshold dynamics, by fitting $\dot{V}(t)$ on the experimental voltage derivative $\dot{V}^{(\text{data})}(t) = [V^{(\text{data})}(t + \Delta T) - V^{(\text{data})}(t)] / \Delta T$, where $\Delta T = 0.05$ ms was given by the experimental sampling frequency. Since adaptation currents directly affect the membrane potential dynamics, we fitted $\dot{V}^{(\text{data})}$ with the following model

$$C\dot{V} = -g_L(V - E_L) + I_{\text{ext}} - \sum_{\hat{t}_j < t} \eta(t - \hat{t}_j - T_{\text{ref}}) + I_{\text{drift}}(t), \quad (2.14)$$

where Equation 2.7 was extended with a spike-triggered current $\eta(s)$ and the additional term $I_{\text{drift}}(t)$ is an unknown current that averages out at zero over time and captures experimental drifts within individual injections. To avoid any *a priori* assumption on the functional shape of the spike-triggered current, we defined $\eta(s)$ as linear combination of basis functions

$$\eta(s) = \sum_{k=1}^K \alpha_k f_k(s), \quad (2.15)$$

where the coefficients α_k control the shape of $\eta(s)$ and $f_k(s) = \text{rect}\left(\frac{s - T_k}{\Delta_k}\right)$ are rectangular functions of width Δ_k and centered at T_k . For GLIF- ξ_L , we used $K=45$ log-spaced non-overlapping bins with Δ_k ranging from 0.5 ms to 4 s. For GLIF- ξ_S , we set $K=30$ and $\Delta_k \in [0.5, 200]$ ms. Similarly, we defined $I_{\text{drift}}(t)$ as a piecewise constant function

$$I_{\text{drift}}(t) = \sum_{l=1}^L \beta_l \cdot \text{rect}\left(\frac{t - \Delta(l - 0.5)}{\Delta}\right). \quad (2.16)$$

For both GLIF- ξ_L and GLIF- ξ_S , we constrained $I_{\text{drift}}(t)$ to vary slowly in time by choosing a small number of $L = 5$ of regularly spaced bins of size $\Delta = 12.8$ s.

As in Paninski (2004) and Mensi et al. (2011), given the injected current I_{ext} , the estimate of the membrane potential obtained after electrode compensation $V^{(\text{data})}$ as well as the spike times $\{\hat{t}_j\}$, optimal parameters (minimizing the SSE between $\dot{V}^{(\text{data})}$ and \dot{V} of Eq. 2.14) were obtained by solving a multilinear regression problem in discrete time. Since GLIF models do not account for the action potential waveform, all the data points $\{t | t \in [\hat{t}_j - 5\text{ms}; \hat{t}_j + T_{\text{ref}}]\}$ were excluded from the fit. Finally, we fixed the absolute refractory period at $T_{\text{ref}} = 2$ ms and obtained the voltage reset by averaging the potential recorded T_{ref} milliseconds after the spikes $V_r = \langle V^{(\text{data})}(\hat{t}_j + T_{\text{ref}}) \rangle_j$. Performing parameters extraction in presence of the term $I_{\text{drift}}(t)$ qualitatively affected the results and slightly improved the predictive power of the model (Fig. 2.12). Note, however, that the term $I_{\text{drift}}(t)$ was not part of the model but was just used in the fitting procedure to *absorb* slow changes in the subthreshold potential that could not be explained by spike-triggered processes.




Figure 2.12: Influence of extending the fitting procedure with the hidden variable I_{drift} . Extending the fitting procedure with an additional hidden variable I_{drift} improves the accuracy of parameter extraction. To assess the influence of the additional variable I_{drift} , we quantified the predictive power of a GLIF- ξ model fitted under the assumption that E_L does not change over time (i.e., the fit is performed without I_{drift}). In the following we will refer to this model as GLIF- ξ_C (where “C” stands for *control*). **a)** Spike-triggered current $\eta(t)$ of the model GLIF- ξ_C (blue). To allow for a comparison, the adaptation filters of GLIF- ξ_S and GLIF- ξ_L are plotted in orange and red, respectively. The inset of panel *a* shows that performing the fit with the auxiliary variable I_{drift} qualitatively changes the functional shape of the resulting spike-triggered current $\eta(t)$. **b)** The moving threshold $\gamma(t)$ is not affected. Colors as in *a*. **c)** Effective adaptation filter $\xi(t)$ obtained by combining the spike-triggered current and the moving threshold. Colors are as in *a*. **d-e)** Figures 2.5h and j are completed with the predictions of the GLIF- ξ_C model (blue). **f)** Control showing that the average firing rate r_0 does not depend on the period of modulation T (c.f. Eq. 2.1). Colors are as in panel *d* and *e*. **g)** The performance of different models in predicting the occurrence of individual spikes with a precision of ± 4 ms was quantified using the similarity measure M_d^* . Performance of GLIF- ξ_L were not significantly different from GLIF- ξ_C ($n = 12$ cells, paired t -test, $t_{11} = 0.58$, $p = 0.59$) and GLIF- ξ_S ($n = 12$ cells, paired t -test, $t_{11} = -0.82$, $p = 0.43$). **h)** Figure 2.5f is completed with the performance of GLIF- ξ_C ($n = 12$ cells, paired t -test, $t_{11} = -1.10$, $p = 0.29$). **i)** Figure 2.5g is completed with the performance of GLIF- ξ_C ($n = 12$ cells, paired t -test, $t_{11} = -4.17$, $p = 0.002$). In all panels, error bars indicate one standard deviation.

Given the subthreshold dynamics, the second step consisted of estimating the parameters of the firing threshold. Since adaptation due to spike-triggered currents was already captured by

the filter $\eta(s)$, the effective threshold of Equation 2.8 was replaced by

$$V_T^{(\text{bio})}(t) = V_T^* + \sum_{\hat{t}_j < t} \gamma(t - \hat{t}_j - T_{\text{ref}}), \quad (2.17)$$

where $V_T^{(\text{bio})}(t)$ describes the *physiological* threshold at which action potentials were initiated *in vitro*. In contrast to $\xi(s)$, $\gamma(s)$ is not a phenomenological model but describes physiological changes of the firing threshold triggered by the emission of previous spikes. Similar to $\eta(s)$, we defined the moving threshold $\gamma(s)$ as a linear combination of rectangular basis function

$$\gamma(s) = \sum_{k=1}^K \delta_k f_k(s), \quad (2.18)$$

with $f_k(s)$ as in Equation 2.15. Finally, the functional shape of $\gamma(s)$, along with the parameters V_T^* and ΔV , were extracted from experimental data by maximizing the *log-likelihood* of the observed spike-train (Brillinger, 1988):

$$\log L(\theta) = \log p(\{\hat{t}_j\} | V; \theta) = \sum_{\hat{t}_j} \log \lambda_{\theta}(\hat{t}_j) - \int_{\Omega} \lambda_{\theta}(s) ds, \quad (2.19)$$

where $\theta = \{\delta_1, \dots, \delta_K, V_T^*, \Delta V\}$ are the threshold parameters, $\Omega = \{t | t \notin [\hat{t}_j, \hat{t}_j + T_{\text{ref}}]\}$ is a set that excludes periods of absolute refractoriness and the conditional firing intensity $\lambda_{\theta}(s)$ is given by

$$\lambda_{\theta}(t) = \lambda_0 \exp\left(\frac{V(t) - V_T^* - \sum_{\hat{t}_j < t} \gamma(t - \hat{t}_j - T_{\text{ref}})}{\Delta V}\right), \quad (2.20)$$

where $V(t)$ was obtained by integrating Equation 2.14 and, without loss of generality, we set $\lambda_0 = \Delta T^{-1}$. With the exponential function in Equation 2.20, the *log-likelihood* to maximize is a concave function of θ (Paninski, 2004). Consequently, we could perform the fit in discrete time using standard gradient ascent methods (Truccolo et al., 2005; Pillow et al., 2008; Mensi et al., 2011).

With this fitting procedure, an inaccurate estimation of the spike-triggered current $\eta(s)$ would affect the measure of the moving threshold $\gamma(s)$. To make sure that the estimation of $\gamma(s)$ we obtained (Fig. 2.2a, red line) can indeed be attributed to a movement of the physiological threshold, we also extracted the threshold parameters using the experimental membrane potential $V^{(\text{data})}$, rather than V (Fig. 2.2a, gray line).

Power-law fit of the effective adaptation filter $\xi(s)$

For GLIF- ξ_{PL} , the effective adaptation filter $\xi_L(s)$ extracted from the intracellular recordings was fitted with a truncated power law $\xi_{\text{PL}}(s)$ (Eq. 2.2). The fit was performed in two steps. In the first step, we estimated the magnitude α_{ξ} and the scaling exponent β_{ξ} using a least-

square linear regression performed in log-log space. For that, data points were logarithmically resampled and excluded from the fit if $\xi_L(s) < 5 \cdot 10^{-3}$ mV or $s < 5$ ms. In the second step, we obtained the cutoff T_ξ by calculating the intercept between the power-law fitted in the first step and the average value of the extracted kernel $\xi_L(s)$ computed on the first 5 ms. A similar procedure (i.e., least-square linear regression in log-log space with logarithmically resampled points) was used for the power-law fit of the spike-triggered current $\eta(s)$ and the spike-triggered movement of the firing threshold $\gamma(s)$ shown in Figure 2.2a.

2.4.7 Performance evaluation

Cross-validation

All the performances reported in this study were evaluated on datasets that have not been used for parameter extraction. For the predictions reported in Figure 2.4 and Figure 2.5a-g, the model fitted on the first half of the experiment (*training set*) was used to predict the responses observed in the second half (*test set*). Since in certain experiments the average firing rates r_0 observed in the *test set* were slightly different than the ones of the *training set*, the parameter V_T^* was readjusted using the firsts 16 s of all the *test set* injections and models were validated on the responses recorded in the remaining 48 s. According to this procedure, models that do not capture SFA on slow timescales were expected to overestimate the average firing rate r_0 . For the predictions reported in Figure 2.5h-j, a *leave-one-out strategy* was used. In this case, models fitted on the responses to five different periods of modulation were used to predict the sixth one.

Spike-train metrics

To evaluate spike time prediction, we used the similarity measure M_d^* introduced in Naud et al. (2011). M_d^* quantifies the similarity between two groups of spike trains generated by two stochastic processes and corrects the bias caused by the small number of available repetitions. M_d^* takes values between 0 and 1, where $M_d^* = 0$ indicates that the model is unable to predict any of the observed spikes and $M_d^* = 1$ means that the two groups of spike trains have the same instantaneous firing rate and are statistically indistinguishable. M_d^* can also be interpreted as the number of spikes correctly predicted (here with a precision of ± 4 ms) divided by an estimate of the number of reliable spikes.

2.4.8 Estimating the statistical properties of the input current received *in vivo* by neocortical pyramidal neurons

To test the hypothesis that power-law adaptation contributes to efficient coding by *whitening* the single neuron output, we measured the power spectrum of the currents $\Delta I(f)$ received as input at the soma of a neocortical pyramidal neuron *in vivo*. According to Equation 2.14, in absence of spikes, the membrane potential $\Delta V(t)$ is a low-pass filtered version of the input current, where the cutoff frequency $f_c = \tau_m^{-1}$ is defined by the membrane timescale.

Consequently, at all frequencies $f \ll f_c$, we have that $\Delta I(f) = \Delta V(f)/R^2$, with $\Delta V(f)$ being the power spectrum of the subthreshold membrane potential fluctuations and R the input resistance.

We estimated $\Delta V(f)$ using 20-second-long whole-cell recordings ($n = 57$) of the synaptically driven membrane potential dynamics obtained from 7 different L2/3 pyramidal neurons of behaving mice (data from Crochet et al. (2011)). All the *in vivo* recordings were performed in primary somatosensory barrel cortex during active whisker sensation. Further details on the experimental protocol can be found in the original paper (Crochet et al., 2011). Spike-triggered currents last for more than 20 seconds and can in principle affect $\Delta V(f)$ even at very low frequencies. For this reason, only trials with low firing rates $r_0 < 0.5$ Hz were used. However, including recordings with $r_0 > 0.5$ Hz did not affect the results.

2.4.9 Simulating the population response to *in vivo* like inputs

To obtain the results reported in Figure 2.7, we simulated a population of $N = 100$ unconnected GLIF- ξ_{PL} neurons in response to 4000-s-long currents $I(t)$ characterized by a power spectrum $\Delta I(f) \propto f^{-\beta_1}$, with $\beta_1 = 0.67$. Model parameters are given in Table 2.1 and input currents were generated as in Wang et al. (2003) by numerically solving the following inverse Fourier transform

$$I(t) = I_0 + Z \cdot \int_{-\infty}^{+\infty} \sqrt{\Delta I(f)} \mathcal{N}(f) e^{i(2\pi f t + \phi(f))} df, \quad (2.21)$$

where $\mathcal{N}(f)$ is a Gaussian white-noise process, the phases $\phi(f)$ were independently sampled from a uniform distribution and the scaling factor Z was adjusted to fit the power spectrum of the subthreshold membrane potential fluctuations observed *in vivo* (see Fig. 2.7a). To avoid unrealistic large power at slow frequencies, we introduced a cutoff $\Delta I(f) = 0$, for $f < 0.025$ Hz. The highest frequency in the signal was determined by the time step $\Delta T = 0.5$ ms used for numerical simulations. The mean input I_0 was adjusted to obtain a plausible average activity of $A_0 = 4$ Hz, which was consistent with the firing rates obtained *in vitro*.

The population activity $A(t)$ was constructed by counting the number of spikes falling in bins of 50 ms and its power spectrum $A(f)$ was finally computed using time series of 40 s.

2.4.10 Statistics

The number of cells used for the analysis ($n = 12$ or $n = 14$) was limited due to experimental constraints. Data analysis only started after complete data collection and no data were excluded. Two-sided t -test was used as standard. Normality was verified using the Anderson-Darling test. Multiple comparison correction was not appropriate and therefore not used.

2.5 Author contributions

C.P and R.N. conceived the study; C.P. designed the experiments, analyzed the data, performed the modeling and wrote the initial draft; S.M. contributed to data analysis and modeling; W.G. supervised the project. All the authors worked on the manuscript. This project was funded by the Swiss National Science Foundation (grant no. 200020 132871/1; C.P. and S.M.) and by the European Community Seventh Framework Program (BrainScaleS, grant no. 269921; S.M. and R.N.).

3 Automated high-throughput characterization of single neurons by means of simplified spiking models

Christian Pozzorini[†], Skander Mensi[†], Olivier Hagens, Richard Naud, Christof Koch and Wulfram Gerstner; [†]equal contribution (submitted to PLOS Computational Biology)

Abstract

Single-neuron models are useful not only for realistic large-scale simulations of neural systems, but also for rapid characterization of electrophysiological recordings. To account for the large variety of behaviors observed in different neuronal types, single-neuron models should be flexible and allow efficient parameter extraction from experimental data. Here we provide a brief overview of commonly used spiking neuron models and, by taking the Generalized Integrate-and-Fire model as a particular example, we review how simplified models capable of predicting both the spiking activity and the subthreshold dynamics of different cell types can be accurately fitted with a limited amount of data. A procedure is proposed that, combined with a recently developed technology for automatic patch-clamp recordings, allows for automated high-throughput characterization of single neurons.

3.1 Introduction

In vitro patch-clamping is the gold standard used to investigate the intrinsic electrophysiological properties of single neurons, but remains labour intensive and requires a trained experimentalist with high technical skills. In the last years, several platforms have been developed that automatize electrophysiological recordings for ion-channel screening and drug discovery (Dunlop et al., 2008). Most of the existing platforms are, however, designed to record from mammalian cell lines or oocytes in which ion-channels of interest are artificially expressed (Xu et al., 2003; Finkel et al., 2006). In the near future, this technology is likely to be transferred to more complex setups, such as *in vitro* brain slices. High-throughput electrophysiology can be pushed forward with *in vivo* whole-cell patch-clamp recordings that are, at

Chapter 3. Single-neuron characterization by means of spiking models

least partially, automatized (Kodandaramaiah et al., 2012). With this technique, three to seven minutes are sufficient for a robot to automatically identify a cell and form a gigaohm seal of the same quality as achieved by a trained electrophysiologist. This technological advance represents an important step towards high-throughput electrophysiology *in vivo* or on *in vitro* brain slices.

To make sense of the large amount of data that automated patch-clamp will make available, adequate computational tools have to be developed. Due to their high-dimensionality, raw data from intracellular recordings are difficult to interpret and cannot be directly compared against each others. For example, during ongoing activity, it is not clear how different electrophysiological types can be identified from a set of voltage traces acquired from a large number of neurons. To solve this problem, traditional protocols rely on current-clamp injections of stimuli (e.g., square current pulses, ramps of current) that are specifically designed to extract a small number of parameters (e.g. membrane time constant, firing threshold). While this is a valid approach, the input currents adopted in these experiments are artificial and strongly differ from the signals that single neurons process *in vivo*. Moreover, the choice of the parameters used for single-neuron characterization is arbitrary and different parameters are generally estimated from different experiments.

An alternative method to extract and summarize the most relevant information contained in an intracellular recording consists of fitting a spiking neuron model with a limited number of parameters to the data. Ideally, a single-neuron model should be sufficiently complex and flexible to capture, by a single change of parameters, the spiking activity of different neurons, but also simple-enough to allow robust parameter estimation (Herz et al., 2006; Gerstner and Naud, 2009). Detailed biophysical models with stochastic ion channel dynamics can in principle account for every aspect of single-neuron activity; however, due to their complexity, they require high computational power (Koch, 1999; Herz et al., 2006; Markram, 2006; Lang et al., 2011). While systematic fitting of detailed biophysical models is possible (Prinz et al., 2003; Gold et al., 2006; Huys et al., 2006; Druckmann et al., 2007; Hay et al., 2011; Vavoulis et al., 2012), most of the existing methods assume the knowledge of all the parameters that determine the dynamics of the ion channels included in the model. Overall, a reliable and efficient fitting procedure for detailed biophysical models is not known (Gerstner and Naud, 2009). In a second class of spiking neuron models, which we call simplified threshold models, the biophysical mechanisms relevant for neural computation are not explicitly modeled, but are accounted for by phenomenological (i.e., effective) descriptions (Gerstner et al., 2014; Izhikevich, 2007). Despite their simplicity, threshold models are surprisingly good at predicting the single-neuron activity (Softky and Koch, 1993; Troyer and Miller, 1997; Keat et al., 2001; Pillow et al., 2005; Paninski et al., 2005; Jolivet et al., 2008b; Kobayashi et al., 2009; Gerstner and Naud, 2009; Dong et al., 2013), at least for the case of single-electrode somatic stimulation (but see Häusser et al. (2000); Larkum et al. (2009)). Nowadays, simplified spiking neuron models are mainly used in large-scale simulations to study the emergent properties of neural circuits (Izhikevich and Edelman, 2008; Gewaltig and Diesmann, 2007). By taking a different perspective, we will argue that the same models can also serve an equally important purpose, namely to characterize the

electrical properties of single neurons. In this view, simplified threshold models are interpreted as computational tools to compress the information contained in a voltage recording into a set of unique and meaningful parameters. Summarizing the information of complex voltage recordings can in turn enable systematic comparisons and clustering of cell types.

By taking the Generalized Integrate-and-Fire model (GIF; see Mensi et al. (2011); Pozzorini et al. (2013)) as a particular example, we discuss how the electrical properties of single neurons can be efficiently characterized by fitting a simplified threshold model to voltage recordings. After demonstrating that a limited amount of data, and little computing time, are sufficient to achieve this goal, we introduce an experimental protocol that, combined with automated patch-clamp technology, could make automated high-throughput single-neuron characterization possible. The validity of our approach is finally demonstrated with two applications: i) *In silico* recordings obtained by simulating the activity of a multi-compartmental conductance-based model; and ii) *In vitro* recordings from layer 5 (L5) pyramidal neurons obtained using manual patch clamping.

3.1.1 Simplified threshold models

Over the last century, several mathematical models have been developed that describe the electrical behavior of single neurons. So far, efforts of quantitative modeling mainly focused on single-electrode experiments, in which the membrane potential is recorded at a specific location (generally the soma) in response to currents delivered at exactly the same place. The reason for that does not simply reside in the experimental difficulties associated with double patch-clamp recordings (Häusser et al., 2000; Larkum et al., 2009), but also arises from the extreme complexity inherent to dendritic processing.

Single-neuron models can be divided in two main categories: *detailed biophysical models* and *simplified threshold models*. The models belonging to the first category are based on the Hodgkin-Huxley formalism (Hodgkin and Huxley, 1952) and explicitly describe the dynamics of different ion-channels. These models are well suited to study how specific channels interact to shape the global electrical properties of nerve cells. However, due to their complexity, detailed biophysical models are analytically intractable, difficult to use in large-scale simulations and rarely provide an intuitive understanding of the computational principles of single neurons. For these reasons, simplified spiking models are often preferred to investigate the dynamics of neural networks and their emergent properties.

The complexity of a neuron model can be significantly reduced by modeling the intricate dynamics responsible for spike generation with a simple threshold process. Simplified threshold models (*integrate-and-fire* models) describe action potentials as all-or-none unitary events and implicitly assume that all of the information is carried by the precise timing, or the rate, at which action potentials are fired. While the original idea of modeling spike emission by a threshold process dates back to the beginning of the last century (Lapicque, 1907; Hill, 1936), detailed studies involving integrate-and-fire models only started around the 1960s,

when theoreticians began to investigate their mathematical and stochastic properties (Stein, 1965; Knight, 1972). Later, simplified threshold models were shown to provide a surprisingly good description of cortical neurons (Softky and Koch, 1993; Troyer and Miller, 1997) and became extremely popular amongst computational neuroscientists studying the dynamical properties of neural networks. Integrate-and-fire models are generally composed of: i) a set of differential equations describing the dynamics of the subthreshold membrane potential; ii) a set of equations describing the dynamics of the voltage threshold for spike initiation; and iii) a mathematically formalized condition for the emission of action potentials (Box 1). Amongst simplified threshold models, the Leaky Integrate-and-Fire (LIF) model is one of the most widely used in the field of computational neuroscience. Despite its popularity, the LIF model is oversimplified and does not reach a high performance in predicting individual spikes with millisecond precision (Jolivet et al., 2008b; Kobayashi et al., 2009). During the last decades, several extensions have been introduced to account for different single-neuron features such as subthreshold resonance (Izhikevich, 2001; Richardson et al., 2003), spike-frequency adaptation (La Camera et al., 2004; Jolivet et al., 2006b), smooth spike initiation (Fourcaud-Trocmé et al., 2003; Latham et al., 2000; Badel et al., 2008), firing threshold dependence on the membrane potential (Mihalaş and Niebur, 2009; Platkiewicz and Brette, 2010; Higgs and Spain, 2011; Fontaine et al., 2014b) and stochastic spike emission (Gerstner and van Hemmen, 1992; Paninski et al., 2005; Jolivet et al., 2006b). In particular, the INCF *Quantitative Single Neuron Modeling Competition* that took place in 2007 and 2008 (Jolivet et al., 2008b,a) demonstrated that including a spike-triggered current (or a spike-triggered movement of the firing threshold) for spike-frequency adaptation dramatically improves the predictive power of simplified threshold models compared to the LIF model (Gerstner and Naud, 2009).

Ideally, a single-neuron model should: i) be analytically tractable; ii) be sufficiently flexible to capture the behavior of different neuronal types by suitable parameter changes; iii) have a low computational cost; iv) predict spikes with millisecond precision; v) predict the subthreshold dynamics of the membrane potential; vi) come with a fitting procedure that enables efficient parameters extraction from intracellular recordings; and vii) allow for a biophysical interpretation of its parameters and components. Amongst others (e.g. Brette and Gerstner (2005); Izhikevich et al. (2003); Gerstner et al. (2014)), the Generalized Integrated-and-Fire model (GIF) we recently introduced (Mensi et al., 2011; Pozzorini et al., 2013) fulfills all of these requirements. The GIF model is a LIF model augmented with a spike-triggered current $\eta(t)$ and a moving threshold $\gamma(t)$ for spike-frequency adaptation, as well as with the exponential escape-rate mechanism for stochastic spike emission (Box 1). Although being more complex than a LIF model, the GIF model remains amenable to analytical treatments that illustrate how specific single-neuron features affect the global network dynamics (Naud and Gerstner, 2012; Deger et al., 2013). Most importantly, as discussed below, the GIF model is particularly well suited to perform single-neuron characterization because all of its parameters can be efficiently extracted from a limited amount of data.

BOX 1: Generalized Integrate-and-Fire (GIF) model
Subthreshold dynamics of the membrane potential

In the GIF model, the subthreshold membrane potential $V(t)$ evolves according to the differential equation:

$$C\dot{V}(t) = -g_L(V(t) - E_L) - \sum_{\hat{t}_j < t} \eta(t - \hat{t}_j) + I(t), \quad (3.1)$$

where the parameters C , g_L and E_L define the passive properties of the neuron, $I(t)$ is the input current and $\{\hat{t}_j\}$ are the spike times. According to Equation 3.1, each time an action potential is fired, an intrinsic current with stereotypical shape $\eta(t)$ is triggered. Currents triggered by different spikes accumulate and produce spike-frequency adaptation if $\eta(t) > 0$, or facilitation if $\eta(t) < 0$. The functional shape of $\eta(t)$ varies among neuron types (Mensi et al., 2011). Therefore the time course of $\eta(t)$ is not assumed *a priori* but is extracted from intracellular recordings. Each time a spike is emitted, the numerical integration is stopped during a short absolute refractory period T_{ref} and the membrane potential is reset to $V_{\text{reset}} = V(\hat{t}_j + T_{\text{ref}})$.

Dynamics of the firing threshold

The dynamics of the firing threshold $V_T(t)$ is given by:

$$V_T(t) = V_T^* + \sum_{\hat{t}_j < t} \gamma(t - \hat{t}_j), \quad (3.2)$$

where V_T^* is a constant and $\gamma(t)$ describes the stereotypical time course of the firing threshold after the emission of an action potential. Since the contribution of different spikes accumulates, the moving threshold defined in Equation 3.2 constitutes an additional source of adaptation (or facilitation). Similar to $\eta(t)$, the functional shape of $\gamma(t)$ is not assumed *a priori* but is extracted from intracellular recordings.


Stochastic spike emission (escape-rate mechanism)

Spikes are produced stochastically according to a point process with conditional firing intensity $\lambda(t|V, V_T)$, which exponentially depends on the momentary difference between the membrane potential $V(t)$ and the firing threshold $V_T(t)$ (Gerstner and van Hemmen, 1992; Paninski et al., 2005; Jolivet et al., 2006b):

$$\lambda(t) = \lambda(t|V, V_T) = \lambda_0 \cdot \exp\left(\frac{V(t) - V_T(t)}{\Delta V}\right), \quad (3.3)$$

where λ_0 has units of s^{-1} , so that $\lambda(t)$ is in Hz and ΔV defines the level of stochasticity. According to Equation 3.3, if $\Delta V > 0$, the probability of a spike to occur at a time $\hat{t} \in [t; t + \Delta t]$ is given by:

$$P(\hat{t} \in [t; t + \Delta t]) = 1 - \exp\left(-\int_t^{t+\Delta t} \lambda(s) ds\right) \approx \lambda(t)\Delta t. \quad (3.4)$$



Predicting the single-neuron activity with a GIF model. **a)** Block representation of the GIF model. The membrane acts as a low-pass filter $\kappa(t) = C^{-1} \cdot \exp(-g_L t/C)$ on the input current $I(t)$ to produce the modeled potential $V(t)$. An exponential nonlinearity transforms this voltage into an instantaneous firing intensity $\lambda(t)$, according to which spikes are generated. Each time a spike is emitted, both a current $\eta(t)$ and a movement of the firing threshold $\gamma(t)$ are triggered. **b)** The response of a L5 pyramidal neuron to a fluctuating input current (top) has been recorded intracellularly (middle, black). The same protocol was repeated nine times to assess the reliability of the response (black raster). The GIF model accurately predicts both the subthreshold (middle, red) and the spiking response (red raster) of the neuron.

3.1.2 Extracting model parameters from experimental recordings

The input-output transformation of a simplified threshold model strongly depends on its parameters θ (Izhikevich, 2004; Naud et al., 2008). To account for the variety of behaviors observed across neuronal types (Markram et al., 2004; Shinomoto et al., 2009), as well as across neurons of the same cellular type, efficient fitting techniques are required that allow parameter extraction from experimental data on a cell-by-cell basis. A fitting procedure for neuron models is formally defined as an algorithm that, given a set of electrophysiological recordings and an error function $E(\theta, D)$ quantifying the mismatch between model prediction and experimental data D , automatically finds the optimal solution θ_{opt} that minimizes $E(\theta, D)$:

$$\theta_{\text{opt}} = \arg \min_{\theta} \{E(\theta, D)\}. \quad (3.5)$$

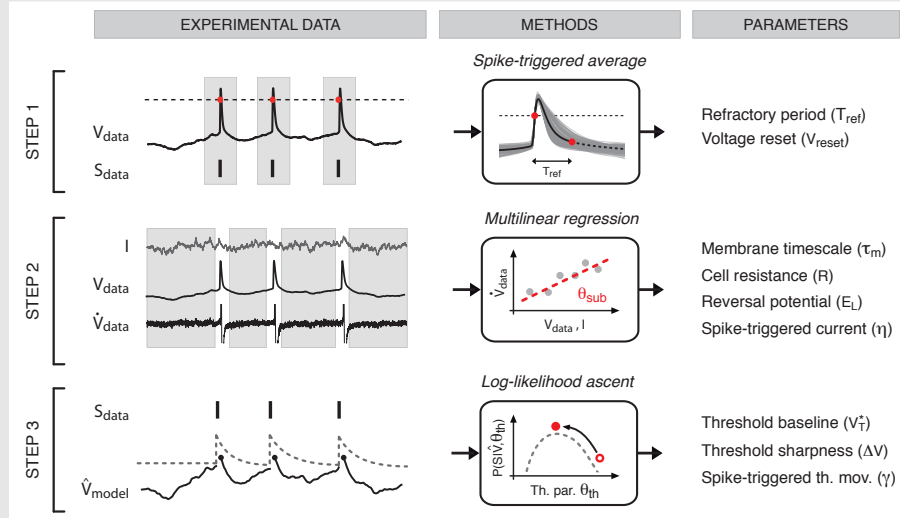
Since the seminal work of Brillinger and colleagues (Brillinger et al., 1976; Brillinger and Segundo, 1979), different fitting methods have been introduced (see Jolivet et al. (2008b,a); Van Geit et al. (2008); Rossant et al. (2011a) and references therein). These methods differ in the choice of both the error function $E(\theta, D)$ and the search algorithm used to solve the optimization problem given by Equation 3.5 (Van Geit et al., 2008). The choice of $E(\theta, D)$ is of particular importance not only because it implicitly determines the optimal solution, but also because it imposes constraints on the search algorithms that can be successfully applied.

A possible solution is to define $E(\theta, D)$ as the mean-square difference between the model and the experimental firing rates observed in response to a set of stationary currents (Rauch et al., 2003; La Camera et al., 2004; Arsiero et al., 2007; Hertäg et al., 2012). However, since neurons communicate by means of action potentials, it seems more natural to define the error function $E(\theta, D)$ as a measure of the dissimilarity between the model and the experimental spike-train. In the case of a deterministic model, because of the threshold process used for spike generation, any error function $E(\theta, D)$ based on a spike-train similarity measure is discontinuous in θ (Brette, 2004). Consequently, the parameters of deterministic models cannot be efficiently extracted using gradient-based methods and *black-box* techniques have to be employed that involve the simulation of the neuron model at each iteration (e.g. evolutionary search algorithms) (Rossant et al., 2010). Overall, since the error function $E(\theta, D)$ is often a non-convex function of θ (i.e., $E(\theta, D)$ contains local minima) and the parameter space is typically large, finding the optimal solution is a difficult problem (Van Geit et al., 2008; Rossant et al., 2011a).

While fitting techniques that entirely rely on the minimization of a spike-train dissimilarity measure are in principle applicable to any neuron model (Rossant et al., 2011a), they are often not efficient and require high computing power. The efficiency of these methods can be significantly improved by pre-estimating part of the model parameters using the information available in the subthreshold membrane potential fluctuations $V_{\text{data}}(t)$ (Jolivet et al., 2004, 2006b; Clopath et al., 2007). Except for some rare cases (e.g., Badel et al. (2008)), fitting

procedures that exploit the knowledge of $V_{\text{data}}(t)$ are performed in two steps.

BOX 2: GIF model parameter extraction



Given the intracellular voltage response $V_{\text{data}}(t)$ evoked *in vitro* by a fluctuating input current $I(t)$, all of the GIF model parameters can be extracted from experimental data using a three-step procedure (Mensi et al., 2011; Pozzorini et al., 2013) (the mathematical details of the procedure are provided in the Materials and Methods section).

Step 1: Extracting the parameters related to absolute refractoriness and voltage reset

First, the experimental spike train $S_{\text{data}} = \{\hat{t}_j\}$ is defined as the collection of instants \hat{t}_j at which $V_{\text{data}}(t)$ crossed a certain threshold from below. The average spike shape $V_{\text{STA}}(t)$ is then obtained by computing the spike-triggered average (STA) of $V_{\text{data}}(t)$. The absolute refractory period T_{ref} is fixed to twice the spike width at half maximum and the reset potential is computed as $V_{\text{reset}} = V_{\text{STA}}(T_{\text{ref}})$.

Step 2: Extracting the parameters related to the subthreshold voltage dynamics

The first-order temporal derivative of the experimental voltage $\dot{V}_{\text{data}}(t)$ is estimated from the data and the parameters determining the membrane potential dynamics are extracted by fitting the model voltage derivative $\dot{V}(t)$ on $\dot{V}_{\text{data}}(t)$. This is done by exploiting the knowledge of the experimental voltage $V_{\text{data}}(t)$ and the external input $I(t)$. To avoid *a priori* assumptions on the functional shape of the spike-triggered current, $\eta(t)$ is expanded in a linear combination of rectangular basis functions. Consequently, optimal subthreshold parameters $\theta_{\text{sub}} = \{C, g_L, E_L, \eta(t)\}$ minimizing the sum of squared errors between $\dot{V}(t)$ and $\dot{V}_{\text{data}}(t)$ can be efficiently obtained by solving a multilinear regression problem (Paninski et al., 2005). Since simplified threshold models do not describe the membrane potential dynamics during action potentials, all the data close to spikes are discarded.

Step 3: Extracting the parameters related to the firing threshold dynamics

The parameters estimated so far are first used to compute the subthreshold membrane potential of the model $\dot{V}_{\text{model}}(t)$. Given $\dot{V}_{\text{model}}(t)$, the parameters $\theta_{\text{th}} = \{V_T^*, \Delta V, \gamma(t)\}$ defining the firing threshold dynamics are then extracted by maximizing the probability (i.e., the *log-likelihood*) of the experimental spike train $S_{\text{data}}(t)$ being produced by the GIF model. Similar to $\eta(t)$, the spike-triggered threshold movement is extracted nonparametrically by expanding $\gamma(t)$ in a linear combination of rectangular basis functions. In the GIF model, the nonlinear function of the escape-rate mechanism (Eq. 3.3) is convex and log-concave. Moreover, its argument is linear in the model parameters θ_{th} . Given these properties, the *log-likelihood* to maximize is guaranteed to be a concave function of the model parameters (Paninski, 2004) and the optimization problem can be solved using standard gradient ascent techniques.

Chapter 3. Single-neuron characterization by means of spiking models

In the first step, all of the parameters θ_{sub} governing to the subthreshold voltage dynamics are estimated by minimizing the error on the membrane potential fluctuations. In the second step, a black-box optimization is performed to find the threshold parameters θ_{th} that maximize a spike-train similarity measure. The main advantage of this strategy resides in the fact that the number of parameters that have to be estimated through black-box optimization is significantly reduced. Interestingly, the winner of the INCF *Quantitative Single Neuron Modeling Competition* used this approach (Kobayashi et al., 2009).

To efficiently exploit the information contained in the subthreshold fluctuations of the membrane potential, a good strategy is to minimize the sum of squared errors between the model and the experimental voltage derivative $\dot{V}_{\text{data}} \approx (V_{\text{data}}(t + \Delta T) - V_{\text{data}}(t))/\Delta T$ (Paninski et al., 2005). This method is particularly convenient because many subthreshold parameters can be efficiently estimated with a simple linear regression (Paninski et al., 2005). Moreover, important processes like linearized subthreshold currents, spike-dependent adaptation and subthreshold nonlinearities can be extracted nonparametrically, thereby limiting the number of *a priori* assumptions (e.g. the number of time scales over which spike-frequency adaptation has to be modeled (Pozzorini et al., 2013)).

Several problems arising during parameters extraction for deterministic models can be avoided with probabilistic spiking models (Paninski et al., 2007; Gerstner et al., 2014). Stochastic models provide a big advantage because an objective function for the optimization problem is naturally provided by the likelihood function of the data $p(D|\theta)$. In this framework, the optimal set of parameters is given by the maximum likelihood estimate θ_{MLE} , which is formally defined as the set of parameters θ under which the probability of observing the experimental data D is maximized:

$$\theta_{\text{MLE}} = \arg \max_{\theta} \{p(D|\theta)\}. \quad (3.6)$$

In contrast to deterministic threshold models, the objective function $p(D|\theta)$ is smooth in θ and the optimization problem defined in Equation 3.6 can be solved using standard gradient-ascent methods that do not involve multiple simulations of the neuron model response. Different maximum likelihood methods have been developed in which the spike generation mechanism is made stochastic using either diffusive noise (Paninski et al., 2004; Pillow et al., 2005; Dong et al., 2011) or the escape-rate mechanism (Paninski, 2004; Truccolo et al., 2005; Pillow et al., 2008). Amongst these models, the Generalized Linear Model (GLM, Truccolo et al. (2005); Pillow et al. (2008)) is particularly convenient because its likelihood function is concave in θ (Paninski, 2004). This property is extremely valuable for parameter extraction because it guarantees that there are no non-global ("local") maxima. Indeed, in the absence of local maxima, standard gradient-ascent techniques are always guaranteed to find the optimal set of parameters.

Despite these advantages, GLMs are typically considered as statistical models for spike trains and their parameters are only loosely related to biophysical properties of the neuron. The

reason for this is that GLM parameter extraction entirely relies on spiking data. If on one hand this fact constitutes a big advantage in case of (multielectrode) extracellular recordings, the standard GLM framework is less appropriate for whole-cell current-clamp data. GLMs do not explicitly model the membrane potential dynamics, do not exploit all the information available in intracellular recordings and, consequently, cannot explain the subthreshold activity of single neurons. To overcome these problems, a fitting method for GIF models (Mensi et al., 2011; Pozzorini et al., 2013) has been recently introduced in which the subthreshold parameters θ_{sub} are first extracted by performing a linear regression on \dot{V}_{data} and the parameters θ_{th} governing the firing threshold dynamics are successively extracted using a maximum likelihood approach (Box 2). With this technique, a biologically plausible model that simultaneously captures the spike response and the subthreshold voltage dynamics of different neuron types can be efficiently fitted to experimental data (Mensi et al., 2011; Pozzorini et al., 2013).

3.1.3 Model validation

In order to build an automated pipeline for single-neuron characterization, a fitting algorithm has to be complemented with a validation protocol designed to automatically detect and discard trials in which the fitting procedure fails. Good spiking neuron models should be able to accurately predict the occurrence of individual action potentials with millisecond precision (Gerstner and Naud, 2009). To take the stochastic nature of single neurons into account (Mainen and Sejnowski, 1995; Jolivet et al., 2006b), the validation protocol should assess the model performance in predicting spike emission probability. To that end, a set of recordings is required in which single neurons are repetitively stimulated with the same time-dependent test current $I_{\text{test}}(t)$. The resulting set of experimental spike trains has then to be compared against a set of spike trains predicted by repetitive simulations of the model.

To assess the model's predictive power, the similarity between the two sets of spike trains has to be quantified. For this purpose, different metrics have been introduced (Kistler et al., 1997; Victor and Purpura, 1997; Rossum, 2001; Paiva et al., 2009). Amongst them, the similarity measure M_{d}^* (Naud et al., 2011) (Box 3) is particularly well suited because it resolves the *small sample bias* known to affect most of the similarity measures when the number of available spike trains is small. Also, in contrast to previous measures based on naive pairwise comparisons (e.g., the Γ -coincidence factor used in Jolivet et al. (2008a)), M_{d}^* does not suffer from the so-called *deterministic bias* known to favor noise-free models (Naud et al., 2011). M_{d}^* takes values between 0 and 1 (where $M_{\text{d}}^* = 0$ indicates that the model is unable to predict any of the experimental spikes and $M_{\text{d}}^* = 1$ indicates a perfect match for a particular level of temporal granularity) and quantifies the similarity between the experimental and predicted spike-emission probability. This measure can be interpreted within an abstract vector space as the normalized distance between the experimental and the predicted spike trains (see Box 3). During a comparison of two peristimulus time histograms (PSTHs), M_{d}^* is sensitive to differences in shape and amplitude of the PSTH whereas correlation-based measures (i.e., *angle measures*) disregard the amplitude.

3.2 Results

3.2.1 A protocol for automated high-throughput single-neuron characterization

By simulating the process of GIF model parameter extraction (Box 2) and validation (Box 3) on artificial datasets of different sizes, we concluded that only a limited amount of experimental data, as well as little computing time, are required to characterize a simplified threshold model (see Section 3.4.2). Based on this result, we designed a high-throughput protocol for the automated fitting and validation of GIF models on *in vitro* intracellular recordings (Fig. 3.1).




Figure 3.1: Schematic representation of the protocol for high-throughput single-neuron characterization. To characterize the properties of the electrode required for data preprocessing by Active Electrode Compensation (AEC, see Box 4), the experimental protocol starts with the injection of a short subthreshold current. While the filtering properties of the electrode are estimated (AEC box, left part), the training dataset is collected. After training set collection, the raw data are preprocessed with AEC to compensate for the bias known to occur when the same electrode is used to both record and stimulate the neuron (AEC box, right part). Then, in parallel with GIF model parameter extraction and successive spike timing prediction, the test dataset is collected by injecting 9 repetitions of the same time-dependent current. Finally, after complete acquisition of the test set, the similarity measure M_d^* between the observed and the predicted spike trains is computed. Overall, GIF model parameter extraction and validation requires around 5 minutes of recording and CPU time.

The protocol is conceptually divided in two parts. In the first part, a *training set* is acquired by recording the single-neuron response to a 100-second fluctuating current $I(t)$. These data are used to perform GIF model parameter extraction. The input current $I(t)$ is generated by an Ornstein-Uhlenbeck process (i.e., a colored noise process) sampled at $\Delta T^{-1} = 20$ kHz by numerically solving the stochastic differential equation $\tau \dot{I} = -I + I_0 + \sqrt{2\tau}\sigma(t)\xi(t)$ in discrete time

$$I(t + \Delta T) = I(t) + \frac{I_0 - I(t)}{\tau} \cdot \Delta T + \sqrt{\frac{2\sigma^2\Delta T}{\tau}} \cdot \mathcal{N}(0, 1), \quad (3.7)$$

where $\xi(t)$ is a Gaussian white-noise process generated by independently drawing from a Normal distribution $\mathcal{N}(0, 1)$, $\tau = 3$ ms is the characteristic timescale on which the input fluctuates, I_0 defines the mean input and $\sigma(t)$ is the time-dependent standard deviation of $I(t)$.


Stationary filtered Gaussian processes have been extensively used to model the input current received *in vivo* at the soma of neocortical neurons (Destexhe and Paré, 1999). In our protocol, the assumption of stationarity is relaxed by modulating the variance of the input with a periodic oscillation (Lundstrom et al., 2008) given by:

$$\sigma(t) = \sigma_0(1 + \Delta\sigma \sin(2\pi f t)), \quad (3.8)$$

where σ_0 and $\Delta\sigma$ are constants and $f = 0.2$ Hz is the modulation frequency. Such an input current with non-stationary statistics drives the neurons through different regimes and generates broad ISI distributions that better constrain the fit of parameters characterizing adaptation processes. At the beginning of the experiment (i.e. before the first part of the protocol starts), the input parameters I_0 , σ_0 and $\Delta\sigma$ are automatically adjusted to obtain an average firing rate of around 10 Hz. In the second part of the protocol (*test set*), 9 repetitive injections of a new 10-second current $I_{\text{test}}(t)$ are performed with an interstimulus interval of 10 seconds, so as to allow the cell to recover. These data are used to quantify the predictive power of the GIF model with the spike-train similarity measure M_d^* (Box 3). The test current $I_{\text{test}}(t)$ is generated according to Equations 3.7-3.8 with the same parameters as in the training set.

Current-clamp experiments in which the same electrode is used both for stimulating and recording from single neurons are biased due to the voltage drop across the electrode (Brette et al., 2008). To remove this bias, intracellular recordings can be preprocessed using a technique called Active Electrode Compensation (AEC, see Box 4 and Materials and Methods). In order to perform AEC, the filtering properties of the electrode have to be estimated. For that, an additional 10-second subthreshold current $I_{\text{sub}}(t)$ is injected before the acquisition of the *training set* (Fig. 3.1). $I_{\text{sub}}(t)$ is generated according to Equation 3.7 with parameters $I_0=0$ pA, $\tau=3$ ms and $\sigma(t) = \sigma$ pA, where σ is a constant calibrated to evoke subthreshold membrane potential fluctuations with a standard deviation of around 2-3 mV.

Since all of the computations required for parameter extraction and model validation can be performed *on the fly*, the whole protocol can be performed in 5 minutes and is suitable for high-throughput *in vitro* characterization of single neurons (Fig. 3.1).

BOX 3: Quantifying the predictive power of a spiking neuron model


Measuring the similarity between two sets of spike-trains. **a)** Two sets of spike trains have to be compared. **b)** Each individual spike train (dashed arrows) can be represented in a vector space. The PSTH vectors v_d (solid arrow, black) and v_m (solid arrow, red) lie in the center of the individual spike trains belonging to the two sets. The two ellipses represent the variability of the model spike trains and the experimental spike trains. To define norms and distances between vectors (i.e., between spike trains), the vector space is equipped with an inner product $\langle S_i, S_j \rangle$ (c.f., Eq. 3.10). Within the vector space, different metrics can be interpreted either as angle measures (orange) or distance measures (blue). In this abstract view, the similarity measure M_d^* is related to the distance between v_m and v_d (Naud et al., 2011). **c)** Schematic representation of the scalar product $\langle S_i, S_j \rangle$ defined as the Kistler coincidence window $K_{\text{Kistler}}(s, s')$. The spike trains $S_i(t)$ and $S_j(t)$ are filtered with a rectangular function of size 2Δ ms and a δ -Dirac function, respectively. The resulting functions are then multiplied together and integrated over time. The Kistler coincidence window counts the spikes $S_j(t)$ falling within $\pm\Delta$ ms of a spike in $S_i(t)$ (green arrows). Non-coincident spikes (red arrows) are not counted.

The similarity measure M_d^*

Given a small set of experimental spike trains $S_i^{(d)} = \sum_f \delta(t - \hat{t}_f)$ recorded in response to N_d repetitive injections of the same input current $I_{\text{test}}(t)$ and a large set of spike trains $S_j^{(m)} = \sum_f \delta(t - \hat{t}_f)$ predicted by $N_m \geq 500$ repeated simulations of a stochastic model, the predictive power of the model can be quantified by the similarity measure M_d^* defined as:

$$M_d^* = \frac{2 \cdot \langle v_d, v_m \rangle}{\frac{2}{N_d(N_d-1)} \sum_{i=1}^{N_d} \sum_{i'=i+1}^{N_d} \langle S_i^{(d)}, S_{i'}^{(d)} \rangle + \langle v_m, v_m \rangle}, \quad (3.9)$$

where $v_d = \frac{1}{N_d} \sum_{i=1}^{N_d} S_i^{(d)}$ is the experimental PSTH, $v_m = \frac{1}{N_m} \sum_{j=1}^{N_m} S_j^{(m)}$ is the model PSTH and $\langle v_m, v_m \rangle$ represents its squared norm (Naud et al., 2011). Due to high-throughput requirements and experimental constraints, only a small number of experimental spike trains are available. For this reason, the squared norm of the experimental PSTH (i.e., $\langle v_d, v_d \rangle$) must be carefully estimated using the unbiased estimator provided by the first term in the denominator of Equation 3.9. The similarity measure M_d^* can be interpreted within a vector space as a normalized measure of the distance between the model and the experimental PSTH. In this abstract view, the operator $\langle \cdot, \cdot \rangle$ plays the role of an inner product given by:

$$\langle S_i, S_j \rangle = \int_0^T \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} K(s, s') S_i(t-s) S_j(t-s') ds ds' dt, \quad (3.10)$$

where $K(s, s')$ is a function defining the degree of coincidence between two spikes which occurred at times s and s' .

Kistler coincidence kernel

While different kernels $K(s, s')$ may be used (see Naud et al. (2011)), a sensible choice is the Kistler coincidence kernel. The Kistler coincidence kernel is mathematically defined as $K_{\text{Kistler}}(s, s') = \delta(s') \cdot \Theta(s+\Delta) \cdot \Theta(-s-\Delta)$, where Δ defines the temporal precision on which spike timing prediction is assessed. With this particular choice, the inner product $\langle S_i, S_j \rangle$ equals the number of spikes in S_i that are within $\pm\Delta$ ms of one of the spikes in S_j and, consequently, M_d^* simplifies to:

$$M_d^* = \frac{2n_{dm}}{n_{dd}^* + n_{mm}}, \quad (3.11)$$

with n_{dm} being the average number of coincident spikes between data (d) and model (m), n_{mm} being the average number of coincident spikes computed across $N_m = 500$ repetitions generated by the model and n_{dd}^* being the bias-corrected average number of coincident spikes between different experimental spike trains (that is, the number of coincident spikes between experimental spike trains $S_i^{(d)}$ and $S_j^{(d)}$ averaged across $(i, j) \in [1, N_d] \times [1, N_d]$ with $i \neq j$, see Eq. 3.10).

3.2.2 Applying the protocol on *in silico* recordings

The most sophisticated single-neuron models available today are known as multi-compartment conductance-based models. In contrast to simplified threshold models, these models account for the intricate morphology of both dendritic and axonal arborizations and explicitly describe the dynamics of a large variety of ion channels mediating active currents. Both aspects are likely to play a role in single-neuron information processing (Koch and Segev, 2000; London and Häusser, 2005). A multi-compartment conductance-based model (Hay model, HM) has recently been proposed that captures several features of L5b thick-tufted pyramidal neurons (Hay et al., 2011). In particular, this model includes active dendrites and describes the interactions between Na^+ spiking at the soma, back-propagating action potentials and Ca^{2+} -spikes generated at the distal apical dendrites.

To validate our procedure for high-throughput single-neuron characterization, the protocol described in Figure 3.1 was tested *in silico* by simulating the HM response to a set of current injections (see Materials and Methods). Since we are mainly interested in automatic somatic patching, all *in silico* experiments were performed by delivering the current at the somatic compartment (Fig. 3.2a). HM recordings were then used to perform GIF model parameter extraction (Fig. 3.2b-d). The passive properties of the membrane were characterized by a relatively short timescale ($\tau_m = 6.7$ ms, s.d. 0.1 ms, Fig. 3.2b). GIF model parameter extraction also revealed the presence of a long-lasting adaptation current (Fig. 3.2c) and a spike-triggered movement of the firing threshold (Fig. 3.2d). Consistent with the tendency of L5b pyramidal neurons to produce bursts of action potentials (see Hay et al. (2011) and Fig. 3.2g), the activation of the spike-triggered current was not instantaneous.




Figure 3.2: Testing the protocol for high-throughput single-neuron characterization on *in silico* recordings. **a)** Reconstructed morphology of the deterministic Hay model (Hay et al., 2011). The recording site is indicated by the red pipette. To reproduce trial-to-trial spike-timing variability, the HM was made stochastic by corrupting all of the input currents $I(t)$ with a source of additive noise modeled as a zero-mean Gaussian white-noise. **b-d)** GIF model parameters extracted from *in silico* recordings obtained by simulating the HM response to a somatic current injection. The filters obtained by averaging the parameters extracted from 5 independent training sets of $T_{tr} = 100$ s each are shown in red. Gray areas indicate one standard deviation. **b)** Membrane filter $\kappa(t)$. Inset: comparison between the membrane timescale τ_m extracted using the GIF model and the GLM (cf, exponential fit of $\kappa_{GLM}(t)$ in panel **e**). Each couple of open circles indicates the timescale extracted from a specific training set. Bar plots represent the mean and one standard deviation across training sets ($\tau_m = 6.7$ ms, s.d. 0.1 ms, GIF; $\tau_m = 8.9$ ms, s.d. 1.3 ms, GLM). **c)** Spike-triggered current $\eta(t)$. Inset: zoom on the first 400 ms. **d)** Spike-triggered movement of the firing threshold $\gamma(t)$. **e-f)** GLM parameters extracted from the same *in silico* recordings used to fit the GIF model. Average filters are shown in blue. Gray areas indicate

one standard deviation across training sets. **e)** Linear filter $\kappa_{\text{GLM}}(t)$ (blue) and exponential fit (dashed black). For comparison, a rescaled version of the membrane filter $\kappa(t)$ is shown in red. Inset: same data displayed on semi-log scale. **f)** Spike-history filter $h_{\text{GLM}}(t)$. For comparison, a rescaled version of the GIF model effective filter $h(t)$ is shown in red. Inset: same data displayed on double-logarithmic scales. **g)** Interspike interval (ISI) distributions computed using the *test set* data (black) the GIF model prediction (red) and the GLM prediction (blue). **h)** Fraction of the input current $I_{\text{test}}(t)$ (top, gray) used for model validation; typical HM response evoked by a single current injection (middle, black); HM spiking activity in response to nine repetitive injections of the same input (bottom, black raster); PSTH constructed by averaging the nine spike trains with a rectangular window of 500 ms (bottom, black line). GIF model and GLM predictions are shown in red and blue, respectively. Dashed black lines represent 0 nA (top) and 0 Hz (bottom). **i-k)** Performance comparison between GIF model and GLM in predicting the HM activity. Parameter extraction and model validation were repeated five times using different datasets. Each couple of open circles indicates the performance obtained by both models on a specific dataset. Bar plots indicate the mean and one standard deviation across repetitions. **i)** Spike-timing prediction as quantified by M_{d}^* with precision $\Delta = 4$ ms (Kistler coincidence window). **j)** Fraction of variance explained ϵ_{V} on subthreshold membrane potential fluctuations. The GLM does not explicitly model the subthreshold membrane potential dynamics and is therefore not applicable (N/A). **k)** GIF model spike-timing prediction as a function of the training set size used for parameter extraction. Increasing the duration of the training set from 100 s to 120 s does not improve the GIF model predictive power ($M_{\text{d}}^* = 0.80$, s.d. 0.01, $T_{\text{tr}} = 100$ s; $M_{\text{d}}^* = 0.80$, s.d. 0.01, $T_{\text{tr}} = 120$ s; $n=10$, paired Student t -test, $t_4 = 0.05$, $p = 0.97$; n.s. > 0.05).

Following our high-throughput protocol, we then assessed the predictive power of the GIF model by simulating the HM response to nine repetitive injections of a new 10-second current (Fig. 3.2h). In agreement with previous results obtained from *in vitro* recordings (Mensi et al., 2011; Pozzorini et al., 2013), the GIF model was able to predict around 80% of the spikes emitted by the HM ($M_{\text{d}}^* = 0.80$, s.d. 0.01, Kistler coincidence window with $\Delta = 4$ ms; Fig. 3.2i) and captured the HM subthreshold voltage fluctuations with an average root mean squared error (RMSE) of 3.4 mV, s.d. 0.03 mV (variance explained $\epsilon_{\text{V}} = 74.3\%$, s.d. 1.1%; Fig. 3.2j). Repeating the entire protocol by varying the duration of the training set T_{tr} confirmed that a 100-second current is sufficient to ensure convergence of the fitting procedure (Fig. 3.2k). Overall, these results demonstrate that, despite their simplicity, modern threshold models are capable of predicting most of the spikes emitted by a detailed biophysical model in response to complex somatic current injections.

3.2.3 Applying the protocol on *in vitro* patch-clamp recordings

To confirm the results reported in the previous section, the protocol for high-throughput single-neuron characterization was further tested using whole-cell current-clamp *in vitro* recordings from L5 pyramidal neurons in mouse somatosensory cortex (see Materials and Methods). At the beginning of each experiment, the input current was calibrated to obtain an average firing rate of 10 Hz with amplitude fluctuations between 7 and 13 Hz.

BOX 4: Data preprocessing - Active Electrode Compensation

When the same patch-clamp electrode is used to simultaneously stimulate and record from a single neuron, the acquired signal $V_{\text{rec}}(t)$ is a biased version of the real membrane potential $V_{\text{data}}(t)$ (Brette et al., 2008; Badel et al., 2008):

$$V_{\text{rec}}(t) = V_{\text{data}}(t) + V_e(t). \quad (3.12)$$

The bias is due to the voltage drop $V_e(t)$ across the patch-clamp electrode and can be removed using a technique called Active Electrode Compensation (AEC, see Materials and Methods and Brette et al. (2008); Badel et al. (2008)).

Active Electrode Compensation (AEC)

In AEC, the electrode is modeled as an arbitrarily complex linear filter $\kappa_e(t)$, which has to be estimated from a subthreshold current-clamp injection performed at the beginning of the experiment. For all subsequent injections, the voltage drop across the electrode $V_e(t)$ is obtained by convolving the input current $I(t)$ with the electrode filter $\kappa_e(t)$. The membrane potential is finally computed by subtracting $V_e(t)$ from the recorded signal:


$$V_{\text{data}}(t) = V_{\text{rec}}(t) - \int_0^\infty \kappa_e(s) I(t-s) ds. \quad (3.13)$$

Estimating the electrode filter $\kappa_e(t)$

To estimate $\kappa_e(t)$, a short experiment is performed in which the single neuron is stimulated with a subthreshold fluctuating current. The optimal linear filter $\kappa_{\text{opt}}(t)$ between the injected current and the recorded potential is first computed. In AEC, the electrode is assumed to operate on sub-millisecond timescales and the slow decay in $\kappa_{\text{opt}}(t)$ is attributed to the cell. The electrode filter $\kappa_e(t)$ is therefore estimated as (Badel et al., 2008):

$$\kappa_e(t) = \kappa_{\text{opt}}(t) - f(t), \quad (3.14)$$

where $f(t)$ is an exponential function fitted on the tail of $\kappa_{\text{opt}}(t)$. Alternatively, the electrode filter $\kappa_e(t)$ can be extracted from $\kappa_{\text{opt}}(t)$ by considering that also the net current flowing through the cell membrane is affected by the electrode properties (see Brette et al. (2008)).



Active Electrode Compensation. a-c) A pyramidal neuron was patched with two electrodes: the *injecting electrode* (orange) used to both inject the input current and record the membrane potential and the *control electrode* (black) only used to record the membrane potential. **b)** The signals recorded with the injecting electrode (orange) significantly differ from the control (black). **c)** After AEC, the signal acquired with the injecting electrode (red) matches the control (black). Insets in panels *b-c* show a zoom on the voltage trajectory preceding an action potential. Scale bars: 10 mV, 3 ms. **d)** Schematic representation of AEC. The artifactual voltage $V_e(t)$ is estimated by filtering the input current $I(t)$ with the electrode filter $\kappa_e(t)$. The intracellular membrane potential $V_{\text{data}}(t)$ is obtained by subtracting the artifactual voltage $V_e(t)$ from the recorded signal $V_{\text{rec}}(t)$. **e)** Typical optimal linear filter $\kappa_{\text{opt}}(t)$. Dashed black: exponential fit performed on the tail of $\kappa_{\text{opt}}(t)$. Inset: magnification of the y-axis. **f)** Typical electrode filter $\kappa_e(t)$ obtained from the optimal linear filter $\kappa_{\text{opt}}(t)$ shown in panel *e*. When recordings are performed using the bridge-balance technique, $\kappa_e(t)$ is characterized by an initial negative peak. Dashed black: exponential fit performed on $\kappa_e(t)$ to estimate the characteristic timescale τ_e of the electrode. Inset: distribution of the electrode timescale τ_e estimated in different recordings (cf Fig. 3.3). Panels *b-c* were adapted from Badel et al. (2008).

All of the *in vitro* recordings were preprocessed using AEC (Brette et al., 2008; Badel et al., 2008) (Box 4). Consistent with previous results (Badel et al., 2008; Pozzorini et al., 2013), the electrode response was characterized by a very rapid timescale $\tau_e = 0.54$ ms, s.d. 0.11 ms (see Box 4, panel *f*).

After AEC, the *in vitro* recordings acquired from 10 different L5 pyramidal neurons (Fig. 3.3a) were used to perform GIF model parameter extraction (Fig. 3.3b-e). All of the extracted parameters were consistent with the ones obtained in an independent study by fitting the GIF model to *in vitro* recordings from L5 pyramidal neurons responding to a mean-modulated input (Pozzorini et al., 2013) (Fig. 3.3e). In particular, we found that spike-frequency adaptation was mediated by both a long-lasting spike-triggered current (Fig. 3.3c) and a movement of the firing threshold (Fig. 3.3d). Testing the predictive power of the GIF model on a new set of recordings in which a test current $I_{\text{test}}(t)$ was repetitively injected (Fig. 3.3i) confirmed the model's ability in predicting around 80% of the spikes ($M_d^* = 0.79$, s.d. 0.04, Kistler coincidence window with $\Delta = 4$ ms; Fig. 3.3j). Also, the GIF model was able to predict the subthreshold membrane potential fluctuations with a RMSE of 3.6 mV, s.d. 0.5 mV (variance explained $\epsilon_V = 80.1$ %, s.d. 4.1 %; Fig. 3.3k). Intriguingly, the GIF model achieved almost identical performances in predicting *in silico* and *in vitro* recordings, indicating that detailed biophysical models could be used in the future to guide the improvement of simplified threshold models.

Comparing the predictive power of different GIF models with parameters extracted from five training sets of different durations further confirmed that 100 seconds of intracellular recordings are sufficient to accurately constrain the GIF model parameters (Fig. 3.3i). Overall, we concluded that our high-throughput protocol based on GIF model parameter extraction and validation is suitable for high-throughput single-neuron characterization.

3.3 Discussion

3.3.1 GIF model limitations and comparison with other spiking models

In contrast to popular nonlinear integrate-and-fire models like the adaptive exponential integrate-and-fire (ADEx, Brette and Gerstner (2005)) or the adaptive quadratic integrate-and-fire (AQIF, Izhikevich et al. (2003)), the GIF model describes the subthreshold dynamics of the membrane potential with a linear differential equation (see Eq. 3.1). Consequently the GIF model does not account for smooth spike-initiation. Theoretical investigations initially concluded that spike initiation was a smooth process (Jack et al., 1975) and emphasized the relevance of nonlinear integrate-and-fire models. However, more recent studies found that, in cortical neurons, the spike onset is almost as in LIF or GIF models (Naundorf et al., 2006; Badel et al., 2008; Rossant et al., 2011a). Further evidence for sharp spike-initiation was provided by other studies showing that cortical neurons can encode fluctuating inputs up to 300 Hz (Tchumatchenko et al., 2011; Ilin et al., 2013). Overall, these results suggest that neglecting smooth spike-initiation does not constitute a major limitation of the model.




Figure 3.3: Testing the protocol for high-throughput single-neuron characterization on *in vitro* patch-clamp recordings. **a)** Staining of a biocytin-filled L5 pyramidal neuron included in this study. **b-e)** GIF model parameters extracted from ten L5 pyramidal neurons. Average filters are shown in red. Gray areas indicate one standard deviation across neurons. **b)** Membrane filter $\kappa(t)$. Inset: comparison between the characteristic timescale $\tau_m=20.9 \pm 6.5$ ms of $\kappa(t)$ and the slow timescale $\tau_{\text{slow}}=22.5 \pm 3.0$ ms of $\kappa_{\text{GLM}}(t)$ (see panel **f**). Each couple of open circles indicates the parameters measured in a single neuron. Bar plots indicate the mean and one standard deviation across neurons. **c)** Spike-triggered current $\eta(t)$ displayed on double-logarithmic scales. **d)** Spike-triggered movement of the firing threshold $\gamma(t)$ displayed on double-logarithmic scales. **e)** Distributions of GIF model parameters extracted from ten L5 pyramidal neurons. From left to right: reversal potential, E_L ; membrane timescale, $\tau_m = C/g_L$; cell resistance, $R = g_L^{-1}$; firing threshold baseline, V_T^* ; firing threshold sharpness, ΔV . For comparison, gray arrows indicate the average parameters obtained in a previous study (Pozzorini et al., 2013). **f-g)** GLM parameters extracted from ten L5 pyramidal neurons. Average filters are shown in blue. Gray areas indicate one standard deviation across neurons. **f)** GLM linear filter $\kappa_{\text{GLM}}(t)$ (blue). For comparison, a rescaled version of the GIF filter $\kappa(t)$ is shown in red. Inset: same data shown on a semi-logarithmic scales. To quantify the slow timescale τ_{slow} (panel **b**, inset) and the fast timescale

$\tau_{\text{fast}} = 1.9 \pm 0.5$ ms of $\kappa_{\text{GLM}}(t)$, we performed a double exponential fit (not shown for clarity) of $\kappa_{\text{GLM}}(t)$. **g**) GLM spike-history filter $h_{\text{GLM}}(t)$ (blue). For comparison, a rescaled version of the effective GIF adaptation filter $h(t)$ (Eq. 3.16) is shown in red. **h**) Distribution of the GLM parameter λ_0 extracted from ten L5 pyramidal neurons. **i**) Input current $I_{\text{test}}(t)$ (top, gray) used for model validation, typical L5 pyramidal neuron response evoked by a single current injection (middle, black); spiking activity observed in response to nine repetitive injections of the same input (bottom, black raster) and PSTH constructed by averaging the nine spike trains within rectangular windows of 500 ms (bottom, black line). GIF model and GLM predictions are shown in red and blue, respectively. **j-l**) Summary data for the performance of the GIF model and the GLM in predicting the responses of 10 L5 pyramidal neurons. Each couple of open circles indicates the performance on an individual cell. Bar plots indicate the mean and one standard deviation across neurons. **j**) Spike-timing prediction as quantified by M_{d}^* with precision $\Delta = 4$ ms (Kistler coincidence window). **k**) Fraction of variance explained ϵ_{v} on subthreshold membrane potential fluctuations. The GLM does not explicitly model the subthreshold membrane potential dynamics and is therefore not applicable (N/A). **l**) GIF model spike-timing prediction as a function of the training set size used for parameter extraction. Increasing the duration of the training set from 100 s to 120 s does not improve the GIF model predictive power ($M_{\text{d}}^* = 0.79$, s.d. 0.04, $T_{\text{tr}} = 100$ s; $M_{\text{d}}^* = 0.79$, s.d. 0.04, $T_{\text{tr}} = 120$ s; $n=10$, paired Student t -test, $t_9 = 0.25$, $p = 0.8$; n.s. > 0.05).

The GIF model describes the passive properties of the membrane with an exponential filter $\kappa(t)$ characterized by a single timescale τ_{m} . According to cable theory (Rall, 1977), the large number of dendritic branches explicitly modeled in the HM (Fig. 3.2a) and present in L5 pyramidal neurons (Fig. 3.3a) is expected to manifest itself in a membrane filter decaying over multiple timescales. To verify the accuracy of the GIF model single-exponential assumption, we used both *in silico* and *in vitro* recordings to fit a Generalized Linear Model (GLM, Truccolo et al. (2005); Pillow et al. (2008)). In the GLM, spikes are emitted stochastically according to a conditional firing intensity $\lambda_{\text{GLM}}(t)$ given by

$$\lambda_{\text{GLM}}(t) = \lambda_0 \cdot \exp \left(\int_0^t \kappa_{\text{GLM}}(s) I(t-s) ds + \sum_{\hat{t}_j < t} h_{\text{GLM}}(t - \hat{t}_j) \right), \quad (3.15)$$

where λ_0 is a constant, $\kappa_{\text{GLM}}(s)$ is a linear filter acting on the input current $I(t)$ and $h_{\text{GLM}}(t)$ is a spike-triggered process describing how the spiking probability is affected by previous action potentials emitted at times $\{\hat{t}_j\}$. Importantly, the functional shapes of $\kappa_{\text{GLM}}(t)$ and $h_{\text{GLM}}(t)$ are not assumed *a priori*, but are extracted from experimental data using a nonparametric maximum-likelihood method (see Truccolo et al. (2005); Pillow et al. (2008) and Materials and Methods). In the case of *in silico* recordings, we found that the GLM filter $\kappa_{\text{GLM}}(t)$ and the membrane filter $\kappa(t)$ of the GIF model were in good agreement (Fig. 3.2e), suggesting that the complex dendritic morphology of the HM weakly affects temporal integration at the somatic compartment, when the current is injected at the soma. Further quantitative evidence was provided by fitting $\kappa_{\text{GLM}}(t)$ with a single exponential function and comparing the resulting timescale against τ_{m} (Fig. 3.2b, inset). Overall, the linear filters $\kappa_{\text{GLM}}(t)$ and $\kappa(t)$ extracted from *in vitro* recordings were also in good agreement (Fig. 3.3b,f). However, the large values observed in the first two bins of $\kappa_{\text{GLM}}(t)$ indicated the presence of a rapid component that the

GIF model was unable to capture (Fig. 3.3f, inset).

In both *in silico* (Fig. 3.2f) and *in vitro* (Fig. 3.3g) recordings, we found that the GLM spike-history filter $h_{\text{GLM}}(t)$ was in good agreement with the effective adaptation filter $h(t)$ of the GIF model computed by merging the spike-triggered current $\eta(t)$ and the threshold movement $\gamma(t)$ (Mensi et al., 2011; Pozzorini et al., 2013):

$$h(t) = \int_0^\infty \kappa(s)\eta(t-s)ds + \gamma(t). \quad (3.16)$$

This result indicates that $h_{\text{GLM}}(t)$ combines, but cannot disentangle, the effects of the adaptation current $\eta(t)$ and the movement of the firing threshold $\gamma(t)$. At first glance, having two spike-dependent processes might seem redundant. However, while the firing threshold only affects spike probability, adaptation currents also alter the dynamics of the subthreshold membrane potential. This explains why the correct distinction between these two forms of adaptation is key to correctly predict the subthreshold response of single neurons.

To compare the GIF model performance against a reference model, we finally used the GLM to predict the test dataset (Fig. 3.2h and 3i). In terms of spike-timing prediction, we found that the GLM performance ($M_d^* = 0.79$, s.d. 0.01, *in silico* recordings; $M_d^* = 0.81$, s.d. 0.04, *in vitro* recordings) was similar to the one achieved by the GIF model (Fig. 3.2i and 3j). This result indicates that, in the case of L5 pyramidal neurons, the small difference observed between the linear filters $\kappa(t)$ and $\kappa_{\text{GLM}}(t)$ has little consequences. However, this might be different for neurons in other brain regions or in the case of more sophisticated stimulation paradigms.

3.3.2 Limitations of point-neuron models

In the GIF model, the morphological complexity of single-neurons is reduced to a single homogeneous compartment (i.e. to a *point*) and dendritic computations are implicitly assumed to be linear. This *point-neuron* assumption probably constitutes the major limitation of the model. Conceptually, the GIF is a model of the soma and its proximal axonal segments. For small, electronically compact neurons, GIF models probably summarize the essential dynamics of the neuron as a whole. Regenerative processes in dendrites present a potential limitation of this approach (Häusser et al., 2000; Larkum et al., 2009). Dendritic dynamics can in principle affect the properties of somatic current injections (Doiron et al., 2001, 2007) and therefore contribute to the effective characterization. Importantly, dendritic spikes will influence input to the dendritic and somatic compartments differentially. An extension of the GIF framework was recently shown to predict the spike timings of independent current injection in both the soma and the apical dendrite, in a L5 pyramidal neuron where Ca^{2+} -spikes are known to occur (Naud et al., 2014). To cast NMDA-based dendritic spikes into an extension of the GIF will require further work, but the successful modeling of Ca^{2+} -spikes and the success of statistical methods for embedded nonlinearities (McFarland et al., 2013) constitute encouraging results. Setting aspects of morphology and threshold aside, biophysically detailed models may still have dynamical features inaccessible to the GIF model. Systematic reduction of detailed

models have indeed indicated that higher-order kernels are required to capture their complete dynamics (Marmarelis et al., 2013; Eikenberry and Marmarelis, 2013).

3.3.3 Conclusion

The intrinsic dynamics of individual neurons strongly differ between cell types and brain areas (Markram et al., 2004). This heterogeneity is increasingly considered as a critical feature of the brain and not as a consequence of biological imprecision (Padmanabhan and Urban, 2010; Tripathy et al., 2013). Taking into account single-neuron variability may be crucial to understand how neural systems support computation. In the near future, automated electrophysiology will likely make available increasingly large amounts of data. Due to their inherent complexity (and their high dimensionality), raw data from patch-clamp recordings are difficult to interpret and cannot be directly clustered to identify electrophysiological types. Simplified spiking models are currently employed by computational neuroscientists mainly to investigate the emergent properties of neural networks. Here, we argued that these models also comprise a powerful tool to cast the information provided by voltage recordings into small sets of model parameters that can be easily interpreted and compared (Mensi et al., 2011).

After reviewing the computational methods for the fit and the validation of simplified spiking models to voltage recordings (Box 1-4), we designed a pipeline in which these tools are combined to perform automated single-neuron characterization (Fig. 3.1). On the experimental side, the proposed protocol relies on *in vitro* somatic injections of rapidly fluctuating currents that mimic the *natural inputs* received *in vivo* at the soma of cortical neurons. On the computational side, the methods are based on GIF model parameter extraction. In GIF models, the morphological complexity of single-neurons is reduced to a single homogeneous compartment that describes the voltage dynamics at the injection site (generally the soma). Moreover, the dynamics of different ion channels is not explicitly modeled, but is accounted for by phenomenological descriptions. The main advantage of this simplification is that model parameters can be efficiently extracted from experimental data using a convex optimization procedure, which is guaranteed to find the optimal (and unique) solution.

We demonstrated the validity of our approach using both *in silico* data (Fig. 3.2) and *in vitro* recordings obtained from mouse L5 pyramidal neurons (Fig. 3.3). In both cases, GIF models with around 50 parameters accurately accounted for 100-second recordings performed at 20 kHz (i.e., $2 \cdot 10^6$ data points). In particular, the GIF model was able to explain around 80% of the variance of the membrane potential fluctuations and correctly predicted around 80% of the spikes with a precision of ± 4 ms. While the precise values of these empirical numbers might vary across cell types, as well as across different biological preparations (e.g. *in vivo* recordings or dendritic stimulations), these results demonstrate that simplified spiking models can be successfully used to cast most of the information of a patch-clamp recording into a simple format. Combined with morphological and transcriptional data (Markram et al., 2004), these parameters could finally allow for enhanced characterization of neuronal types. In conclusion,

simplified spiking models can be extremely useful to make sense of large amounts of otherwise difficult to quantify data.

3.4 Materials and Methods

3.4.1 GIF model parameter extraction

Given the intracellular membrane potential $V_{\text{data}}(t)$ measured at a sampling frequency ΔT^{-1} in response to a known input current $I(t)$, as well as the spike times $\{\hat{t}_j\}$ defined as instants at which $V_{\text{data}}(t)$ crosses 0 mV from below, all of the GIF model parameters are extracted following a three-step procedure (Mensi et al., 2011; Pozzorini et al., 2013) (see Box 2).

Step 1: The absolute refractory period T_{ref} is fixed to an arbitrary value and the voltage reset is estimated by the average membrane potential recorded T_{ref} milliseconds after a spike $V_{\text{reset}} = \langle V_{\text{data}}(\hat{t}_j + T_{\text{ref}}) \rangle_j$. In the GIF model, a period of absolute refractoriness can alternatively be implemented by setting the first milliseconds of the spike-triggered threshold movement $\gamma(t)$ to very large values. For this reason, as long as T_{ref} remains smaller than the shortest interspike interval (ISI) observed in the data, its precise value is not critical. A sensible choice is to set T_{ref} about twice the spike width at half maximum. For pyramidal neurons, we fixed $T_{\text{ref}} = 4$ ms.

Step 2: The parameters determining the subthreshold dynamics of the membrane potential are extracted. To allow convex optimization, the spike-triggered current $\eta(t)$ is expanded as a linear combination of basis functions (Paninski et al., 2005):

$$\eta(t) = \sum_{k=1}^K \eta_k f^{(k)}(t), \quad (3.17)$$

where $\{\eta_k\}$ is a set of parameters controlling the time course of $\eta(t)$. The subthreshold parameters $\theta_{\text{sub}}^T = C^{-1} \cdot [g_L, E_L g_L, \eta_1, \dots, \eta_K, 1]$ are then extracted by minimizing the sum of squared errors between the observed voltage derivative \dot{V}_{data} and that of the model (i.e., Eq. 3.1). Since all of the subthreshold parameters θ_{sub} act linearly on the observables, this optimization problem can be efficiently solved by computing the following multilinear regression (Paninski et al., 2005; Huys et al., 2006):

$$\hat{\theta}_{\text{sub}} = (X^T X)^{-1} X^T \dot{V}_{\text{data}}, \quad (3.18)$$

where X is a matrix whose rows are given by the vectors

$$x_t^T = \left[-V_{\text{data}}(t), 1, -\sum_j f^{(1)}(t - T_{\text{ref}} - \hat{t}_j), \dots, -\sum_j f^{(K)}(t - T_{\text{ref}} - \hat{t}_j), I(t) \right], \quad (3.19)$$

and \dot{V}_{data} is a column-vector containing the voltage first-order derivative estimated by finite differences $\dot{V}_{\text{data}}(t) = (V_{\text{data}}(t + \Delta T) - V_{\text{data}}(t)) / \Delta T$. Since the GIF model does not capture the

subthreshold dynamics during spike initiation, all the data points close to action potentials $\{t | t \in [\hat{t}_j - 5 \text{ ms}; \hat{t}_j + T_{\text{ref}}]\}$ are excluded from the regression.

Step 3: The parameters defining the dynamics of the firing threshold are extracted. To determine the functional shape of the spike-triggered movement of the firing threshold, we first expand $\gamma(t)$ as a sum of basis functions:

$$\gamma(t) = \sum_{p=1}^P \gamma_p f^{(p)}(t). \quad (3.20)$$

Given the parameters obtained in the first two steps and the spike times observed in the experiment, the subthreshold membrane potential $\hat{V}_{\text{model}}(t)$ is computed by numerical integration of Equation 3.1. Without loss of flexibility, the parameter λ_0 is fixed to 1 Hz and all threshold parameters $\theta_{\text{th}}^T = \Delta V^{-1} \cdot [1, V_T^*, \gamma_1, \dots, \gamma_P]$ are finally extracted by maximizing the log-likelihood of the experimental spike train (Brillinger, 1988; Truccolo et al., 2005; Pillow et al., 2008):

$$\hat{\theta}_{\text{th}} = \underset{\theta_{\text{th}}}{\text{argmax}} \left\{ \sum_{t \in \{\hat{t}_j\}} y_t^T \theta_{\text{th}} - \Delta T \cdot \sum_{t \in \Omega} \exp(y_t^T \theta_{\text{th}}) \right\}, \quad (3.21)$$

where $\Omega = \{t | t \notin [\hat{t}_j, \hat{t}_j + T_{\text{ref}}]\}$ is a set that excludes all the data points falling in the absolute refractory periods and the vectors y_t^T are defined as

$$y_t^T = \left[\hat{V}_{\text{model}}(t), -1, -\sum_j f^{(1)}(t - T_{\text{ref}} - \hat{t}_j), \dots, -\sum_j f^{(P)}(t - T_{\text{ref}} - \hat{t}_j) \right]. \quad (3.22)$$

With the exponential function in Equation 3.3, the log-likelihood to maximize is guaranteed to be a convex function of θ_{th} (Paninski, 2004) and both its gradient and Hessian can be formulated analytically. Consequently, the optimization problem of Equation 3.21 can be efficiently solved using the Newton-Raphson method. Alternatively, *Step 3* can be performed using the recorded potential $V_{\text{data}}(t)$ instead of $\hat{V}_{\text{model}}(t)$ in Equation 3.22. Since small inaccuracies in *Step 2* can be compensated in *Step 3*, performing the fit using $\hat{V}_{\text{model}}(t)$ generally improves spike-timing prediction.

3.4.2 Determining the amount of data required to perform GIF model parameter extraction and validation

To estimate the amount of data required to perform GIF model parameter extraction and validation, we tested our fitting procedure (Box 2) on an artificial training set generated by simulating the response of a GIF model to a fluctuating current $I(t)$. The choice of reference parameters (Fig. 3.4a-d, black) was based on previous results (Pozzorini et al., 2013). In particular, both the spike-triggered current $\eta(t)$ and the threshold movement $\gamma(t)$ were defined as a linear combination of $K = 26$ log-spaced rectangular basis functions approximating a power-law decay over 5 seconds (Pozzorini et al., 2013; Lundstrom et al., 2008). Overall, the

Chapter 3. Single-neuron characterization by means of spiking models

reference model had 59 parameters: 31 were related to the subthreshold dynamics and 28 to the firing threshold. The input current $I(t)$ used to build the artificial dataset was generated at $\Delta T^{-1} = 20$ kHz according to Eqs. 3.7-3.8 with parameters I_0 , σ_0 and $\Delta\sigma$ adjusted to obtain an average firing rate of around 10 Hz oscillating over 5 seconds between 7 and 13 Hz.

The fitting procedure illustrated in Box 2 was then applied to recover the reference parameters of the GIF model used to generate the artificial dataset (Fig. 3.4a-d, black). To estimate the amount of data required to guarantee a high degree of accuracy, this operation was repeated several times by varying the size of the training set T_{tr} (that is, the duration of the input current $I(t)$ used to generate the artificial training set). Fig. 3.4a-d shows a comparison between the reference parameters and the results obtained by fitting a training set of $T_{\text{tr}} = 10$ seconds (gray) and $T_{\text{tr}} = 100$ seconds (red). Overall, we found that 100 seconds were sufficient to accurately recover the reference parameters. To quantify the accuracy of the fit, we defined the mean error ϵ_{param} on model parameters as

$$\epsilon_{\text{param}} = \left\langle \frac{\Delta\theta_i}{|\theta_i|} \right\rangle_i, \quad (3.23)$$

where $\langle \cdot \rangle_i$ denotes the average across parameters i and $\Delta\theta_i = |\theta_i - \hat{\theta}_i|$ is the $L1$ -error between the estimated parameter $\hat{\theta}_i$ and the reference parameter θ_i . Computing ϵ_{param} as a function of T_{tr} revealed that 100 seconds were sufficient to limit the error to $\epsilon_{\text{param}} < 2.0\%$ (Fig. 3.4e, top). The great accuracy with which the fitted model was able to predict the spiking activity of the reference model ($M_{\text{d}}^* = 0.998$) confirmed the goodness of this fit (Fig. 3.4e, middle). To achieve high-throughput and perform parameter extraction *on the fly*, it is crucial to minimize the computing time (CPU time) required for the fit. Measuring the CPU time as a function of the training set duration T_{tr} (Fig. 3.4e, bottom) revealed that accurate parameter extraction from a training set of $T_{\text{tr}} = 100$ seconds requires around 60 seconds of computing. CPU times were obtained using an IntelCore i7 CPU920 at 2.67GHz with 24 GB RAM.

A second time-consuming procedure that has to be analyzed is the validation protocol. To quantify the predictive power of the fitted model, the reference model was stimulated with repetitive injections of a test current $I_{\text{test}}(t)$ generated according to Equations 3.7-3.8. To estimate the number of repetitions n_{test} and the duration T_{test} of the test current required to obtain a reliable estimate of the model predictive power, the similarity measure M_{d}^* was computed multiple times using different values of n_{test} and T_{test} (Fig. 3.4f). On average, the value of M_{d}^* was independent of both the input current duration and the number of repetitions, confirming that the spike-train metric M_{d}^* successfully eliminates the *small sample bias* (Fig. 3.4f, dashed black; see also Naud et al. (2011)). We measured the variability of M_{d}^* across validation procedures performed with different realizations of $I_{\text{test}}(t)$ and found that the reliability of M_{d}^* increased with both the number of repetitions n_{test} and the duration of the test current T_{test} (Fig. 3.4f, gray lines). Spike-triggered processes can last for several seconds (Pozzorini et al., 2013; Lundstrom et al., 2008). This sets a constraint on the minimal duration of both the test current $I_{\text{test}}(t)$ and the interstimulus interval. By taking into account these

constraints, we concluded that, while respecting high-throughput constraints, a validation protocol based on 9 injections of a 10-second current guarantees a reliable estimation of the model's predictive power.




Figure 3.4: Estimating the amount of data required to perform accurate GIF model parameter extraction and validation. a-d) GIF model parameters used to generate the artificial data (black) and recovered using a training set of $T_{tr} = 10$ s (gray) and $T_{tr} = 100$ s (red). Error bars and shaded areas represent one standard deviation obtained using 5 different datasets. In case of perfect agreement, black lines, gray lines and shaded areas are not visible. **a)** Membrane filter $\kappa(t)$. Inset: membrane timescale $\tau_m = C/g_L$. **b)** Spike-triggered current $\eta(t)$. **c)** Spike-triggered movement of the firing threshold $\gamma(t)$. **d)** Reversal potential (E_L , top left); cell resistance ($R = g_L^{-1}$, top right); threshold baseline (V_T^* , bottom left) and threshold sharpness (ΔV , bottom right). **e)** Estimation error ϵ_{param} on model parameters (upper panel), performance on spike-timing prediction M_d^* (middle panel) and computing time required for parameter extraction (lower panel) as a function of the training set size T_{tr} . Gray areas indicate one standard deviation across different artificial datasets generated using the same reference parameters. Gray and red arrows indicate the performance obtained with a training set of 10 s and 100 s, respectively. **f)** Reliability of the validation procedure as a function of the number of repetitions n_{test} and the duration T_{test} of the test current. For different values of n_{test} and T_{test} , M_d^* was computed 1000 times using different test currents. Consistent with the result that M_d^* corrects the small-sample bias, the mean value $M_d^* = 0.998$ (dashed line) obtained across repetitions of different test currents did not depend on n_{test} and T_{test} . Continuous lines represent the 0.25-quantiles of the M_d^* distribution obtained with $n_{test} = \{3, 6, 9, 12, 15\}$ (from dark to light gray) and indicate that the reliability of the measure increases with n_{test} and T_{test} .

3.4.3 Active electrode compensation

All of the *in vitro* recordings included in this study were preprocessed using AEC (Badel et al., 2008; Brette et al., 2008). AEC was implemented according to the following four-step procedure.

Step 1: We recorded the intracellular response $V_{\text{sub}}(t)$ evoked by the injection of a 10-second subthreshold current $I_{\text{sub}}(t)$. The input was generated according to Equation 3.7 with parameters $I_0=0$ pA, $\sigma(t)=75$ pA and $\tau=3$ ms and evoked small-amplitude subthreshold fluctuations around the resting potential \bar{V} . With this parameter choice, the standard deviation of $V_{\text{sub}}(t)$ was around 2-3 mV.

Step 2: The optimal linear filter $\kappa_{\text{opt}}(t)$ between the subthreshold input $I_{\text{sub}}(t)$ and the recorded signal $V_{\text{sub}}(t)$ was computed. To reduce the computing time, $\kappa_{\text{opt}}(t)$ was defined over a finite interval $[0, 200$ ms] as

$$\kappa_{\text{opt}}(t) = \sum_{m=1}^M b_m f^{(m)}(t), \quad (3.24)$$

with $\{f^{(m)}(t)\}$ being a set of $M=202$ rectangular basis functions of linearly increasing width. The parameters $b = [b_1, \dots, b_M]$ determining the shape of $\kappa_{\text{opt}}(t)$ were then estimated by solving the following multilinear regression:

$$b = (Z^T Z)^{-1} Z^T V, \quad (3.25)$$

where, using the discrete-time notation $x_t = x(t\Delta T)$ and by removing the subscripts *sub* for clarity, V is a vector whose t -th element is given by the mean-normalized membrane potential $V_t = V(t\Delta T) - \bar{V}$ and Z is a matrix made of vectors z_t^T defined as

$$z_t^T = \left[\sum_{s=0}^t f_s^{(1)} I_{t-s}\Delta T, \dots, \sum_{s=0}^t f_s^{(M)} I_{t-s}\Delta T \right], \quad (3.26)$$

with $I_t = I(t\Delta T)$. A typical filter $\kappa_{\text{opt}}(t)$ obtained from a L5 pyramidal neuron is shown in Box 4 (panel f).

Step 3: An exponential function $f(t; a_1, a_2) = a_1 \exp(-t/a_2)$ was fitted to the tail of $\kappa_{\text{opt}}(t)$ by minimizing the error $E(a_1, a_2) = \int_{T_{\text{min}}}^{\infty} (\kappa_{\text{opt}}(s) - f(s; a_1, a_2))^2 ds$ (see dashed line in Box 4, panel e). In AEC, the electrode is assumed to operate on fast timescales (< 1 ms) and the slow decay in $\kappa_{\text{opt}}(t)$ is attributed to the cell. For this reason the fit was performed with $T_{\text{min}} = 5$ ms, and the electrode filter was estimated as

$$\kappa_e(t) = \kappa_{\text{opt}}(t) - f(t; \hat{a}_1, \hat{a}_2), \quad (3.27)$$

with \hat{a}_1 and \hat{a}_2 being the optimal parameters minimizing $E(a_1, a_2)$. To improve accuracy, Steps 2 and 3 were repeated 15 times by resampling from the available data and the final electrode

filter used for AEC was obtained by averaging the results across repetitions. A typical filter $\kappa_e(t)$ obtained from a L5 pyramidal neuron is shown in Box 4 (panel f).

Step 4: For all subsequent current-clamp injections, the membrane potential $V_{\text{data}}(t)$ was estimated as:

$$V_{\text{data}}(t) = V_{\text{rec}}(t) - \int_0^{\infty} \kappa_e(s) I_{\text{ext}}(t-s) ds, \quad (3.28)$$

where $I_{\text{ext}}(t)$ is the injected current, $V_{\text{rec}}(t)$ is the recorded signal and the convolution integral on the right-hand side of Equation 3.28 approximates the voltage drop $V_e(t)$ across the electrode.

The subthreshold injection used to estimate the properties of the electrode is performed shortly before the acquisition of the *training dataset* (see Fig. 3.1). Overall, Steps 1-3 require around 62 seconds and can therefore be executed during the acquisition of the training set (Fig. 3.1). Step 4 requires less than 1 second and can be performed after training set collection without compromising high-throughput (Fig. 3.1). Since our validation protocol only relies on spike-timing prediction, AEC only needs to be applied to the training dataset.

3.4.4 Multi-compartmental model simulations

In silico recordings shown in Figure 3.2 were performed by simulating a multi compartmental model of a L5b pyramidal neuron (Hay model, Hay et al. (2011)). The model was obtained from Model DB (accession number 139653) and all simulations were performed in Neuron (Carnevale and Hines, 2006). Similar to the *in vitro* experiments, input currents $I(t)$ were generated according to Equations 3.7-3.8 (with sampling frequency $\Delta T^{-1} = 20$ kHz) and were delivered at the somatic compartment. To obtain an average firing rate fluctuating between 7 and 13 Hz, the input parameters were set to $I_0 = 520$ pA, $\sigma_0 = 320$ pA and $\Delta\sigma = 0.5$. To reproduce spike timing variability between responses to repetitive injections of the same current $I(t)$, a source of noise was included in the model by adding a zero-mean white-noise signal $\xi_{\text{w.n.}}(t)$ to $I(t)$. In order to capture the cross-correlation function between spike trains recorded *in vitro* in response to repetitive injections of the same test set current, the magnitude of the noise was set to $\sqrt{\langle \xi_{\text{w.n.}}(t)^2 \rangle} = 160$ pA. The same amount of noise was also used to generate the test dataset. GIF model and GLM parameter extraction were performed by treating the noise current $\xi_{\text{w.n.}}(t)$ as being unknown.

3.4.5 Electrophysiological recordings

All procedures in this study were conducted in conformity with the Swiss Welfare Act and the Swiss National Institutional Guidelines on Animal Experimentation for the ethical use of animals. The Swiss Cantonal Veterinary Office approved the project following an ethical review by the State Committee for Animal Experimentation.

Chapter 3. Single-neuron characterization by means of spiking models

In vitro electrophysiological recordings reported in Figure 3.3 were performed on 300 μm thick parasagittal acute slices from the right hemispheres of male P13-15 C57Bl/6J mouse brains, which were quickly dissected and sliced (HR2 vibratome, Sigmann Elektronik, Germany) in ice-cold artificial cerebrospinal fluid (ACSF) (in mM: NaCl 124, KCl 2.5, MgCl_2 10, NaH_2PO_4 1.25, CaCl_2 0.5, D-(+)-Glucose 25, NaHCO_3 25; pH 7.3 ± 0.1 , aerated with 95% O_2 , 5% CO_2), followed by a 15 minute incubation at 34 $^\circ\text{C}$ in standard ACSF (in mM: NaCl 124, KCl 2.5, MgCl_2 1, NaH_2PO_4 1.25, CaCl_2 2, D-(+)-Glucose 25, NaHCO_3 25; pH 7.4, aerated with 95% O_2 , 5% CO_2), equally used as bath solution. Cells were visualized using infrared differential interference contrast video microscopy (VX55 camera, Till Photonics, Germany and BX51WI microscope, Olympus, Japan). Somatic whole-cell current clamp recordings of layer 5 pyramidal cells in the primary somatosensory cortex were performed at 32 ± 1 $^\circ\text{C}$ with an Axon Multiclamp 700B Amplifier (Molecular Devices, USA) using 6.5-7.5 $\text{M}\Omega$ borosilicate pipettes, containing (in mM): K^+ -gluconate 110, KCl 10, ATP- Mg^{2+} 4, Na_2^+ -phosphocreatine 10, GTP- Na^+ 0.3, HEPES 10, biocytin 5 mg/ml; pH 7.3, 300 mOsm. To ensure intact axonal and dendritic arborisation, recordings were conducted in slices cut parallel to the apical dendrites.

Data were acquired at $\Delta T^{-1} = 20$ kHz using an ITC-18 digitising board (InstruTECH, USA) controlled by a custom-written software module operating within IGOR Pro (Wavemetrics, USA). Voltage signals were low-pass filtered (Bessel, 10 kHz) and not corrected for the liquid junction potential. Only cells with an access resistance < 25 $\text{M}\Omega$ (20.2 ± 3.2 $\text{M}\Omega$, $n = 10$), which was compensated throughout the recording, and a drift in the resting membrane potential < 2.5 mV (1.2 ± 0.8 mV, $n = 10$) between the start and the end of the recording were retained for further analysis. All the *in vitro* recordings included in this study were preprocessed using AEC (Brette et al. (2008); Badel et al. (2008), see Box 4).

3.4.6 Generalized Linear Model

All GIF model performance included in this study were compared against the ones of a standard Generalized Linear Model (GLM) (Truccolo et al., 2005; Pillow et al., 2008). In the GLM, spikes are emitted stochastically according to the following conditional intensity

$$\lambda_{\text{GLM}}(t|I, \{\hat{t}_j\}) = \lambda_0 \cdot \exp \left(\int_0^t \kappa_{\text{GLM}}(s) I(t-s) ds - \sum_{\hat{t}_j} h_{\text{GLM}}(t - \hat{t}_j) \right). \quad (3.29)$$

In contrast to the membrane filter of the GIF model, the linear filter $\kappa_{\text{GLM}}(t)$ is not assumed to be exponential, but is extracted from experimental data through linear expansion in rectangular basis functions. Moreover, the GLM accounts for spike-history effects with a unique filter $h_{\text{GLM}}(t)$.

To obtain a fair comparison, $h_{\text{GLM}}(t)$ was expanded using the same basis functions as used for $\gamma(t)$ in the GIF model and the number of basis functions used for $\kappa_{\text{GLM}}(t)$ was such that, in total, the two models had the same number of parameters. Given the input current $I(t)$ and the observed spike train $S_{\text{data}}(t)$, GLM parameters $\theta_{\text{GLM}} = \{\lambda_0, \kappa_{\text{GLM}}(t), h_{\text{GLM}}(t)\}$ were

extracted with standard methods (Truccolo et al., 2005; Pillow et al., 2008) by maximizing the model log-likelihood $L(\theta_{\text{GLM}}) = \log p(S_{\text{data}}|I, \theta_{\text{GLM}})$.

3.5 Author contributions

C.P., S.M., W.G. and C.K. conceived the study; C.P. wrote the initial draft; S.M. analyzed the data; O.H. performed the electrophysiological experiments; R.N., C.K. and W.G. supervised the project. All of the authors worked on the manuscript. This project was funded by the Swiss National Science Foundation (SNSF, grant numbers 200020 132871/1 and 200020 147200; C.P. and S.M.), by the European Community's Seventh Framework Program (BrainScaleS, grant no. 269921, S.M.), the Human Brain Project (grant agreement 604102; S.M. and C.P.). O.H. was funded by a grant from the EPFL to the LNMC. R.N. was funded by FQRNT.

4 Enhanced temporal coding by nonlinear threshold dynamics in neocortical neurons

Christian Pozzorini[†], Skander Mensi[†], Olivier Hagens and Wulfram Gerstner; [†]equal contribution (in preparation)

4.1 Introduction

Over which timescale do cortical neurons integrate their synaptic inputs? The answer to this question is of fundamental importance because it determines which coding strategies can be supported by the brain (Ratté et al., 2013). If the timescale of integration is longer compared to the mean interspike interval of afferent spikes trains, then multiple action potentials received from the same presynaptic neuron can lead to a single output spike and neurons operate as temporal integrators supporting a rate code in which the temporal pattern of spikes conveys little information (Shadlen and Newsome, 1994). Alternatively, if the effective timescale of integration is shorter than the average interspike interval, output spikes are primarily driven by the coincident arrival of presynaptic spikes sent by different cells and single neurons operate as coincidence detectors supporting a temporal code in which most of the information is carried by the precise timing of action potentials (Abeles, 1982; König et al., 1996). Despite decades of research, whether *in vivo* cortical neurons operate as temporal integrators or coincidence detectors remains a controversial topic (Softky and Koch, 1993; Rossant et al., 2011b; Bruno, 2011; Stanley, 2013; Ratté et al., 2013).

Theoretical studies of integrate-and-fire models responding to *in vivo*-like fluctuating currents concluded that a smooth transition between coincidence detection and temporal integration occurs when the average strength μ_1 of the input current reaches rheobase (Gerstner and Kistler, 2002). In agreement with this theoretical result, the average firing rate of cortical neurons have initially been shown to lose sensitivity to rapid input fluctuations at increasing driving currents (Rauch et al., 2003; Giugliano et al., 2004; La Camera et al., 2006). This view has recently been challenged by *in vitro* recordings demonstrating that the output firing rate of pyramidal neurons from prefrontal cortex (PFC) (Arsiero et al., 2007), somatosensory cortex

(SSC) (Higgs et al., 2006) and hippocampus (Fernandez et al., 2011) always increases with the amplitude of rapid input fluctuations σ_I . These results indicate that, in response to strong depolarizing offsets, pyramidal neurons adapt their intrinsic dynamics in such a way as to operate as coincidence detectors over a broad range of input statistics. A similar behavior can be qualitatively reproduced by Hodgkin-Huxley models with decreased Na^+ conductance (Lundstrom et al., 2008) or augmented with slow Na^+ -channel inactivation (Arsiero et al., 2007; Fleidervish et al., 1996). While these studies indicate that the firing threshold dynamics plays a crucial role, the mechanisms by which single neurons maintain sensitivity to rapid input fluctuations remain unclear.

The membrane potential at which spikes are initiated is highly variable both *in vitro* and *in vivo* (Azouz and Gray, 2000, 2003). Threshold variability is not random but probably results from a complicated dynamics which likely plays a functional role (Brette, 2011; Carandini and Ferster, 2000; Wilent and Contreras, 2005; Cardin et al., 2010). Many studies have demonstrated that the voltage threshold for spike initiation depends not only on the average value of the membrane potential (Azouz and Gray, 2003), but also on previous interspike intervals (Henze and Buzsaki, 2001; Chacron et al., 2007; Jolivet et al., 2006a; Badel et al., 2008; Mensi et al., 2011) and on the depolarization rate before a spike (that is, the speed at which the firing threshold is approached) (Azouz and Gray, 2000; Henze and Buzsaki, 2001; Azouz and Gray, 2003; Wilent and Contreras, 2005). When single neurons are stimulated with current ramps of different slopes, rapid rates of depolarization are often associated with lower thresholds (Ferragamo and Oertel, 2002; Higgs and Spain, 2011). While in rat pyramidal neurons this phenomenon results from the activation of low-threshold Kv1 channels (Higgs and Spain, 2011), theoretical studies suggest that a fast coupling between voltage threshold and subthreshold membrane potential could also result from fast Na^+ -channel inactivation (Hodgkin and Huxley, 1952; Platkiewicz and Brette, 2011; Wester and Contreras, 2013).

In the present study, we recorded the *in vitro* response of layer 5 pyramidal (Pyr) neurons to fluctuating currents generated by systematically varying μ_I and σ_I . In agreement with previous results, we found that the average firing rate response remained sensitive to rapid input fluctuations over a broad range of depolarizing offsets μ_I . By independently fitting different Generalized Linear Models (GLM) to the spiking activity evoked by currents of increasing μ_I , we found that Pyr neurons progressively shortened their effective timescale of integration. Moreover, the GLM spike-history filter also changed with μ_I . To identify and understand the mechanisms underlying these complex forms of adaptation, we designed a new spiking model in which the firing threshold is nonlinearly coupled with the subthreshold membrane potential and also linearly dependent on the spike history. Results obtained by extracting model parameters from intracellular recordings demonstrate that a nonlinear interaction between spike-dependent and voltage-dependent threshold movements dynamically regulates the timescale of somatic integration in such a way as to enhance coincidence detection over a wide range of input statistics.

4.2 Results

In vivo, neocortical neurons constantly receive barrages of excitatory and inhibitory inputs. To understand how synaptic inputs are transformed into output spike trains, single neurons can be tested *in vitro* with somatic injections of rapidly fluctuating currents (Destexhe et al., 2001; Rauch et al., 2003; Richardson, 2004). *In vivo*-like fluctuating currents $I(t)$ mimic the net input received at the soma of a postsynaptic neuron and are generally described by Ornstein-Uhlenbeck processes (i.e., filtered Gaussian processes) characterized by a mean μ_I , corresponding to the average intensity of $I(t)$, and a standard deviation σ_I , defining the magnitude of the input fluctuations (Rauch et al., 2003; Arsiero et al., 2007). Since *in vivo* both the strength and the synchrony of afferent spike-trains can vary over time, single neurons are likely to receive inputs with varying statistics (Tan et al., 2014).

4.2.1 Enhanced sensitivity to rapid input fluctuations

To study single-neuron computation over a broad range of input statistics, we intracellularly recorded the response of L5 Pyr neurons evoked *in vitro* by a set of 5-second currents generated by independently varying the parameters μ_I and σ_I (Fig. 4.1a-c). *In vivo*-like fluctuating currents were generated according to Equation 4.9 and injected at the soma of L5 Pyr neurons of mouse SSC (see Materials and Methods). In agreement with previous results from Pyr neurons of rat PFC (Arsiero et al., 2007; Thurley et al., 2008) and hippocampus (Fernandez et al., 2011), we found that augmenting the magnitude of input fluctuations σ_I significantly increased the output firing rate over the entire range of depolarizing steps μ_I that were tested (Fig. 4.1d,e). This result is surprising because, according to standard theories, rapid input fluctuations are expected to affect the output firing rate only in the subthreshold regime, that is when the mean input μ_I by itself is not sufficiently strong to evoke action potentials (Ricciardi, 1977). In this regime, spikes are always driven by rapid input fluctuations and single neurons operate as temporal detectors of coincidence inputs. The results shown in Figure 4.1 indicate that L5 Pyr neurons never lose sensitivity to input fluctuations and adapt their intrinsic dynamics in order to act as coincidence detectors even in response to strong depolarizing offsets μ_I .

To characterize the mechanisms underlying enhanced sensitivity to input fluctuations, we independently fitted a Generalized Linear Model (GLM, see Truccolo et al. (2005); Pillow et al. (2008)) to different datasets obtained by splitting our intracellular recordings according to μ_I (see Materials and Methods). In the GLM, spikes are generated stochastically according to a point process whose firing intensity depends on the input current as well as on previous action potentials (Fig. 4.2a). Briefly, the input current is first passed through a linear filter $\kappa_{\text{GLM}}(t)$ and then transformed into a firing intensity by an exponential nonlinearity. Each time an action potential is fired, an adaptation process $h_{\text{GLM}}(t)$ is triggered that accounts for spike-history dependence. When a GLM is used to characterize the response of a single neuron *in vivo*, the linear filter $\kappa_{\text{GLM}}(t)$ is generally interpreted as a receptive field. In the case




Figure 4.1: Pyr neurons maintain sensitivity to rapid input fluctuations over a wide range of depolarizing offsets. **a)** Staining of a biocytin-filled L5 Pyr neuron. **b)** Neurons were tested with a set of 5-second fluctuating currents generated by systematically varying the magnitude of the DC component, μ_I (dashed black), and the magnitude of input fluctuations, σ_I . **c)** Typical response of a L5 Pyr neuron to a somatic injection of the current shown in panel **b** ($\mu_I=0.27$ nA, $\sigma_I=50$ pA). **d)** Steady-state firing rate f of a typical Pyr neuron as a function of the mean input μ_I (i.e., $f - \mu_I$ curve). Different lines indicate different levels of input fluctuations σ_I . For each combination of input parameters (μ_I , σ_I), three different 5-s recordings were performed. Firing rates were estimated by discarding the transient response observed during the first second. Error bars indicate one standard deviation across repetitions and are often not visible reflecting high reliability. **e)** Summary data obtained in different Pyr neurons ($n=6$) by computing the percentage change in steady-state firing rate obtained by increasing the input standard deviation σ_I . Changes were computed with respect to $\sigma_I=0$ pA by only considering the strongest input $\mu_I^{(8)}$. Each set of open circles represents data from a particular cell. Bar plots represent mean and standard deviation across cells. Increasing σ_I from 50 to 100 pA ($n=6$, paired Student t -test, $t=4.4$, $p=6.7 \cdot 10^{-3}$) and from 100 to 150 pA ($n=6$, paired Student t -test, $t=4.5$, $p=6.5 \cdot 10^{-3}$) significantly increased the firing rate.

of *in vitro* current-clamp injections, $\kappa_{\text{GLM}}(t)$ defines the feature of the somatic input current to which the neuron is most sensitive. In particular, the timescale τ_{GLM} over which the filter $\kappa_{\text{GLM}}(t)$ decays defines the temporal interval over which two incoming spikes are considered to be coincident.

By comparing the GLM filters extracted under different stimulus conditions, we found that both $\kappa_{\text{GLM}}(t)$ and $h_{\text{GLM}}(t)$ changed with μ_I (Fig. 4.2b,c). In agreement with the ability of Pyr neurons to maintain sensitivity to rapid input fluctuations over a broad range of depolarizing offsets, increasing μ_I resulted in a progressive shortening of the linear filter $\kappa_{\text{GLM}}(t)$ (Fig. 4.2b).




Figure 4.2: A comparison between GLM filters extracted from the spiking response to different input statistics reveals complex forms of adaptation. **a)** Schematic representation of the GLM. The input current $I(t)$ is first low-pass filtered through $\kappa_{\text{GLM}}(t)$ and then transformed by an exponential nonlinearity into a firing intensity $\lambda_{\text{GLM}}(t)$. Spikes are fired stochastically according to $\lambda_{\text{GLM}}(t)$. Each time an action potential is fired, a feedback process $h_{\text{GLM}}(t)$ is triggered that accounts for spike-history dependence. **b-c)** Average GLM filters $\kappa_{\text{GLM}}(t)$ and $h_{\text{GLM}}(t)$ extracted from 6 Pyr neurons by splitting the experimental data in eight groups according to $\mu_I = \{\mu_I^{(1)}, \mu_I^{(2)}, \dots, \mu_I^{(8)}\}$. For clarity, only three filters are shown: low input ($\mu_I^{(2)} = 64 \pm 7$ pA, blue), medium input ($\mu_I^{(4)} = 190 \pm 30$ pA, gray), strong input ($\mu_I^{(7)} = 390 \pm 50$ pA, red). Shaded areas indicate one standard deviation across neurons. **b)** Average linear filters $\kappa_{\text{GLM}}(t)$. Inset: timescale τ_{GLM} of somatic integration quantified by fitting the filters $\kappa_{\text{GLM}}(t)$ with an exponential function. **c)** Average spike-history filters $h_{\text{GLM}}(t)$. **d)** The average membrane timescale τ_m (solid black) and the effective timescale τ_{GLM} of somatic integration (dashed black) are plotted as function of the input strength μ_I . The shaded areas represent one standard deviation across neurons. Colored lines represent the timescales τ_{GLM} extracted from the three filters shown in panel *b*.

We quantified this result by fitting the GLM linear filters with a single-exponential function and found that, with increasing μ_I , the timescale τ_{GLM} over which Pyr neurons integrate their somatic inputs was drastically reduced (Fig. 4.2d and inset of Fig. 4.2b).

Single neurons can in principle shorten the timescale of somatic integration by increasing their total membrane conductance (Destexhe et al., 2003; Badel et al., 2008). To test this hypothesis,

we measured the membrane timescale τ_m as a function of μ_I by fitting the intracellular response to different depolarizing offsets with a Generalized Integrate-and-Fire (GIF) model that we previously introduced (see Materials and Methods and also Mensi et al. (2011) and Pozzorini et al. (2013)). Surprisingly, we found that changes in the passive properties of the membrane only accounted for part of the reduction observed in τ_{GLM} (Fig. 4.2d). This result indicates that increasing μ_I not only induced a net increase of the membrane conductance, but also recruited a different cellular mechanism that further enhanced sensitivity to rapid input fluctuations without affecting τ_m .

4.2.2 The voltage threshold for spike initiation depends on the input statistics

Previous studies found that the particular input-output function we experimentally observed in Pyr neurons (see Fig. 4.1d) can be qualitatively reproduced by Hodgking-Huxley models in which the ratio between sodium and potassium maximal conductances is low (Arsiero et al., 2007; Lundstrom et al., 2008). These results indicate that enhanced sensitivity to input fluctuations might be mediated by a mechanism that reduces the number of Na^+ -channels available for spike initiation or, more generally, by a process that regulates the voltage threshold for spike initiation. An active role of the firing threshold in enhancing sensitivity to input fluctuations would additionally explain the discrepancy between the membrane timescale τ_m and the effective timescale of somatic integration τ_{GLM} .

Using standard methods, we extracted from intracellular recordings the voltages at which individual action potentials were initiated (see Materials and Methods). Consistent with earlier results showing that the firing threshold adaptively rises after the emission of previous action potentials (Henze and Buzsaki, 2001; Badel et al., 2008; Mensi et al., 2011), the voltage threshold for spike initiation always increased with the mean input μ_I (Fig. 4.3a,e-f). However, at odds with the hypothesis of a threshold dynamics entirely governed by positive spike-triggered changes, we also found that larger input fluctuations, while evoking higher firing rates, significantly reduced the firing threshold (Fig. 4.3a-b, g-h). Since, in our protocol, increased input fluctuations σ_I translates into faster currents (c.f. Eq. 4.9), we reasoned that this negative correlation could reflect a dependence of the firing threshold on the rate of rise of the membrane potential preceding a spike (i.e., the speed dV/dt at which the membrane potential reaches the firing threshold) (Azouz and Gray, 2000, 2003). Confirming this hypothesis, we found that the firing threshold was negatively correlated with the membrane depolarization rate (Fig. 4.3c, see Materials and Methods). A theoretical study has recently demonstrated that the latter dependence could in principle arise from a nonlinear coupling between firing threshold and subthreshold membrane potential (Platkiewicz and Brette, 2010). In agreement with this theoretical result, we also found a nonlinear relationship between firing threshold and average subthreshold membrane potential (Fig. 4.3d, see Materials and Methods).




Figure 4.3: Standard analysis of intracellular recordings suggests an intricate dependence of the firing threshold on the input statistics. **a)** Average voltage threshold for spike initiation in a typical Pyr neuron as a function of μ_I . Different gray levels indicate different values of σ_I (see legend in panel **d**). Dashed lines show four linear regressions performed on experimental data with different σ_I . The regression slopes s_{LR} observed in 6 different neurons were significantly larger than zero ($n=24$, $s_{LR} = 22.4 \pm 2.7$ mV/nA, Student t -test, $t=39.9$, $p < 10^{-6}$). **b)** Summary data from six L5 Pyr neurons showing the relative change in voltage threshold induced by an increase in σ_I . For each cell, changes were computed with respect to $\sigma_I = 0$ pA by averaging the results obtained for all depolarizing offsets μ_I . Each set of open circles represents data from a particular cell. Bar plots represent mean and standard deviation across cells. Increasing σ_I from 50 to 100 pA ($n=6$, paired Student t -test, $t=7.5$, $p=6.8 \cdot 10^{-4}$) and from 100 to 150 pA ($n=6$, paired Student t -test, $t=6.4$, $p=1.3 \cdot 10^{-3}$) significantly reduced the voltage threshold for spike initiation. **c)** Average voltage threshold as a function of the membrane potential depolarization rate preceding a spike. Only trials in which the average firing rate was ≤ 10 Hz are shown. The dashed line indicates a linear regression. The regression slopes s_{LR} observed in 6 different cells were significantly smaller than zero ($n=6$, $s_{LR} = -6.2 \pm 2.0$ ms, Student t -test, $t=-7.1$, $p = 8.6 \cdot 10^{-4}$). **d)** Average voltage threshold as a function of the mean subthreshold membrane potential. Different gray levels indicate different σ_I . The dashed line shows the identity diagonal delimitating the area (gray) in which the membrane potential is larger than the firing threshold. **e-f)** Subset of raw data showing that the firing threshold is positively correlated with μ_I . **e)** For each action potential shown in panel **f**, the voltage derivative $\dot{V}_{data}(t)$ is plotted as a function of $V_{data}(t)$. The voltage threshold for spike initiation was defined as the voltage where \dot{V}_{data} crossed 10 mV/ms from below (dashed black). Different colors indicate different values of μ_I (color code as in panel **f**). **f)** Intracellular recordings obtained with three different values of μ_I (and $\sigma_I = 100$ pA). Colored dots indicate the membrane potential at spike initiation. For clarity the y-axis was truncated at -20 mV. **g-h)** Subset of raw data showing that the firing threshold is negatively correlated with σ_I . The same analysis shown in panels

e-f was performed on intracellular recordings obtained in response to three values of σ_I ($\mu_I = 0.2$ nA). Different colors indicate different values of σ_I (color code as in panel *h*). In panels *a*, *c* and *d*, each data point *j* was computed by averaging the results obtained by analyzing the responses to three different 5-second currents with the same statistics ($\mu_I^{(j)}$, $\sigma_I^{(j)}$). As in Figure 4.1, action potentials recorded in the first second of each repetition were not considered. Error bars indicate one standard deviation across repetitions. Data are from the same cell shown in Figure 4.1.

Overall, the results reported in Figure 4.3 demonstrate that the voltage threshold for spike initiation is highly variable. This variability unlikely results from channel noise and can in part be explained by different covariates. We concluded that the firing threshold evolves according to a nontrivial dynamics, whose understanding requires mathematical modeling and could explain the origin of enhanced sensitivity to rapid input fluctuations.

4.2.3 Modeling the firing threshold dynamics

In the standard Hodgkin-Huxley model (HH, see Hodgkin and Huxley (1952)), the sodium current I_{Na} responsible for spike initiation is gated by two independent variables, m and h , that describe Na^+ -channel activation and inactivation, respectively (Fig. 4.4a). Sodium activation occurs on very short timescales and can therefore be considered as instantaneous (Fourcaud-Trocmé et al., 2003; Badel et al., 2008; Platkiewicz and Brette, 2010). It follows that, at spike onset, I_{Na} is accurately approximated by an exponential function of the membrane potential (Platkiewicz and Brette, 2011):

$$I_{\text{Na}} \propto h \exp\left(\frac{V - V_{\text{T}}^*}{k_{\text{a}}}\right) = \exp\left(\frac{V - (V_{\text{T}}^* - k_{\text{a}} \log h)}{k_{\text{a}}}\right), \quad (4.1)$$

where V_{T}^* is a constant, k_{a} is a biophysical parameter describing the sharpness of sodium channel activation $m_{\infty}(V)$ and $\theta = V_{\text{T}}^* - k_{\text{a}} \log h$ defines a smooth threshold for spike initiation. Since in the HH model Na^+ -channel inactivation follows a first-order kinetics $\tau_{\text{h}} \dot{h} = -h + h_{\infty}(V)$, an accurate model of the firing threshold $\theta(t)$ is given by the following differential equation (Platkiewicz and Brette, 2011):

$$\tau_{\theta} \dot{\theta} = -\theta + \theta_{\infty}(V), \quad (4.2)$$

where $\tau_{\theta} = \tau_{\text{h}}$. By modeling the steady-state inactivation curve $h_{\infty}(V)$ with an inverse sigmoidal function $h_{\infty}(V) = \left(1 + \exp\left(\frac{V - V_{\text{i}}}{k_{\text{i}}}\right)\right)^{-1}$, Platkiewicz and Brette (2011) further predicted that the coupling between firing threshold and membrane potential resulting from fast Na^+ -channel inactivation should be correctly described by a smooth linear rectifier function $\theta_{\infty}^{\text{Na}}(V)$ defined as (Fig. 4.4b):

$$\theta_{\infty}^{\text{Na}}(V) = V_{\text{T}}^* - k_{\text{a}} \log h_{\infty}(V) = V_{\text{T}}^* + k_{\text{a}} \log \left(1 + \exp\left(\frac{V - V_{\text{i}}}{k_{\text{i}}}\right)\right). \quad (4.3)$$

Depending on the half-inactivation voltage V_i and on the asymptotic slope $\theta_{\text{slope}}^{\text{Na}} = k_a/k_i$, Equations 4.2 and 4.3 can provide a theoretical explanation for the negative correlation between the firing threshold and the depolarization rate of the membrane potential preceding a spike (Platkiewicz and Brette, 2011; Higgs and Spain, 2011). Indeed, if $\theta_{\text{slope}}^{\text{Na}} > 0$ and $V_i < V_T^*$, all membrane depolarizations occurring on a slower rate than the characteristic timescale τ_θ on which Na^+ -channels inactivate will reduce the number of Na^+ -channels available for spike initiation. Consequently, compared to fast inputs, slow currents will evoke action potentials that are initiated at larger voltages (Fig. 4.4b,c and g). On the other hand, if $h_\infty(V)$ is shifted towards more depolarized potentials such that $V_i - V_T^* \gg k_i$ and $k_a > 0$ (Fig. 4.4d), Na^+ -channels do not inactivate at subthreshold potentials and spikes initiate at the same voltage threshold V_T^* , independently of the depolarization rate (Fig. 4.4e-g). In summary, to effectively modulate the voltage threshold for spike initiation, the inactivation profile of Na^+ -channel has to be such that, in the $V - \theta$ plane, the voltage-dependent part of $\theta_\infty^{\text{Na}}(V)$ is not *masked* on the right-hand side of the diagonal $V = \theta$ (compare Fig. 4.4b and Fig. 4.4e and see also Platkiewicz and Brette (2011)).

With the final goal of understanding whether the adaptation mechanisms underlying enhanced sensitivity to rapid input fluctuations are mediated by the firing threshold dynamics, we fitted our intracellular recordings using a new spiking model obtained by extending the GIF model (Mensi et al., 2011; Pozzorini et al., 2013) with a nonlinear coupling between firing threshold and membrane potential (Platkiewicz and Brette, 2011). We refer to this model as iGIF, where *i* stands for *inactivating* (Fig. 4.5a, see Materials and Methods). In the iGIF model, spikes are produced stochastically according to a firing intensity which exponentially depends on the instantaneous difference between the membrane potential V and firing threshold V_T (Gerstner and van Hemmen, 1992; Jolivet et al., 2006a). As in the GIF model, the dynamics of the membrane potential is modeled as a leaky integrator augmented with a spike-triggered adaptation current I_A . To capture the reduction in the membrane timescale observed in Figure 4.2, the adaptation current I_A was defined by the following conductance-based model (Dayan and Abbott, 2001):

$$I_A(t) = \sum_{\hat{t}_j < t} \eta(t - \hat{t}_j) \cdot (V - E_R), \quad (4.4)$$

where E_R is a reversal potential and $\eta(t - \hat{t}_j)$ describes the time course of the conductance change triggered by the emission of an action potential at time \hat{t}_j . In the iGIF model, the firing threshold V_T is given by:

$$V_T(t) = \theta(t) + \sum_{\hat{t}_j < t} \gamma(t - \hat{t}_j), \quad (4.5)$$

where the dynamics of $\theta(t)$ is as in Equation 4.2 and implements a nonlinear coupling between V_T and V . In Pyr neurons, the firing threshold has been previously shown to adaptively increase after the emission of previous spikes. To capture this phenomenon, the iGIF model features a spike-triggered movement of the firing threshold $\gamma(t)$. Similar to $\eta(t)$, this function




Figure 4.4: Simplified integrate-and-fire model illustrating firing threshold modulation by fast Na^+ -channel inactivation. **a)** Steady-state activation (black) and inactivation (red) functions describing the dynamics of Na^+ -channels responsible for spike initiation in the HH model. **b)** Phase plane analysis illustrating the dynamics of the inactivating leaky integrate-and-fire model (iLIF Platkiewicz and Brette (2011)) consisting of a standard LIF model augmented with a dynamic threshold θ defined as in Equations 4.2 and 4.3. In the iLIF model, a spike is deterministically emitted each time the membrane potential V crosses the firing threshold θ (dashed black line). After each spike, both V and θ are reset to low values (open circle). The steady-state firing threshold $\theta_\infty^{\text{Na}}(V)$ (red) of the iLIF model was obtained by transforming $h_\infty(V)$ in panel *a* according to Equation 4.3. The iLIF model responses to different ramps of current (see panel *g*) are represented in the phase plane (gray lines). Different gray levels correspond to ramps of current with different slopes (color code as in panels *c* and *g*). As soon as a trajectory entered the right hand side of the phase-plane, the membrane potential was artificially set to a large value to represent the emission of an action potential. **c)** In the iLIF model shown in panel *b*, current ramps of different slopes (see panel *g*) elicit spikes at different voltage thresholds (dashed black). Gray dots correspond to the trajectories shown in panel *b*. **d-e)** Same analysis as in panels *a-c*, but for the case in which Na^+ -channel inactivation is shifted towards more depolarized potentials. Since the voltage-dependent part of $\theta_\infty^{\text{Na}}(V)$ (blue) lies on the right-hand side of the diagonal $V = \theta$, responses to current ramps with different slopes are on top of each other and spikes originate at around the same voltage threshold V_T^* . **g)** iLIF model response to different ramps of current. Top: input current. Middle: iLIF model response for the case in which Na^+ -channel inactivation starts before Na^+ -channel activation (see panels *a-c*). Black: membrane potential $V(t)$; red: firing threshold $\theta(t)$. The same trajectories $\{V(t), \theta(t)\}$ are represented in panel *b*. Since Na^+ -channels inactivation is not instantaneous, but occurs on a characteristic timescale τ_θ , fast inputs (left) elicit spikes that, in comparison to slow inputs (right), initiate at lower membrane potentials (see panel *b*). Bottom: iLIF model response for the case in which Na^+ -channel inactivation starts after Na^+ -channel activation (see panels *d-f*).

describes changes in V_T induced by previous action potentials and constitutes an additional mechanism for spike-frequency adaptation. A spike-dependent movement of the firing threshold could in principle be accounted for by increasing $\theta(t)$ after each spike (Fontaine et al., 2014a). However, such a model would implicitly assume that spike-dependent and voltage-dependent threshold changes occur on the same timescale τ_θ . To avoid this assumption, the variable $\theta(t)$ is reset to V_T^* after each spike and all threshold changes induced by previous spikes are included into $\gamma(t)$.

Importantly, the functional shape of $\eta(t)$, $\gamma(t)$ and $\theta_\infty(V)$, along with all the other iGIF model parameters, were extracted from intracellular recordings using a new nonparametric fitting procedure (see Materials and Methods). In what follows, we refer to the iGIF model with parameters extracted using the nonparametric method as iGIF-NP, where NP stands for *nonparametric*.

4.2.4 iGIF model parameter extracted from intracellular recordings reveals a nonlinear coupling between membrane potential and firing threshold

In agreement with the results obtained by measuring the membrane timescale τ_m as a function of the mean input μ_1 , the passive properties of the membrane were characterized by a timescale $\tau_0 = 35.3 \pm 8.6$ ms (Fig. 4.5b). Moreover, spike-triggered conductance changes were always positive (Fig. 4.5c) and associated with a low reversal potential $E_R = -57.0 \pm 3.9$ mV. When displayed on log-log scales, the decay of the spike-triggered threshold movement $\gamma(t)$ was approximatively linear over several orders of magnitude (Fig. 4.5d). This result is in agreement with previous findings that, in Pyr neurons, spike-frequency adaptation does not have a preferred timescale, but is characterized by a power-law decay (Lundstrom et al., 2008; Pozzorini et al., 2013). While avoiding *a priori* assumptions about the existence of a coupling between firing threshold and membrane potential, as well as about the underlying biophysical mechanisms, our nonparametric method allowed us also to extract $\theta_\infty(V)$ directly from intracellular recordings (see Materials and Methods). We found that, in Pyr neurons, firing threshold and subthreshold membrane potential were indeed nonlinearly coupled (Fig. 4.5e, black). Moreover, the functional shape of $\theta_\infty(V)$ was in striking agreement with the theoretical prediction based on fast Na^+ -channel inactivation (Platkiewicz and Brette, 2011).

Since the value of the coupling timescale $\tau_\theta = 8.6 \pm 3.0$ ms (Fig. 4.5f) was also consistent with previous measurements of fast Na^+ -channel inactivation (see, e.g., McCormick and Huguenard (1992)), we used the intracellular recordings to fit a new iGIF model, called iGIF-Na, in which $\theta_\infty(V)$ was assumed *a priori* to be the smooth rectifier function $\theta_\infty^{\text{Na}}(V)$ defined in Equation 4.3. For that, a different maximum likelihood procedure was used allowing for the extraction of the biophysical parameters k_a , k_i and V_i , along with all the other iGIF-Na model parameters (see Materials and Methods). Both the spike-triggered movement of the firing threshold $\gamma(t)$ (Fig. 4.5d, red) and the nonlinear coupling $\theta_\infty^{\text{Na}}(V)$ (Fig. 4.5e, red) extracted by fitting the iGIF-Na model confirmed the results obtained with the nonparametric method (Fig.

Chapter 4. Enhanced temporal coding by nonlinear threshold dynamics

4.5d,e, black). Moreover, the asymptotic slope $\theta_{\text{slope}}^{\text{Na}} = k_a/k_i$ of the threshold coupling was very close to one (Fig. 4.5g, bottom).




Figure 4.5: Inactivating Generalized Integrate-and-Fire model (iGIF) with parameters extracted from intracellular recordings in Pyr neurons. **a)** Schematic representation of the iGIF model. The input current is first low-pass filtered by the *Passive membrane filter* $\kappa_0(t) = \Theta(t)C^{-1}e^{-\frac{t}{\tau_0}}$. The resulting signal models the subthreshold membrane potential $V(t)$ and, after subtraction of the firing threshold $V_T(t)$, is transformed into a firing intensity $\lambda(t)$ by the exponential *Escape-rate nonlinearity*. Spikes are emitted stochastically and elicit both a *Spike-triggered conductance* $\eta(t)$ and a *Spike-triggered threshold movement* $\gamma(t)$. In the iGIF model, but not in the GIF model, the firing threshold $V_T(t)$ is coupled to the subthreshold membrane potential (dashed circuit). For that, the membrane potential $V(t)$ is first passed through the nonlinear *Threshold coupling* function $\theta_\infty(V)$ and then low-pass filtered by the *Threshold filter* $\kappa_\theta(t) = \Theta(t)\tau_\theta^{-1}e^{-\frac{t}{\tau_\theta}}$. **b-e)** Average parameters extracted from 6 Pyr neurons. Black: iGIF-NP, red: iGIF-Na. Gray areas indicate one standard deviation across cells. **b)** Passive membrane filter $\kappa_0(t)$. Inset: passive membrane timescale τ_0 . Open circles: results from individual cells. Bar plot:

mean and standard deviation. **c)** Spike-triggered conductance $\eta(t)$. Inset: same data on log-log scales. **d)** Spike-triggered threshold movement $\gamma(t)$. **e)** Nonlinear threshold coupling $\theta_\infty(V)$. In the absence of spikes, the voltage-dependent part of $\theta_\infty(V)$ lies on the right hand side of the spiking boundary $V = \theta$ (dashed black). Spike-triggered movements of the firing threshold shift the spiking boundary to the right (dashed gray, θ). Inset: percentage increase in model log-likelihood (LL) computed with respect to the GIF model and on the same data used for the fit. Open circles: model performance on the same cell. Bar plots: mean and standard deviation across neurons. **f)** Top: LL percentage increase of the iGIF-NP model as a function of τ_θ . Red circle: optimal timescale τ_θ . Data are shown from a typical neuron. Bottom: optimal timescales τ_θ extracted from 6 Pyr neurons. **g)** Top: LL percentage increase of the iGIF-Na model as a function of k_i and V_i . The LL increases from dark to light red. Bottom: optimal parameters extracted from 6 different Pyr neurons. The mean and the standard deviation across cells are shown in red.

Despite having less parameters, the log-likelihood of the iGIF-Na model was not significantly different from that of the iGIF-NP model (Fig. 4.5e, inset). This result provides additional evidence for the hypothesis that the biophysical mechanism underlying the nonlinear coupling between firing threshold and membrane potential is fast Na^+ -channel inactivation. In the followings, we just work with the iGIF-Na model, which, for simplicity, will be referred to as iGIF model.

4.2.5 The iGIF model captures enhanced sensitivity to rapid input fluctuations and predicts spikes with millisecond precision

To verify whether the iGIF model was able to capture enhanced sensitivity to rapid input fluctuations, we repeated our experimental paradigm *in silico* by testing the iGIF model with a set of 5-second currents generated by systematically varying the input parameters μ_I and σ_I (Fig. 4.6a-c). We compared the steady-state firing rate response of the model against experimental data and found that, despite its relative simplicity, the iGIF model accurately captured the behavior of Pyr neurons over a broad range of input parameters (Fig. 4.6a). In particular, the iGIF model exhibited enhanced sensitivity to input fluctuations throughout the entire set of depolarizing currents μ_I that were tested and reproduced the average firing rate response with an accuracy of $\epsilon_{\text{rate}} = 1.0 \pm 0.2$ Hz. Notably, the iGIF model also captured the complex dependence of the firing threshold on input statistics. In particular, the voltage threshold at which spikes were initiated was positively correlated with μ_I (Fig. 4.6b) and negatively correlated with σ_I (Fig. 4.6b,c).

To further appreciate the importance of modeling the nonlinear coupling between membrane potential and firing threshold, we also fitted the experimental data with our previous Generalized Integrate-and-Fire model (GIF, Mensi et al. (2011)). The GIF model differs from the iGIF model simply because its firing threshold dynamics only depends on the spike-history and not on the membrane potential (see Materials and Methods). As expected, the GIF model could not capture the firing rate dependence on σ_I and was less accurate in reproducing the firing rates observed at steady-state ($\epsilon_{\text{rate}} = 1.7 \pm 0.3$ Hz see Fig. 4.6d). Finally, the strong mismatch

Chapter 4. Enhanced temporal coding by nonlinear threshold dynamics

between the firing threshold measured in the experiments and produced by the GIF model (Fig. 4.6e-f), demonstrates that a spiking model in which the firing threshold dynamics simply depends on previous action potentials is not sufficient to capture the spiking activity of Pyr neurons over a wide range of input statistics.




Figure 4.6: The iGIF model captures enhanced sensitivity to input fluctuations and predicts spikes with millisecond precision. **a-c)** The iGIF model response (red) is compared against data (gray). **a)** Comparison between steady-state $f - \mu_I$ curves observed in a typical Pyr neuron and produced by the iGIF model. Different colors and gray levels indicate the magnitude of input fluctuations σ_I (see legends in panel *a* and *b*). **b)** The average firing threshold observed in a typical cell and produced by the iGIF model. Model thresholds were mean-normalized. **c)** Summary data of results obtained in six Pyr neurons. Top: percentage change in steady-state firing rate obtained in response to the strongest depolarizing offset $\mu_I^{(8)}$ by increasing the input standard deviation σ_I (data are presented as in Fig. 4.1e). Bottom: average change in voltage threshold obtained by increasing the input standard

deviation σ_1 (data are presented as in Fig. 4.3b). **d-f**) As a control, the GIF model response (yellow) is compared against data (gray). Results are presented as in panels *a-c*. **g-j**) The predictive power of the iGIF (red) and the GIF model (yellow) was assessed on a new dataset (*test dataset*, black), which was not used for parameter extraction. **g**) Segment of the 20-second current used to build the *test dataset*. **h**) Spiking response of a Pyr neuron (black) to 9 repetitive injections of the current shown in panel *g*. The predictions of the iGIF and GIF model are shown in red and yellow, respectively. **i**) PSTHs computed by filtering the spike trains shown in panel *h* with a 500 ms rectangular window. The dashed line indicates 0 Hz. **j**) Typical intracellular response (black), as well as typical iGIF (red) and GIF (yellow) model prediction, to a single presentation of a 1-s segment of the current shown in panel *g*. **k-l**) Summary data showing the performance of the iGIF (red) and the GIF (yellow) model in predicting the *test dataset*. Filled bars and empty bars show the performance of models trained on the *f-I dataset* and the *training dataset*, respectively. Error bars represent one standard deviation across six neurons. **k**) Spike-timing prediction as quantified by the similarity measure M_d^* . The iGIF model significantly outperforms the GIF model with parameters extracted from the *f-I dataset* ($M_d^*=0.75$, s.d. 0.03, iGIF; $M_d^*=0.49$, s.d. 0.08, GIF; $n = 6$, paired Student *t*-test, $t_5 = -8.44$, $p = 3.8 \cdot 10^{-4}$) and from the *training dataset* ($M_d^*=0.83$, s.d. 0.02, iGIF; $M_d^*=0.76$, s.d. 0.05, GIF; $n = 6$, paired Student *t*-test, $t_5 = -3.25$, $p = 0.022$). **l**) The prediction error ϵ_{PSTH} on the PSTH (see panel *i*) was quantified by computing the root mean square error between data and model prediction. The iGIF model significantly outperforms the GIF model with parameters extracted from the *f-I dataset* ($\epsilon_{\text{PSTH}}=1.93$, s.d. 0.71 Hz, iGIF; $\epsilon_{\text{PSTH}}=3.22$, s.d. 0.91 Hz, GIF; $n = 6$, paired Student *t*-test, $t_5 = 6.33$, $p = 1.5 \cdot 10^{-3}$) but not when the parameters are extracted from the *training dataset* ($\epsilon_{\text{PSTH}}=1.41$, s.d. 0.28 Hz, iGIF; $\epsilon_{\text{PSTH}}=1.52$, s.d. 0.30 Hz, GIF; $n = 6$, paired Student *t*-test, $t_5 = 0.95$, $p = 0.38$). **m**) Comparison between GIF the iGIF model stochasticity. The iGIF model is significantly less stochastic than the GIF model ($\Delta V=0.59$, s.d. 0.12 mV, iGIF; $\Delta V=2.74$, s.d. 0.50 mV, GIF; $n = 6$, paired Student *t*-test, $t_5 = 9.68$, $p=2.0 \cdot 10^{-4}$, *f-I dataset*; $\Delta V=0.60$, s.d. 0.05 mV, iGIF; $\Delta V=1.28$, s.d. 0.29 mV, GIF; $n = 6$, paired Student *t*-test, $t_5 = 5.71$, $p=2.3 \cdot 10^{-3}$, *training dataset*).

Good single-neuron models predict the occurrence of individual spikes with millisecond precision (Gerstner and Naud, 2009). The iGIF model does not simply capture the average firing rate observed in Pyr neurons, but also reproduces the fine temporal structure of the spiking response. To take into account the fact that single neurons are stochastic (Mainen and Sejnowski, 1995) and to avoid problems related to overfitting, we assessed spike-timing prediction on a new experimental dataset (*test dataset*). This dataset was collected by recording the response to nine repetitive injections of a new fluctuating current $I_{\text{test}}(t)$ that was not used for parameter extraction. In order to test the model's ability of capturing the single-neuron response to different levels of input fluctuations, the standard deviation of the current $I_{\text{test}}(t)$ was modulated by a slow sinusoidal function (Fig. 4.6g, see Materials and Methods). On average, the iGIF model with parameters extracted from the dataset used to compute the $f - \mu_I$ curves (*f-I dataset*, see Fig. 4.1) was able to predict $75.1 \pm 3.2\%$ of the spikes with a precision of ± 4 ms (Fig. 4.6h,k). The iGIF model performed well also in predicting the slow fluctuations of the firing rate induced by the sinusoidal input modulation (Fig. 4.6i,l) as well as the rapid dynamics of the subthreshold membrane potential (Fig. 4.6j). As expected, the performance achieved by the GIF model were significantly lower. In particular, the GIF model predicted only $48.8 \pm 9.2\%$ of the spikes (Fig. 4.6h,k) and poorly captured the slow firing-rate fluctuations induced by the modulation of the input fluctuations (Fig. 4.6i,l).

In previous studies (Mensi et al., 2011; Pozzorini et al., 2013), we found that the GIF model was able to predict around 80% of the spikes observed in Pyr neurons responding to non-stationary inputs. At first glance, the low performance achieved here might therefore seem surprising. This result can however be understood by comparing the degree of stochasticity of the GIF model and the iGIF model (Fig. 4.6m). In both models, the parameter ΔV regulates the level of stochasticity of the spiking process (cf, Eq. 4.12). In particular, both models are deterministic if $\Delta V = 0$ and tend to an homogeneous Poisson process if $\Delta V \rightarrow \infty$. In the ideal case of a perfect model, ΔV is optimally tuned to capture trial-to-trial variability. In reality, a lack of flexibility in the model can bias the estimation of ΔV towards large values. Indeed, in an oversimplified model, all of the single-neuron features that can not be explained are interpreted as a manifestation of randomness. While the level of stochasticity observed in the iGIF model was weak ($\Delta V = 0.59$ mV, s.d. 0.13 mV), the values obtained by fitting the GIF model to the *f-I dataset* were always very high ($\Delta V = 2.74$ mV, s.d. 0.54 mV), explaining the low performance achieved in predicting individual spikes.

To make sure that the success of the iGIF model does not simply result from the aberrant level of stochasticity in the GIF model, we reassessed spike-timing prediction in both models by extracting model parameters from a third dataset (*training dataset*, see Materials and Methods) obtained by injecting a 120-s current having the same statistics as the *test dataset* (Fig. 4.7k-m, empty bars). As expected, the level of stochasticity in the GIF model dramatically decreased (Fig. 4.7m), recovering the GIF model's performance to expected levels ($M_d^* = 76.2\%$, s.d. 5.1%). Notably, the iGIF model with parameters extracted from the *training dataset* significantly outperformed the GIF model by predicting 82.8% of spikes ($M_d^* = 82.8\%$, s.d. 2.0%; Fig. 4.6k). Overall, these results demonstrate that the iGIF model is an excellent spiking neuron model capable of predicting individual spikes with millisecond precision and capturing the activity of Pyr neurons over a wide range of input statistics.

But the question remains: how do Pyr neurons adapt their coding strategy to maintain sensitivity to rapid input fluctuations, regardless of μ_I ? To answer this question, the next two sections present a detailed analysis of the iGIF model dynamics.

4.2.6 Enhanced sensitivity to input fluctuations results from a nonlinear interaction between spike-dependent and voltage-dependent threshold adaptation

The results obtained by fitting the iGIF model on intracellular recordings indicate that, in L5 Pyr neurons, the firing threshold depends on the emission of previous spikes as well as on the subthreshold voltage dynamics. While the timescale τ_θ of the voltage coupling is relatively short and is consistent with the hypothesis of fast Na^+ -channel inactivation, threshold changes induced by the emission of previous spikes are characterized by a much slower decay that possibly reflects slow Na^+ -channel inactivation (Fleidervish et al., 1996). Due to this separation of timescales, it is convenient to consider positive spike-triggered movements of the firing

threshold in the $V - \theta$ phase-plane as spike-triggered horizontal shifts of the spike initiation boundary (see arrows in Fig. 4.5e and Eq. 4.5). Since the coupling between membrane potential and firing threshold is nonlinear, this picture suggests that spike-dependent and voltage-dependent threshold adaptation might interact in a nontrivial way.

To understand this interaction, as well as its implications for single-neuron coding, we analyzed the dynamics of the iGIF model in response to three fluctuating currents with different offsets μ_I and a fixed standard deviation $\sigma_I = 100$ pA (Fig. 4.7). In response to a weak input $\mu_I = 90$ pA, the membrane potential fluctuated near the resting potential and action potentials were occasionally driven by large input fluctuations (Fig. 4.7a). In this particular regime, the evoked firing rate $f \approx 2$ Hz was very low and threshold movements induced by different action potentials did not build up significantly. Indeed, spike-dependent threshold adaptation $\phi(t) = \sum_{\hat{t}_j} \gamma(t - \hat{t}_j)$ mainly acted as a refractory process by reducing the firing probability after the emission of a previous spike. Given the modest average contribution $\bar{\phi} = T^{-1} \int_0^T \phi(t) dt$ of spike-dependent threshold adaptation (Fig. 4.7a, dashed gray), action potentials evoked in response to weak inputs were initiated with a low firing threshold (Fig. 4.7b). Membrane potential fluctuations were consequently confined to relatively low voltages, where the strength of the coupling between firing threshold and subthreshold membrane potential is weak. That is, where the coupling gain $G_\theta(V) = \frac{d}{dV} \theta_\infty^{\text{Na}}(V)$ is close to zero (Fig. 4.7b). As a result, in this regime, the dynamics of $\theta(t)$ primarily implemented an additional source of refractoriness by transiently increasing the firing threshold as a direct consequence of high-voltage after-spike reset (see Fig. 4.5e). We concluded that, in Pyr neurons responding to weak inputs, the coupling between membrane potential and firing threshold is only weakly recruited.

This result can alternatively be understood by analyzing the dynamics of the iGIF model in the $V - \theta$ plane, where spike-dependent changes of the firing threshold constantly move the spike-initiation boundary $V = \theta + \phi(t)$ (i.e., the region of the phase plane where $V = V_T$). Due to $\phi(t)$, this boundary is highly dynamic. Insights on the adaptive behavior of the iGIF model can however be gained by simply considering its average position $V = V_T + \bar{\phi}$. As shown in Figure 4.7a, for weak inputs evoking low firing rates, the average contribution of spike-dependent threshold adaptation is close to zero, meaning that, on average, the phase-plane region from where spikes originate is approximately delimited by the identity diagonal $V = \theta$ (Fig. 4.7c, blue). In all Pyr neurons included in this study we systematically found that $V_i > V_T^*$ (Fig. 4.5g). Consequently, at low firing rates, the voltage-dependent part of the coupling function $\theta_\infty^{\text{Na}}(V)$ was *masked* on the right-hand side of the diagonal $V = \theta + \bar{\phi}$, explaining why, in this regime, the coupling between firing threshold and membrane potential was only weakly recruited (Fig. 4.7c, see also Fig. 4.4).

Increasing the input strength to $\mu_I = 230$ pA (Fig. 4.7d-f) resulted in a mean firing rate of around 10 Hz and shifted the membrane potential distribution towards more depolarized potentials, where the threshold coupling strength G_θ becomes significant (Fig. 4.7e). As a result of increased spike-dependent threshold adaptation, the diagonal $V = \theta + \bar{\phi}$ was shifted towards more depolarized potentials, partially *unmasking* the voltage-dependent part of




Figure 4.7: Nonlinear interaction between spike-dependent and voltage-dependent threshold adaptation. Dynamics of the iGIF model (with parameters extracted from a typical Pyr neuron) responding to three fluctuating currents with $\sigma_I = 100$ pA, $\tau_I = 3$ ms and $\mu_I = 90$ pA (panels *a-c*), $\mu_I = 230$ pA (panels *d-f*) and $\mu_I = 450$ pA (panels *g-i*). **a)** Input current $I(t)$ (gray), membrane potential $V(t)$ (black), voltage-dependent threshold adaptation $\theta(t)$ (red), spike-dependent threshold adaptation $\phi(t) = \sum_{\hat{t}} \gamma(t - \hat{t})$ (blue) and average spike-dependent threshold adaptation $\bar{\phi} = T^{-1} \cdot \int_0^T \phi(s) ds$ (dashed gray, bottom) are shown as a function of time. The four dotted lines indicates (from top to bottom): $I = 0$ nA, $V = E_0$, $\theta = V_T^*$ and $\phi = 0$ mV. **b)** Distribution of subthreshold membrane potential fluctuations $P(V)$ (gray) and of voltages at which spikes were initiated $P(V|\text{spike})$ (black). The gain of the threshold coupling $G_\theta(V) = \frac{d}{dV} \theta_\infty^{\text{Na}}(V)$ is shown in red. **c)** Phase plane $\theta - V$. In response to a weak input μ_I , the threshold coupling $\theta_\infty^{\text{Na}}(V)$ (red) is *masked* on the right-hand side of the average spiking boundary $V = \theta + \bar{\phi}$ (blue line). For comparison, the diagonal $V = \theta$ is also shown (dashed black). The gray distribution (copied from panel *c*) shows $P(V)$. Inset: The theoretical filter $\kappa_{\text{eff}}(t)$ of somatic integration (black, Eq. 4.8) is compared against the passive membrane filter $\kappa_0(t)$ (dotted black) and the GLM filter $\kappa_{\text{GLM}}(t)$ (gray) extracted as a control by fitting the spiking response of the iGIF model. **d-f)** iGIF model dynamics in response to a fluctuating input current with mean $\mu_I = 230$ pA (results are presented as in panels *a-c*). In this regime, the nonlinear coupling $\theta(t)$ is partially recruited. **g-i)** iGIF model dynamics in response to a fluctuating input current with mean $\mu_I = 450$ pA (results are presented as in panels *a-c*). In this regime, the nonlinear coupling $\theta(t)$ is strongly recruited.

the nonlinear coupling $\theta_{\infty}^{\text{Na}}(V)$ (Fig. 4.7f). Consequently, in this regime, the coupling $\theta(t)$ between firing threshold and membrane potential started to play a role that goes beyond the refractoriness induced by the after-spike reset (Fig. 4.7d, red). Further increasing the input strength to $\mu_I = 450$ pA (Fig. 4.7g-i) made the iGIF model fire at around 18 Hz, where L5 Pyr neurons start losing sensitivity to the mean drive μ_I and manifest enhanced sensitivity to input fluctuations (see Fig. 4.1). Notably, in this regime, threshold movements triggered by different spikes accumulated in such a way as to further *unmask* the coupling function $\theta_{\infty}^{\text{Na}}(V)$ (Fig. 4.7i). The strong input μ_I pushed the membrane potential fluctuations to very depolarized values, where the threshold coupling reaches its maximal strength $G_{\theta} \approx \theta_{\text{slope}}^{\text{Na}}$ (Fig. 4.7h). Consequently, following the after-spike reset $\theta \rightarrow V_T^*$, the coupling variable $\theta(t)$ rapidly increased to large values where, in contrast to the transient behavior observed in response to weak inputs, it fluctuated until the next spike was emitted (compare the red traces in Fig. 4.8g and Fig. 4.8a,d).

Overall, the results reported in Figure 4.7 provide evidence for the existence of a non-trivial interplay between spike-dependent and voltage-dependent threshold movements. In particular, we found that the increased contribution of spike-dependent mechanisms induced by large firing rates, progressively *unmasked* (i.e., activated) the coupling between membrane potential and firing threshold (Fig. 4.7c,f,i), thereby enhancing single-neuron sensitivity to rapid signals. The functional implications of this interaction are further investigated in the next sections.

4.2.7 The iGIF model captures and explains the complex forms of adaptation revealed by the GLM-based analysis

In order to understand the mechanisms underlying the complex forms of adaptation observed in Figure 4.2 by fitting a GLM to the spiking activity of L5 Pyr neurons, we performed a systematic reduction of the iGIF model (see Materials and Methods). More precisely, we derived an analytical formula describing the GLM filters $\kappa_{\text{GLM}}(t)$ and $h_{\text{GLM}}(t)$ that best approximate the iGIF model response to fluctuating currents around different offsets μ_I . Since GLM parameter extraction only relies on spiking data, the linear filter $\kappa_{\text{GLM}}(t)$ is expected to combine the effect of the passive properties of the membrane, of the conductance change triggered by previous spikes and of the coupling between firing threshold and membrane potential. For the same reason, the GLM spike-history filter $h_{\text{GLM}}(t)$ is expected to model the combined action of all the biophysical processes that mediate spike-dependent adaptation.

In the case of a standard GIF model (Mensi et al., 2011; Pozzorini et al., 2013), where the coupling between membrane potential and firing threshold is absent, the linear filter $\kappa_{\text{GLM}}(t)$ of a GLM that best captures the iGIF model response corresponds to the membrane filter $\kappa_m(t)$ obtained by combining the passive properties of the membrane and the conductance

changes triggered by previous spikes (Richardson, 2004):

$$\kappa_m(t) = \Theta(t) \frac{1}{C} \exp\left(-\frac{t}{\tau_m}\right), \quad (4.6)$$

where $\Theta(t)$ is the Heaviside step function, C is the membrane capacitance and $\tau_m = C/(g_0 + \bar{g}_\eta)$ is the membrane timescale computed by taking into account both the average contribution of the spike-triggered conductance $\bar{g}_\eta = T^{-1} \int_0^T \sum \hat{i}_j \eta(t - \hat{t}_j) dt$ and the passive leak g_0 (see Materials and Methods).

In the presence of a nonlinear coupling $\theta_\infty(V)$ between membrane potential and firing threshold, the mapping between iGIF model and GLM is more difficult and additionally involves the linearization of Equation 4.2 (see Materials and Methods). We simplified the threshold dynamics of the iGIF model by taking the first-order approximation $\theta_\infty^{\text{Na}}(V) \approx \bar{C}_\theta + \bar{G}_\theta \cdot V$, with \bar{C}_θ being a constant and \bar{G}_θ being the average coupling strength computed with respect to the membrane potential distribution $P(V)$:

$$\bar{G}_\theta = \int_{-\infty}^{\infty} G_\theta(V) P(V) dV. \quad (4.7)$$

Consequently, since in the iGIF model the spiking probability only depends on the difference between membrane potential and firing threshold (see Fig. 4.5a and Eq. 4.12), the GLM linear filter $\kappa_{\text{GLM}}(t)$ that best approximates the response of an iGIF model is given by the effective linear filter $\kappa_{\text{eff}}(t)$ defined as (Platkiewicz and Brette, 2011):

$$\kappa_{\text{eff}}(t) = \kappa_m(t) - \bar{G}_\theta \cdot \int_{-\infty}^{\infty} \kappa_\theta(s) \kappa_m(t-s) ds, \quad (4.8)$$

where $\kappa_\theta(t) = \Theta(t) \frac{1}{\tau_\theta} \exp\left(-\frac{t}{\tau_\theta}\right)$ is the threshold-coupling filter and \bar{G}_θ depends on the input statistics via the membrane potential distribution $P(V)$ (see Eq. 4.7 and Materials and Methods).

Using the parameters of the iGIF model extracted from intracellular recordings, we computed the theoretical filters $\kappa_m(t)$ and $\kappa_{\text{eff}}(t)$ for a set of fluctuating currents characterized by different average intensities μ_I (Fig. 4.8). In agreement with our experimental findings, increasing the DC component of the input progressively augmented the membrane conductance $g_L = g_0 + \bar{g}_\eta$, thereby shortening the membrane filter $\kappa_m(t)$ (Fig. 4.8a). We concluded that part of the reduction observed in the timescale of somatic integration resulted from a spike-dependent increase of the membrane conductance (Fig. 4.8a, inset). In response to a low input μ_I , the membrane potential fluctuations of the iGIF model were confined in a region where the gain $G_\theta(V)$ of the threshold coupling is vanishing, such that $\bar{G}_\theta \approx 0$ and, according to Equation 4.8, $\kappa_{\text{eff}}(t) \approx \kappa_m(t)$. This result indicates that in response to weak inputs, the coupling between membrane potential and firing threshold is not recruited and somatic integration is entirely controlled by the membrane filter $\kappa_m(t)$. Explaining the adaptive changes observed in the linear filters $\kappa_{\text{GLM}}(t)$ extracted from the data, the iGIF-Na model correctly predicts

that, in response to increasing depolarizing offsets μ_1 , the coupling strength \bar{G}_θ progressively augments, thereby shunting the membrane filter and further shortening the temporal window of somatic integration $\kappa_{\text{eff}}(t)$ (Fig. 4.8b). These results, whose accuracy was confirmed by fitting GLMs to artificial data generated with the iGIF model (Fig. 4.7c,f,i, insets), suggest that both a spike-dependent conductance and an intricate dynamics of the firing threshold actively control somatic integration in order to allow for coincidence detection, regardless of the statistical properties of the input. Notably, the predictions of our theory were not only qualitative. By measuring the timescales of the theoretical filters $\kappa_m(t)$ and $\kappa_{\text{eff}}(t)$, we were indeed able to accurately predict the reduction observed in both the membrane timescale τ_m and the timescale τ_{GLM} of somatic integration (Fig. 4.8c).

Results obtained by fitting GLMs to artificial data generated by stimulating the iGIF model with fluctuating currents of different means μ_1 demonstrate that the iGIF model also captured the adaptive changes that we experimentally observed in the spike-history filter h_{GLM} (Fig. 4.8d). Finally, analytical results obtained following a derivation similar to the one used for Eq. 4.8 (see Materials and Methods) indicate that changes in $h_{\text{GLM}}(t)$ are induced by a shunting effect of the nonlinear threshold coupling on the adaptation current I_A (Fig. 4.8d).

4.2.8 L5 Pyr neurons feature two distinct forms of adaptation

In order to study the temporal dynamics of single-neuron adaptation, we finally performed a switching experiment in which the iGIF model was stimulated with a fluctuating current, whose mean μ_1 periodically switched between a low and a high value, with cycle period $T_{\text{cycle}} = 10$ s (Fig. 4.8e). In response to a sudden increase in μ_1 , the output firing rate transiently raised and then decayed over multiple timescales, confirming that in the iGIF model the combined action of the spike-triggered conductance and the spike-triggered movement of the firing threshold mediate spike-frequency adaptation. Similarly, in response to a sudden decrease of μ_1 , the output firing rate initially dropped and then recovered.

By computing $\kappa_{\text{eff}}(t)$ at different moments in time relative to the cycle, we also found that, in contrast to spike-frequency adaptation, adaptive changes in the timescale τ_{eff} of somatic integration were almost instantaneous (Fig. 4.8e, red and Fig. 4.8f). This result can be understood by noting that, in Eqs. 4.7-4.8, the strength \bar{G}_θ of the shunting effect induced by the threshold dynamics on the membrane filter is controlled by the voltage distribution, which, in response to a sudden switch in the input statistics, changes rapidly. The results presented in Figure 4.8e,f are reminiscent of the adaptive behavior previously observed both in retinal ganglion cells (Baccus and Meister, 2002) and motion sensitive neurons (Fairhall et al., 2001b) responding *in vivo* to external stimuli and demonstrate how intrinsic cellular mechanisms can support complex forms of adaptation that in the past have been attributed to network effects.




Figure 4.8: The iGIF model captures and explains complex forms of adaptation. **a)** Average membrane filters $\kappa_m(t)$ (Eq. 4.6) computed with iGIF model parameters extracted from 6 Pyr neurons by increasing μ_I from 0.05 nA (blue) to 0.5 nA (red, see colorbars in panel *d*). The passive membrane filter $\kappa_0(t)$ (dashed black) is shown for comparison. Inset: average effective conductance $g_L = g_0 + \bar{g}_\eta$ as a function of μ_I . In the absence of spikes, the effective conductance equals g_0 (dashed black). The gray area indicates one standard deviation across neurons. **b)** Same results as in panel *a*, but for the average effective filter $\kappa_{\text{eff}}(t)$ (Eq. 4.8). Inset: average coupling strength \bar{G}_θ as a function of the mean input μ_I . Conventions are as in panel *a*. **c)** The timescales τ_m (red) and τ_{eff} (blue) predicted by the iGIF model with parameters extracted from six neurons match the experimental data (black). Colored lines and gray areas indicate the mean and one standard deviation across neurons. Experimental data are copied from Figure 4.2d. The predicted timescales τ_{eff} (red) were obtained by fitting single-exponential functions to $\kappa_{\text{eff}}(t)$. **d)** The iGIF model explains the adaptive changes in $h_{\text{GLM}}(t)$ (see Fig. 4.2c). Left: average spike-history filter $h_{\text{GLM}}(t)$ obtained by fitting a GLM to artificial data generated by simulating the iGIF model response to fluctuating currents of increasing μ_I (see colorbar). Right: Average the-

oretical filters $h_{\text{eff}}(t)$ computed according to Eq. 4.30 using iGIF model parameters extracted from 6 Pyr neurons. Errors are due to the fact that our theory neglects the correlations between voltage and conductance (see Materials and Methods). **e-f**) Switching experiment performed in a iGIF model with parameters extracted from a typical cell to study the temporal evolution of single-neuron adaptation induced by a sudden change in μ_1 . **e**) Top: fluctuating current (gray) generated by periodically switching μ_1 (dark gray) between 0.1 nA and 0.27 nA, with cycle period $T_{\text{cycle}} = 10$ s (only one cycle is shown). Middle: effective timescale of integration τ_{eff} as a function of time. Bottom: output firing rate. While spike-frequency adaptation occurs on both fast and slow timescales, changes in τ_{eff} triggered by a switch in μ_1 are almost instantaneous. Horizontal black lines indicate: 0 nA, 0 ms and 0 Hz. **f**) Comparison between effective linear filters $\kappa_{\text{eff}}(t)$ estimated at different moments in time during the switching experiment (see arrows in panel *e*). The filters estimated at steady-state (late low, late high) closely resemble the ones estimated right after the stimulus switch (early low, early high), indicating that adaptive changes in $\kappa_{\text{eff}}(t)$ are almost instantaneous. The passive membrane filter $\kappa_0(t)$ (dashed black) is shown for comparison. In all panels, input currents were generated according to Eq. 4.9 with $\sigma_1 = 100$ pA and $\tau_1 = 3$ ms.

Overall, our results indicate that L5 Pyr neurons respond to a sudden change in the input statistics by adapting both their output firing rate and the temporal window over which incoming spikes are considered to be coincident. The high speed at which the timescale of somatic integration adapts indicates that, regardless of the input statistics, L5 Pyr neurons can operate as coincidence detectors, thereby supporting a temporal code in which the precise timing of action potentials conveys relevant information.

4.3 Discussion

Despite years of research, it remains unclear whether cortical neurons operate as temporal integrators or coincidence detectors. The answer to this question is of crucial importance because of its profound implications for neural coding (Ratté et al., 2013; König et al., 1996). By invoking constraints imposed by the passive membrane time constant, previous studies questioned the possibility of a temporal code in which the fine structure of spike trains conveys relevant information (Shadlen and Newsome, 1994). Combining *in vitro* electrophysiological recordings with modeling, we found that, in L5 Pyr neurons of the mouse SSC, the effective timescale over which afferent spikes are somatically integrated is not entirely controlled by the passive properties of the membrane, but adapts to the input statistics as a result of the firing threshold dynamics. In particular, we found that by increasing the DC component of a fluctuating input current, the firing threshold becomes progressively coupled to the subthreshold membrane potential dynamics in such a way as to shorten the temporal window determining whether two afferent spikes are coincident. Suggesting that the cellular mechanisms discussed here are also at work in biologically relevant situations, previous *in vivo* studies reported that the dynamics of the firing threshold enhances coincidence detection in cat visual cortex (Azouz and Gray, 2000, 2003) and sharpens feature selectivity in rat barrel cortex (Wilent and Contreras, 2005).

By extending our previous GIF model (Mensi et al., 2011; Pozzorini et al., 2013) with a nonlinear coupling between membrane potential and firing threshold (Platkiewicz and Brette, 2011), we were able to explain why the output firing rate of L5 Pyr neurons remains sensitive to rapid input fluctuations over a broad range of depolarizing offsets μ_1 . During the last years, enhanced sensitivity to rapid input fluctuations has also been observed in Pyr neurons of the rat PFC (Arsiero et al., 2007) and hippocampus (Fernandez et al., 2011), indicating that our results might equally apply to neurons of other species and cortical areas.

The iGIF model does not simply capture the average firing rate observed in Pyr neurons, but also predicts the occurrence of individual spikes with millisecond precision. Moreover, in agreement with our experimental measurements, the firing threshold of our iGIF model was positively correlated with the DC component of the input and was reduced by rapid input fluctuations. Finally, confirming the validity of our main findings, a semi-analytical reduction of the iGIF model allowed us to accurately predict and explain experimental differences between the membrane timescale and the effective timescale of somatic integration.

4.3.1 Biophysical implementation of nonlinear threshold dynamics

The iGIF model provides a phenomenological description of single neuron dynamics. Consequently, our results did not allow us to identify with certainty the biophysical mechanisms underlying nonlinear threshold dynamics. Previous results indicate that, in rat Pyr neurons, a voltage dependence of the firing threshold is mediated by Kv1 channels (Higgs and Spain, 2011). However, these channels seem not to be expressed by L5 Pyr neurons of the mouse SSC (Goldberg et al., 2008; Miller et al., 2008).

In a recent theoretical study based on the standard HH model, Platkiewicz and Brette (2011) proposed that a nonlinear coupling between subthreshold membrane potential and firing threshold could result from fast Na^+ -channel inactivation. In agreement with this hypothesis, we found that the threshold-coupling $\theta_\infty(V)$ was correctly described by a smooth rectifier function. This result is noteworthy because here, in contrast to previous studies (Higgs and Spain, 2011; Fontaine et al., 2014b,a), the functional shape of $\theta_\infty(V)$ was not assumed *a priori*, but was extracted from intracellular recordings using a nonparametric approach. Also, the values estimated for the coupling timescale τ_θ were consistent with the hypothesis of fast Na^+ -channel inactivation. Providing further indirect evidence, the threshold sharpness $k_a = 2.2 \pm 1.1$ mV extracted by fitting the data under the assumption of Na^+ -channel inactivation was very close to previous measurements that were independently obtained by fitting an exponential integrate-and-fire model to intracellular recordings (Badel et al., 2008).

In our iGIF model, the firing threshold dynamics also depends on the spike-history. In agreement with previous results (Pozzorini et al., 2013), we indeed found that each action potential triggers a change in the firing threshold that lasts for several seconds and decays according to a power-law over several orders of magnitude. Overall, our phenomenological model is consistent with a biophysical model of Na^+ -channels in which inactivation is independently

controlled by fast and slow gating variables (Fleidervish et al. (1996), see also the discussion in Platkiewicz and Brette (2011)). While fast inactivation enhances sensitivity to rapid inputs by *absorbing* slow membrane potential fluctuations, slow inactivation acts as an homeostatic mechanisms by increasing the average firing threshold in response to strong inputs. Confirming the intuition of Platkiewicz and Brette (2011), our results demonstrate that slow and fast Na^+ -channel inactivation interact in a nontrivial way to dynamically adapt the effective timescale of somatic integration to the input statistics. Indeed, by increasing the average threshold, slow inactivation allows for fast inactivation to occur at subthreshold potentials. A similar interaction has been previously shown to occur in a Morris-Lecar model (Morris and Lecar, 1981) stimulated with a conductance based input (Prescott et al., 2006). In contrast to what we found, Prescott et al. (2006) proposed that reductions in the effective timescale of somatic integration are mediated by the activation of M-currents at subthreshold voltages, which become possible in situations where the voltage threshold for spike initiation is increased due to shunting effects induced by presynaptic spikes.

4.3.2 Are simplified spiking models getting complicated?

During the last decades, a number of simplified spiking models, including our previous GIF model (Mensi et al., 2011; Pozzorini et al., 2013), have been shown to accurately predict the spiking response evoked *in vitro* by stationary (or quasi-stationary) currents (Gerstner and Naud, 2009). Cortical neurons feature a strong nonlinear behavior. Consequently, when tested on experimental data acquired in response to currents with varying statistics, the performance of simplified threshold models typically drops. Overall, designing and fitting a spiking model capable of predicting the electrical activity of cortical neurons operating in different regimes remains a big challenge. Indeed, increasing the complexity of a spiking neuron model rapidly makes parameter estimation a difficult problem.

Here, we introduced a new spiking model called inactivating Generalized Integrate-and-Fire (iGIF) and obtained by extending the standard leaky integrate-and-fire model with: the escape-rate model for stochastic spike generation, a spike-triggered conductance and spike-triggered movement of the firing threshold for spike-frequency adaptation and a nonlinear coupling between membrane potential and firing threshold for enhanced sensitivity to input fluctuations. Despite its relative simplicity, the model successfully captures the spiking activity of L5 Pyr neurons over a broad range of input statistics and outperforms our previous GIF model in predicting the occurrence of individual spikes with millisecond precision. Despite its relative complexity, model parameters can be robustly extracted with a new two-step procedure, which extends our previous method (Mensi et al., 2011; Pozzorini et al., 2013). In the first step, the parameters governing the subthreshold dynamics of the membrane potential are extracted by minimizing the sum of squared errors on the rate of change (i.e., the first order derivative) of the voltage fluctuations. Except for the parameter E_R , which defines the reversal potential associated with the spike-triggered conductance, all the parameters can be extracted with a simple multilinear regression. In the second step, the parameters governing

the dynamics of the firing threshold are obtained with a maximum likelihood approach similar to the one used to fit GLMs to spiking data (Truccolo et al., 2005; Pillow et al., 2008). Given the parameter τ_θ controlling the threshold-coupling timescale, the likelihood of the model is convex in its parameters, which can therefore be extracted using standard gradient ascent methods.

The iGIF model remains amenable to analytical treatment and can be reduced to a simpler model, thereby providing insights on the functional role of its multiple adaptation mechanisms. Finally, in contrast to complex models obtained by extending the standard Hodgkin-Huxley model (Hodgkin and Huxley, 1952) with additional conductances, the iGIF model accounts for different aspect of the neural dynamics by the means of phenomenological mechanisms that, while being related to known biophysical processes, can be understood from a functional perspective.

4.3.3 Connection to sensory adaptation

In response to a sudden change in the input statistics, both retinal ganglion cells and motion-sensitive neurons in the fly feature two forms of adaptation (Fairhall et al., 2001b; Baccus and Meister, 2002). Right after a stimulus change, these neurons rapidly modify the shape of their receptive field, thereby adapting the stimulus feature to which they are responsive. While this mechanism is very fast, the same neurons also feature a slower form of adaptation that manifests itself in a decay of the output firing rate over multiple timescales. This second mechanism, known as spike-frequency adaptation, does not induce further changes in the receptive field, but simply reduces the overall excitability of the neuron. Since both the timescale and the net effect of these two adaptation processes are different, it has been hypothesized that changes in feature selectivity and output firing rate are controlled by two independent mechanisms. While in retinal ganglion cells these forms of sensory adaptation have been shown to emerge from network effects, the origin of sensory adaptation in motion sensitive neurons of the fly remains unclear. Our results indicate that both forms of adaptation can be supported by intrinsic cellular mechanisms.

4.3.4 Conclusion

Intracellular patch-clamp recordings are widely used both *in vitro* and *in vivo* to study how neurons process information. As a result of an intricate dynamics of the firing threshold, we found that the subthreshold and the spiking response of single neurons are related in a nontrivial manner. Consequently, drawing conclusions about the neural code directly from subthreshold voltage recordings requires care.

In particular, our results demonstrate and explain why measurements of the membrane timescale provide a biased estimate of the timescale over which synaptic inputs are temporally integrated. Indeed, a nonlinear coupling between membrane potential and firing threshold

adaptively shortens the effective timescale of somatic integration (without affecting the membrane timescale) in such a way as to enhance coincidence detection over a broad range of input statistics.

4.4 Materials and Methods

4.4.1 Electrophysiological recordings

All procedures in this study were conducted in conformity with the Swiss Welfare Act and the Swiss National Institutional Guidelines on Animal Experimentation for the ethical use of animals. The Swiss Cantonal Veterinary Office approved the project following an ethical review by the State Committee for Animal Experimentation.

Somatic whole-cell *in vitro* current clamp recordings were performed on 300 μm thick parasagittal acute slices from the right hemispheres of male P13-P15 C57Bl/6J wild-type mice. Brains were quickly dissected and sliced (HR2 vibratome, Sigmann Elektronik, Germany) in ice-cold artificial cerebrospinal fluid (ACSF) (in mM: NaCl 124.0, KCl 2.50, MgCl_2 10.0, NaH_2PO_4 1.25, CaCl_2 0.50, D-(+)-Glucose 25.00, NaHCO_3 25.00; pH 7.3 ± 0.1 , aerated with 95% O_2 / 5% CO_2), followed by a 15 minute incubation at 34 $^\circ\text{C}$ in standard ACSF (in mM: NaCl 124.0, KCl 2.50, MgCl_2 1.00, NaH_2PO_4 1.25, CaCl_2 2.00, D-(+)-Glucose 25.00, NaHCO_3 25.00; pH 7.40, aerated with 95% O_2 / 5% CO_2). To ensure intact axonal and dendritic arborisation, electrophysiological recordings were conducted in slices cut parallel to the apical dendrites. Recordings in Layer 5 of the primary somatosensory cortex were performed at 32 ± 1 $^\circ\text{C}$ in standard ACSF with an Axon Multiclamp 700B Amplifier (Molecular Devices, USA) using 5 - 7 $\text{M}\Omega$ borosilicate pipettes, containing (in mM): K^+ -gluconate 110.00, KCl 10.00, ATP- Mg^{2+} 4.00, Na^{2+} -phosphocreatine 10.00, GTP- Na^+ 0.30, HEPES 10.00, biocytin 5.00 mg/ml; pH 7.30, 300 mOsm. Cells were visualized using infrared differential interference contrast video microscopy (VX55 camera, Till Photonics, Germany and BX51WI microscope, Olympus, Japan).

Data were acquired with sampling frequency $\Delta T^{-1} = 10$ kHz using an ITC-18 digitizing board (InstruTECH, USA) controlled by a custom-written software module operating within IGOR Pro (Wavemetrics, USA). Voltage signals were low-pass filtered (Bessel, 10 kHz) and not corrected for the liquid junction potential. Only cells with an access resistance ≤ 20 $\text{M}\Omega$ (17.7 ± 2.3 $\text{M}\Omega$, $n = 6$) were retained for further analysis.

4.4.2 Current injections

In all the experiments included in this study, neurons were stimulated with *in vivo*-like fluctuating currents $I(t)$ generated according to an Ornstein-Uhlenbeck process:

$$\tau_1 \dot{I}(t) = -I(t) + \mu_I + \sqrt{2\tau_1\sigma_I} \cdot \psi(t), \quad (4.9)$$

where $\psi(t)$ is a Gaussian white-noise process with zero mean and unitary variance, τ_I is the correlation timescale, μ_I is the mean current and σ_I defines the magnitude of the fluctuations (that is, the standard deviation of the current). The temporal correlation of the input was fixed to $\tau_I = 3$ ms and input currents $I(t)$ were generated at a sampling rate $\Delta T^{-1} = 10$ kHz.

To measure the impact of input fluctuations on the single-neuron input-output transfer function (i.e., the f - μ_I curve), we somatically injected a set of 5-second currents with different means μ_I and standard deviations σ_I (see Eq. 4.9). To let the cell recover, injections were performed with interstimuli intervals of 25 seconds. Similar protocols have already been applied in previous studies (Rauch et al., 2003; Higgs et al., 2006; Arsiero et al., 2007). Here, to exhaustively explore the parameter space (μ_I, σ_I) and to accurately estimate the experimental f - μ_I curves, we considered four different standard deviations $\sigma_I \in \{0, 50, 100, 150\}$ pA and eight different means $\mu_I \in [0, \mu_{\max}]$ nA, with μ_{\max} begin cell-dependent. Each neuron was stimulated with 32 different inputs that were presented randomly. The entire protocol was repeated 3 times. When stimulated with strong inputs, pyramidal neurons undergo spike failure and can not sustain repetitive firing for long periods of time (see, e.g., Fleidervish et al. (1996)). At the beginning of each experiment, the maximum current μ_{\max} was defined in such a way as to reach saturation of the steady state firing rate while preventing spike failures. For that, neurons were tested with 6-s-long noiseless currents (i.e., $\sigma_I = 0$) of increasing magnitude μ_I . Cells that could not sustain continuous firing for input currents $\mu_I < 0.4$ nA were discarded. The maximal mean input μ_{\max} was comprised between 0.4 and 0.55 nA.

To evaluate model performance in predicting the occurrence of individual spikes, a different set of experiments was performed. Currents were generated according to Equation 4.9, but in this case, the stochastic process used to generate the input was made non-stationary by modulating the standard deviation σ_I with a sinusoidal function

$$\sigma_I(t) = \sigma_0 \left(1 + \frac{1}{2} \cdot \sin \left(\frac{2\pi}{T} \cdot t \right) \right), \quad (4.10)$$

where $T = 5$ s is the modulation period. For each cell, input parameters were calibrated to obtain an average firing rate of 10 Hz oscillating between 7 and 13 Hz, approximatively. After calibration, input parameters were in the following ranges: $\mu_I \in [120, 190]$ pA, $\sigma_0 \in [120, 190]$ pA. Since the spiking responses of both neurons and GIF models are stochastic, spike-timing prediction was quantified on a *test set* obtained by 9 repetitive injections of the same (i.e. *frozen-noise*) 20-s current generated according to Equations 4.9-4.10. For parameter extraction, a *training set* was used in which single neurons were stimulated with a single 120-s current having the same statistics as the *test set*, but in which a different realization of the white-noise process $\psi(t)$ was used. All the injections were performed with inter-stimuli intervals of 25 seconds.

4.4.3 Data preprocessing

When acquired with the same electrode used to inject the external input $I(t)$, current-clamp recordings $V_{\text{rec}}(t)$ are biased versions of the membrane potential $V_{\text{data}}(t)$ (Badel et al., 2008). This bias can in principle be removed using series resistance or bridge balance compensation. However, perfect calibration of these methods is technically difficult to achieve. Moreover, during long experiments, the electrode properties, and in particular the series resistance R_e , are subject to change (Pozzorini et al., 2013). Quantitative comparison between membrane potentials evoked by input currents having different offsets μ_1 requires accurate electrode compensation. Indeed, a non-neutralized series resistance \tilde{R}_e would lead, on average, to a mean input-dependent bias $V_{\text{bias}}(\mu_1) = \tilde{R}_e \mu_1$ (Pozzorini et al., 2013). To avoid this and others problems, for all the *in vitro* recordings included in this study, online series resistance compensation was complemented by offline Active Electrode Compensation (AEC) (Brette et al., 2008; Badel et al., 2008). For that, the same procedure applied in Pozzorini et al. (2013) was used. In case of long experiments, estimating the electrode properties at different moments in time can improve the quality of the data by removing drifts due to slow changes in the electrode properties (Pozzorini et al., 2013). For this reason, electrode filters used for AEC were extracted from 10-s subthreshold injections performed before the *training set*, before the *test set* and every sixteen injections in the protocol used to measure the f - μ_1 curves. Subthreshold input currents were generated according to Equation 1 with $\mu_1 = 0$ nA, $\sigma_1 = 75$ pA and $\tau_1 = 3$ ms.

4.4.4 Extracting voltage threshold for spike initiation from *in vitro* recordings

Bifurcation analysis of neuron models capturing smooth spike initiation demonstrates that the concept of voltage threshold (that is, the largest membrane potential that a neuron can reach without emitting a spike) does not have a univocal definition. However, different possible definitions produce very similar results that mainly differ by a shift (Platkiewicz and Brette, 2010). For this reason, our analysis is based on relative variations between the voltage threshold in different conditions, rather than on absolute values. In practice, given an intracellular recording, different methods exist to estimate the voltage threshold (Sekerli et al., 2004). Here, for each spike in the dataset, the voltage threshold was estimated by measuring the membrane potential at which the depolarization rate dV/dt became larger than 10 mV/ms. As shown in Fig. 4.3e,g, voltage threshold detection is only weakly affected by the precise choice of this parameter.

In Figure 4.3c, for each input condition, the depolarization rate preceding the emission of an action potential was extracted from the average spike shape (i.e., from the spike-triggered average of the membrane potential) by performing a linear regression on the time interval $[\hat{t} - 3.5 \text{ ms}, \hat{t} - 0.5 \text{ ms}]$, with \hat{t} denoting the time at which action potentials were initiated. In Figure 4.3d, the average subthreshold membrane potential was computed by discarding all the data points $\{t | t \in [\hat{t}_j - 2 \text{ ms}, \hat{t}_j + 10 \text{ ms}]\}$ that were too close to action potentials $\{\hat{t}_j\}$.

4.4.5 Generalized Linear Model (GLM)

In the GLM (Truccolo et al., 2005; Pillow et al., 2008), spikes are generated stochastically with firing intensity $\lambda_{\text{GLM}}(t)$ defined as:

$$\lambda_{\text{GLM}}(t) = \lambda_0 \cdot \exp \left(E_0 + \int_0^\infty \kappa_{\text{GLM}}(s) I(t-s) ds + \sum_{\hat{t}_j < t} h_{\text{GLM}}(t - \hat{t}_j) \right), \quad (4.11)$$

where $\lambda_0 = 1$ Hz, E_0 is a constant, $\kappa_{\text{GLM}}(t)$ is an arbitrarily-shaped filter through which the input is integrated and $h_{\text{GLM}}(t)$ accounts for all spike-triggered processes that make the single-neuron activity history-dependent. GLM parameter extraction is performed using the standard maximum likelihood method described in Truccolo et al. (2005) and Pillow et al. (2008). For that, both $\kappa_{\text{GLM}}(t)$ and $h_{\text{GLM}}(t)$ were expanded in linear combinations of rectangular basis functions.

4.4.6 Inactivating Generalized Integrate-and-Fire model (iGIF)

The GIF model introduced in Mensi et al. (2011) and Pozzorini et al. (2013) was augmented with a nonlinear coupling between the subthreshold membrane potential $V(t)$ and the firing threshold $V_T(t)$. Such a coupling has recently been shown to occur in Hodgkin-Huxley models in which Na^+ -channels start to inactivate at subthreshold voltages (Platkiewicz and Brette, 2011). For this reason, we call our model iGIF, where *i* stands for *inactivating*. A list of model variables and parameters is provided in Table 4.1.

In the model, spikes are produced stochastically according to the conditional firing intensity $\lambda(t)$ defined by the exponential escape-rate model (Gerstner and van Hemmen, 1992; Jolivet et al., 2006a):

$$\lambda(t) = \lambda_0 \exp \left(\frac{V(t) - V_T(t)}{\Delta V} \right), \quad (4.12)$$

where $\lambda_0 = 1$ Hz and ΔV defines the level of stochasticity. In the limit $\Delta V \rightarrow 0$, the model becomes deterministic and action potentials are fired when the firing threshold is reached. The dynamics of the subthreshold membrane potential is modeled as a leaky integrator augmented with a spike-triggered conductance $\eta(t)$ that describes the time course of the conductance change after a spike. Mathematically, the membrane potential evolves according to the following differential equation:

$$C\dot{V} = -g_0(V - E_0) + I - \sum_{\hat{t}_j < t} \eta(t - \hat{t}_j) \cdot (V - E_R), \quad (4.13)$$

where C , g_0 and E_0 describe the passive properties of the membrane, $\tau_0 = C/g_0$ is the passive membrane timescale, $\{\hat{t}_1, \hat{t}_2, \hat{t}_3, \dots\}$ are the spike times, E_R is a reversal potential and I is the external input. Conductance changes triggered by different spikes accumulate and produce

spike-frequency adaptation (or facilitation). The functional shape of $\eta(t)$ is not assumed *a priori*, but is extracted from experimental data (see below). After each spike, the membrane potential is reset to V_{reset} and the numerical integration only restarts after an absolute refractory period T_{ref} .

The dynamic threshold V_T is modeled as follows

$$V_T(t) = \theta(t) + \sum_{\hat{t}_j < t} \gamma(t - \hat{t}_j), \quad (4.14)$$

where $\gamma(t)$ is a function of time describing the movement of the firing threshold after the emission of an action potential. Since the latter term can only account for spike-dependent effects, the model is augmented with an additional state variable $\theta(t)$ implementing a coupling between the dynamics of the firing threshold and that of the subthreshold membrane potential. Based on theoretical results obtained by a systematic reduction of the Hodgkin-Huxley model, it has recently been proposed that this coupling might be nonlinear and could take different forms depending on the underlying biophysical mechanism (Platkiewicz and Brette, 2010). In the iGIF model, the dynamics of the variable θ is defined by a general differential equation given by

$$\tau_\theta \dot{\theta} = -\theta + \theta_\infty(V), \quad (4.15)$$

where τ_θ is the characteristic timescale on which the threshold reacts to changes in the membrane potential and $\theta_\infty(V)$ is the voltage-dependent steady-state towards which θ converges. Depending on the particular shape of $\theta_\infty(V)$, Equation 4.15 can in principle capture a dependency of the firing threshold on the depolarization rate preceding a spike (Platkiewicz and Brette, 2011) (Fig. 4.4). To avoid *a priori* assumptions on the biophysical processes underlying the coupling, $\theta_\infty(V)$ is defined as an arbitrary function of the membrane potential and is extracted from experimental data using a new non-parametric maximum likelihood approach (see below). Finally, it is worth noting that a spike-triggered movement of the firing threshold could in principle be implemented by incrementing the value of θ after the emission of a spike. However, the timescale on which spike-triggered effects occur might be different from τ_θ . For this reason, spike-dependent movements of the firing threshold are modeled by $\gamma(t)$ and the state variable θ is reset to V_T^* after each spike.

The iGIF-Na model is defined exactly as the iGIF model except for the fact that the dynamics of $\theta(t)$ is as in Equations 4.2-4.3. The GIF model (Mensi et al., 2011) is a particular instance of the iGIF model obtained by removing the coupling between subthreshold membrane potential and firing threshold (i.e., by replacing $\theta(t)$ in Eq. 4.14 with a constant V_T^*).

4.4.7 iGIF model parameter extraction

Given the input current $I(t)$, the intracellular membrane potential $V_{\text{data}}(t)$, its first-order derivative $\dot{V}_{\text{data}}(t) = [V_{\text{data}}(t + \Delta T) - V_{\text{data}}(t)] / \Delta T$ and the experimental spike train $\{\hat{t}_j\}$, iGIF

model parameters are obtained with a new two-steps procedure developed by extending the methods introduced in Mensi et al. (2011) and Pozzorini et al. (2013).

In the first step, all the parameters describing the subthreshold dynamics are extracted by minimizing the sum of squared errors between the voltage derivative observed in the experiment and the one predicted by the model (see Eq. 4.13). To allow for convex optimization and avoid *a priori* assumptions on the timescales of adaptation, the spike-triggered conductance was expanded in a linear combination of basis functions $\eta(t) = \sum_{i=1}^K \eta_i b_i^\eta(t)$, where $\{b_i^\eta(t)\}$ is a set of $K = 40$ log-spaced non-overlapping rectangular functions and the parameters $\{\eta_i\}$ define the shape of $\eta(t)$. As in Paninski et al. (2005); Mensi et al. (2011); Pozzorini et al. (2013), the least-square estimate of the subthreshold parameters $\beta_{\text{sub}}^T(E_R) = C^{-1} \cdot [g_0, E_0 g_0, \eta_1, \dots, \eta_K, 1]$ is obtained by solving a multilinear regression problem:

$$\hat{\beta}_{\text{sub}}(E_R) = (X^T X)^{-1} X^T \dot{V}_{\text{data}}, \quad (4.16)$$

where X is a matrix made of vectors

$$X_t^T(E_R) = \left[-V_{\text{data}}(t), 1, \sum_j b_1^\eta(t - \hat{t}_j)(V - E_R), \dots, \sum_j b_K^\eta(t - \hat{t}_j)(V - E_R), I(t) \right], \quad (4.17)$$

and \dot{V}_{data} is a vector containing the membrane potential first-order derivative. Since the model does not capture the voltage trajectory during a spike, all the data points close to action potentials $\{t | t \in [\hat{t}_j - 5 \text{ ms}; \hat{t}_j + T_{\text{ref}}]\}$ were excluded from the fit. As in Mensi et al. (2011), the optimal reversal potential \hat{E}_R (defined as the E_R minimizing the residuals of the regression in Eq. 4.16) is extracted by an exhaustive search on the interval $[-100, -40]$ mV. Finally, the absolute refractory period was set to $T_{\text{ref}} = 4$ ms and the voltage reset was estimated by computing the average membrane potential after a spike (i.e. $V_{\text{reset}} = \langle V(\hat{t}_j + T_{\text{ref}}) \rangle_j$). Since a period of absolute refractoriness can also be implemented by setting the first milliseconds of the spike-triggered threshold movement $\gamma(t)$ to high values, the particular choice of T_{ref} is not crucial.

In the second step, an estimate of the subthreshold membrane potential $\hat{V}(t)$ is obtained by numerically solving Equation 4.13 and the voltage threshold parameters are extracted by extending the non-parametric maximum-likelihood approach of Mensi et al. (2011); Pozzorini et al. (2013). Again, to avoid *a priori* assumptions on the timescales of spike-dependent adaptation and on the shape of the coupling between firing threshold and subthreshold membrane potential, the two functions $\gamma(t)$ and $\theta_\infty(V)$ were expanded in linear combinations of non-overlapping rectangular basis functions $\gamma(t) = \sum_{i=1}^K \gamma_i b_i^{(\gamma)}(t)$ and $\theta_\infty(V) = V_T^* + \sum_{i=1}^M \theta_i b_i^{(\theta)}(V)$. For the spike-triggered movement of the firing threshold $\gamma(t)$, the same log-spaced rectangular functions already used for $\eta(t)$ were chosen. For $\theta_\infty(V)$, $M = 11$ regularly spaced rectangular functions $\{b_i^{(\theta)}(V)\}$ were chosen that covered the interval of voltages $[\min_j \{\hat{V}(\hat{t}_j)\}, \max_j \{\hat{V}(\hat{t}_j)\}]$ at which action potentials were initiated. Consequently, after integration of Equation 4.15, the

time-dependent voltage threshold is given by

$$V_T(t) = V_T^* + \sum_{i=1}^K \gamma_i \cdot \sum_{\hat{t}_j < t} b_i^{(\gamma)}(t - \hat{t}_j - T_{\text{ref}}) + \sum_{i=1}^M \theta_i f_i(t; \tau_\theta), \quad (4.18)$$

with $f_i(t; \tau_\theta) = \int_{\hat{t}_{\text{last}}}^t \tau_\theta^{-1} e^{-\frac{s}{\tau_\theta}} \cdot b_i^{(\theta)}(V(t-s)) ds$ and \hat{t}_{last} denoting the time of the last spike before t . With the exponential function in Equation 4.12, and assuming that the timescale τ_θ is known, the model log-likelihood is a convex function of the threshold parameters $\beta_{\text{th}}^T = \Delta V^{-1} \cdot [1, V_T^*, \gamma_1, \dots, \gamma_K, \theta_1, \dots, \theta_M]$ and can be written as follows (Paninski, 2004):

$$L(\beta_{\text{th}}; \tau_\theta) = \log p(\{\hat{t}_j\} | \hat{V}(t); \beta_{\text{th}}, \tau_\theta) = \sum_{t \in \{\hat{t}_j\}} Y_t(\tau_\theta) \cdot \beta_{\text{th}} - \Delta T \cdot \sum_{t \in \Omega} \exp(Y_t(\tau_\theta) \cdot \beta_{\text{th}}), \quad (4.19)$$

with $\Omega = \{t | t \notin [\hat{t}_j, \hat{t}_j + T_{\text{ref}}]\}$ being a set that excludes all the points falling in the period of absolute refractoriness and $Y_t(\tau_\theta)$ being a vector of observables that implicitly depends on the parameter τ_θ :

$$Y_t(\tau_\theta) = \left[\hat{V}(t), -1, -\sum_j b_1^{(\gamma)}(t - \hat{t}_j), \dots, -\sum_j b_K^{(\gamma)}(t - \hat{t}_j), -f_1(t; \tau_\theta), \dots, -f_M(t; \tau_\theta) \right]. \quad (4.20)$$

Given τ_θ , the maximum likelihood estimate of the other threshold parameters β_{th} can be obtained as in Mensi et al. (2011); Pozzorini et al. (2013):

$$\hat{\beta}_{\text{th}}(\tau_\theta) = \underset{\beta_{\text{th}}}{\operatorname{argmax}} \{L(\beta_{\text{th}}; \tau_\theta)\} \quad (4.21)$$

by maximizing Equation 4.19 with standard gradient-ascent methods. The optimal timescale of the coupling between threshold and membrane voltage $\hat{\tau}_\theta = \underset{\tau_\theta}{\operatorname{argmax}} \{L(\hat{\beta}_{\text{th}}(\tau_\theta); \tau_\theta)\}$ is then obtained by systematically searching in the range $\tau_\theta \in [0.5 \text{ ms}, 15 \text{ ms}]$ the value for which the log-likelihood is maximized. Finally, it is worth noting that even if we do not dispose of a proof of joint convexity, the landscape of the log-likelihood function $L(\hat{\beta}_{\text{th}}(\tau_\theta); \tau_\theta)$ was smooth in τ_θ and always contained a unique maximum in the explored range (see Fig. 4.5f).

Fitting procedure for the iGIF-Na model

The iGIF-Na model parameters were extracted from experimental data using a maximum likelihood approach closely resembling to the nonparametric method described in the previous section. Briefly, the log-likelihood $L(\beta_{\text{th}}^{\text{Na}}; \tau_\theta, k_i, V_i)$ of the iGIF-Na model is convex in $\beta_{\text{th}}^{\text{Na}} = \Delta V^{-1} \cdot [1, V_T^*, \gamma_1, \dots, \gamma_K, k_a]$. Consequently, given the nonlinear parameters k_i, V_i and τ_θ , all the other threshold parameters can be easily extracted by solving a convex optimization problem:

$$\hat{\beta}_{\text{th}}^{\text{Na}}(\tau_\theta, k_i, V_i) = \underset{\beta_{\text{th}}^{\text{Na}}}{\operatorname{argmax}} \{L(\beta_{\text{th}}^{\text{Na}}; \tau_\theta, k_i, V_i)\}. \quad (4.22)$$

On the other hand, extracting the optimal parameters \hat{k}_i , \hat{V}_i and $\hat{\tau}_\theta$ requires the solution of the following nonlinear optimization problem:

$$(\hat{k}_i, \hat{V}_i, \hat{\tau}_\theta) = \underset{(\tau_\theta, k_i, V_i)}{\operatorname{argmax}} \{L(\hat{\beta}_{\text{th}}^{\text{Na}}(\tau_\theta, k_i, V_i); \tau_\theta, k_i, V_i)\}. \quad (4.23)$$

Performing an exhaustive search on a three-dimensional space is possible. Model parameters were however extracted by first fixing the coupling timescale τ_θ to the optimal value previously obtained by fitting the iGIF-NP model and then performing an exhaustive search for k_i and V_i (see Fig. 4.5g).

4.4.8 Extracting the effective membrane timescale from intracellular recordings

In order to extract the effective membrane timescale τ_m from experimental data (see Fig. 4.2d), intracellular recordings were split in different datasets according to μ_1 and independently fitted with a leaky integrate-and-fire model equipped with a spike-triggered current (i.e., a model obtained by dropping the term $(V - E_R)$ from Eq. 4.13) by performing a linear regression similar to Equation 4.16. By fitting this model to data, the average conductance increase mediated by spike-dependent processes is automatically absorbed by the leak conductance g_0 . Consequently, the membrane timescale is directly given by $\tau_m = C/g_0$.

4.4.9 iGIF model linearization

The subthreshold dynamics of the iGIF model features a spike-triggered conductance. Following Richardson (2004), Equation 4.13 can however be approximated by a leaky integrator equipped with a spike-triggered current $\eta_C(t)$:

$$C\dot{V} = -g_L(V - E_L) + I(t) - \left(\sum_{\hat{t}_j} \eta_C(t - \hat{t}_j) - \bar{I}_A \right) \quad (4.24)$$

where the total conductance $g_L = g_0 + \bar{g}_\eta$ accounts for the average contribution $\bar{g}_\eta = T^{-1} \cdot \int_0^T \sum_{\hat{t}_j} \eta(t - \hat{t}_j) dt$ of spike-triggered adaptation, $E_L = (g_0 E_0 + \bar{g}_\eta E_R) / g_L$ is the effective resting potential, $\bar{I}_A = \bar{g}_\eta (\bar{V} - E_R)$ is a constant accounting for the average current mediated by the spike-triggered conductance and the spike-triggered current is given by $\eta_C(t) = (\bar{V} - E_R) \eta(t)$, with \bar{V} being the average membrane potential.

The subthreshold dynamics of the membrane potential defined in Equation 4.24 can be rewritten in its integral form as (Gerstner and Kistler, 2002):

$$V(t) = E_L + \frac{\bar{I}_A}{g_L} + \int_{-\infty}^{\infty} \kappa_m(s) I(t-s) ds - \sum_{\hat{t}_j < t} \eta_V(t - \hat{t}_j), \quad (4.25)$$

where $\kappa_m(t)$ is the effective membrane filter defined in Equation 4.6 and the spike-triggered

filter $\eta_V(t)$ given by

$$\eta_V(t) = \int_{-\infty}^{\infty} \kappa_m(s) [\eta_C(t-s) + C(\bar{V}_{\text{spikes}} - V_{\text{reset}})\delta(t-s)] ds \quad (4.26)$$

describes the influence on the membrane potential of both the spike-triggered current $\eta_C(t)$ and the spike-after reset $V \rightarrow V_{\text{reset}}$, with $\bar{V}_{\text{spikes}} = \langle V(\hat{t}_j) \rangle_j$ being the average value of V at spike times $\{\hat{t}_j\}$. Since this approximation neglects the correlations between membrane potential fluctuations and conductance fluctuations, we expect $\eta_V(t)$ to underestimate the real voltage change induced by the spike-triggered conductance during the first τ_m milliseconds after a spike.

In order to linearize the firing threshold dynamics, Equation 4.15 is simplified by taking the first-order approximation $\theta_{\infty}(V) \approx \bar{C}_{\theta} + \bar{G}_{\theta}V$, with \bar{G}_{θ} being the average gain of the coupling between membrane potential and firing threshold defined in Equation 4.7 and $\bar{C}_{\theta} = \int_{-\infty}^{\infty} \theta_{\infty}^{\text{Na}}(V)P(V)dV - \bar{G}_{\theta}\bar{V}$ being a constant. By integrating Equation 4.15 over time, the linearized threshold dynamics of the iGIF model reads:

$$V_T(t) = \bar{C}_{\theta} + \bar{G}_{\theta} \cdot \int_{-\infty}^{\infty} \kappa_{\theta}(s)V(t-s)ds + \sum_{\hat{t}_j < t} \gamma_{\text{tot}}(t - \hat{t}_j), \quad (4.27)$$

where $\kappa_{\theta}(t) = \frac{1}{\tau_{\theta}} \exp\left(-\frac{t}{\tau_{\theta}}\right)$ is the linearized threshold-coupling filter and the spike-triggered filter $\gamma_{\text{tot}}(t)$ given by

$$\gamma_{\text{tot}}(t) = \gamma(t) - (\bar{\theta}_{\text{spike}} - V_T^*)\tau_{\theta}\kappa_{\theta}(t) \quad (4.28)$$

combines the effects of both the spike-triggered threshold $\gamma(t)$ and the after-spike reset $\theta \rightarrow V_T^*$, with $\bar{\theta}_{\text{spike}} = \langle \theta(\hat{t}_j) \rangle_j$ being the average value of θ at spike times $\{\hat{t}_j\}$.

In the iGIF model, the spiking probability depends on the difference between membrane potential and firing threshold (Eq. 4.12). Thus, the different terms appearing in Equations 4.27-4.25 can be combined to obtain a compact expression for the linearized iGIF model intensity $\lambda_{\text{lin}}(t)$:

$$\lambda_{\text{lin}}(t) = \lambda_0 \cdot \exp\left(\frac{E_0 + \int_{-\infty}^{\infty} \kappa_{\text{eff}}(s)I(t-s)ds + \sum_{\hat{t}_j < t} h_{\text{eff}}(t - \hat{t}_j)}{\Delta V}\right), \quad (4.29)$$

where E_0 is a constant, $\kappa_{\text{eff}}(t)$ is the effective filter defined in Equation 4.8 and $h_{\text{eff}}(t)$ is an effective spike-history filter that phenomenologically accounts for all the spike-triggered mechanisms in the iGIF model:

$$h_{\text{eff}}(t) = \gamma_{\text{tot}}(t) + \int_{-\infty}^{\infty} [\delta(s) - \bar{G}_{\theta}\kappa_{\theta}(s)] \eta_V(t-s)ds. \quad (4.30)$$

Equation 4.30 indicates that the coupling between membrane potential and firing threshold shunts the spike-triggered adaptation mediated by $\eta_V(t)$. Since the shunting strength is

controlled by the average coupling gain \bar{G}_θ , this theoretical result explains our experimental finding that increasing μ_I results in a shrinkage of the GLM spike-history filter $h_{\text{GLM}}(t)$ (compare Fig. 4.2c and Fig. 4.8c). A list of parameters and variables appearing in this derivation is provided in Table 4.2.

4.4.10 Performance evaluation

To avoid problems related to *overfitting* and allow for a comparison between models that differ in the total number of parameters, the performances reported in this study were, unless specified otherwise, evaluated on separate data sets that were not used for parameter extraction. A quantitative measure of the quality of both the GIF and the iGIF model is provided by the log-likelihood:

$$LL_{\text{model}} = \sum_{t \in \{\hat{t}\}} \log \lambda_{\text{model}}(t) - \int_0^T \lambda_{\text{model}}(t) dt \quad (4.31)$$

where $\lambda_{\text{model}}(t)$ is the conditional firing intensity of the model after parameter optimization, $\{\hat{t}\}$ is the experimental spike train and T is the total duration of the experiment on which the model performance were evaluated. All of the log-likelihoods reported in this study were normalized with respect to a Poisson process with constant intensity defined by the experimental firing rate $\bar{r} = N_{\text{spikes}}/T$, as well as with respect to the total number of spikes N_{spikes} (Pillow et al., 2008):

$$LL = \frac{1}{\log(2) \cdot N_{\text{spikes}}} (LL_{\text{model}} - N_{\text{spikes}} (\log \bar{r} - 1)), \quad (4.32)$$

such that units are in bit per spike.

Spike-timing prediction was quantified using the spike-train similarity measure M_{d}^* (Naud et al., 2011). As in our previous studies (Mensi et al., 2011; Pozzorini et al., 2013), M_{d}^* was computed using the Kistler coincidence window with a temporal granularity of $\Delta = \pm 4$ ms.

4.5 Author contributions

C.P had the original idea; C.P and S.M. designed the study, conceived the model, developed the fitting procedure, derived the analytical results and analyzed the model dynamics; C.P wrote the manuscript and analyzed the experimental data; S.M. implemented and performed the fits and the numerical simulations; O.H. performed the experiments; W.G. supervised the project.

We thank David Kastner, Laureline Lagiaco, Alex Seeholzer and Carlos Stein for helpful discussions. This research was supported by the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 604102 (Human Brain Project) and by the Euro-

4.5. Author contributions

pean Research Council under grant agreement number 268689 (MultiRules) O.H. was funded by a grant from the EPFL to the LNMC.

Chapter 4. Enhanced temporal coding by nonlinear threshold dynamics

Term	Description	Definition
Subthreshold membrane potential dynamics		
$V(t)$	Membrane potential	see Eq. 4.13
$\kappa_0(t)$	Passive membrane filter	$\kappa_0(t) = \Theta(t)C^{-1} \exp(-t/\tau_0)$
τ_0	Passive membrane timescale	$\tau_0 = C/g_0$
C	Cell capacitance	
g_0	Passive leak conductance	
E_0	Passive reversal potential	
$I_A(t)$	Adaptation current	$I_A(t) = \sum_{\hat{t}_j < t} \eta(t - \hat{t}_j)(V(t) - E_R)$
$\eta(t)$	Spike-triggered conductance	
E_R	Reversal potential of adaptation current	
ΔV	Threshold sharpness	
T_{ref}	Absolute refractory period	
V_{reset}	Voltage reset	
Firing threshold dynamics		
$V_T(t)$	Firing threshold	$V_T(t) = \theta(t) + \phi(t)$
$\phi(t)$	Spike-dependent threshold adaptation	$\phi(t) = \sum_{\hat{t}_j < t} \gamma(t - \hat{t}_j)$
$\bar{\phi}$	Mean threshold contribution by $\phi(t)$	$\bar{\phi} = \langle \phi(t) \rangle_t$
$\gamma(t)$	Spike-triggered threshold movement	
$\theta(t)$	Voltage-dependent threshold adaptation	$\tau_\theta \dot{\theta} = -\theta + \theta_\infty^{\text{Na}}(V)$
$\kappa_\theta(t)$	Threshold coupling filter	$\kappa_\theta(t) = \Theta(t)\tau_\theta^{-1} \exp(-t/\tau_\theta)$
τ_θ	Threshold coupling timescale	
Firing threshold coupling		
$\theta_\infty^{\text{Na}}(V)$	Steady-state threshold coupling	$\theta_\infty^{\text{Na}}(V) = V_T^* + k_a \log\left(1 + \exp\left(\frac{V - V_i}{k_i}\right)\right)$
V_T^*	Threshold baseline	
V_i	Na ⁺ -channel half-inactivation	
k_a	Slope of Na ⁺ -channel activation	
k_i	Slope of Na ⁺ -channel inactivation	
$G_\theta(V)$	Local gain of threshold coupling	$G_\theta(V) = \frac{d}{dV} \theta_\infty^{\text{Na}}(V)$
$\theta_{\text{slope}}^{\text{Na}}$	Asymptotic gain of threshold coupling	$\theta_{\text{slope}}^{\text{Na}} = \lim_{V \rightarrow \infty} G_\theta(V) = k_a/k_i$

Table 4.1: Parameters and variables of the iGIF-Na model.

Term	Description	Definition
Effective somatic integration		
$\kappa_{\text{eff}}(t)$	Effective filter of somatic integration	$\kappa_{\text{eff}}(t) = \kappa_{\text{m}}(t) - \bar{G}_{\theta} \int \kappa_{\theta}(s) \kappa_{\text{m}}(t-s) ds$
$\tau_{\text{eff}}(t)$	Timescale of $\kappa_{\text{eff}}(t)$	extracted from $\kappa_{\text{eff}}(t)$
\bar{G}_{θ}	Average threshold-voltage gain	$\bar{G}_{\theta} = \langle G_{\theta}(V) \rangle_{P(V)}$
Membrane filter		
$\kappa_{\text{m}}(t)$	Membrane filter	$\kappa_{\text{m}}(t) = \Theta(t) C^{-1} \exp(-t/\tau_{\text{m}})$
τ_{m}	Membrane timescale	$\tau_{\text{m}} = C/g_{\text{L}}$
g_{L}	Average total conductance	$g_{\text{L}} = g_0 + \bar{g}_{\eta}$
\bar{g}_{η}	Mean conductance increase by $\eta(t)$	$\bar{g}_{\eta} = \left\langle \sum_{\hat{t}_j < t} \eta(t - \hat{t}_j) \right\rangle_t$
E_{L}	Reversal potential	$E_{\text{L}} = (g_0 E_0 + \bar{g}_{\eta} E_{\text{R}}) / g_{\text{L}}$
\bar{I}_{A}	Mean adaptation current	$\bar{I}_{\text{A}} = \bar{g}_{\eta} (\bar{V} - E_{\text{R}})$
Spike-history dependence		
$h_{\text{eff}}(t)$	Effective spike-history filter	see Eq. 4.30
$\eta_{\text{C}}(t)$	Spike-triggered current	$\eta_{\text{C}}(t) = (\bar{V} - E_{\text{R}}) \eta(t)$
$\eta_{\text{V}}(t)$	Effect of $\eta_{\text{C}}(t)$ and reset $V \rightarrow V_{\text{reset}}$	see Materials and Methods
$\gamma_{\text{tot}}(t)$	Effect of $\gamma(t)$ and reset $\theta \rightarrow V_{\text{T}}^*$	see Materials and Methods
Generalized linear model		
$\kappa_{\text{GLM}}(t)$	GLM linear filter	
$\tau_{\text{GLM}}(t)$	Timescale of $\kappa_{\text{GLM}}(t)$	extracted from $\kappa_{\text{GLM}}(t)$
$h_{\text{GLM}}(t)$	GLM spike-history filter	

Table 4.2: Parameters and variables of the reduced iGIF-Na model and the GLM.

Bibliography

- L. Abbott and C. van Vreeswijk. Asynchronous states in networks of pulse-coupled oscillators. *Physical Review E*, 48(2):1483, 1993.
- L. Abbott, J. Varela, K. Sen, and S. Nelson. Synaptic depression and cortical gain control. *Science*, 275(5297):221–224, 1997.
- M. Abeles. Role of the cortical neuron: integrator or coincidence detector? *Israel journal of medical sciences*, 18(1):83–92, 1982.
- E. D. Adrian and Y. Zotterman. The impulses produced by sensory nerve-endings part ii. the response of a single end-organ. *The Journal of physiology*, 61(2):151–171, 1926.
- R. B. Anderson. The power law as an emergent property. *Memory & cognition*, 29(7):1061–8, 2001.
- M. Arsiero, H.-R. Lüscher, B. N. Lundstrom, and M. Giugliano. The impact of input fluctuations on the frequency–current relationships of layer 5 pyramidal neurons in the rat medial prefrontal cortex. *The Journal of neuroscience*, 27(12):3274–3284, 2007.
- J. J. Atick. Could information theory provide an ecological theory of sensory processing? *Network: Computation in neural systems*, 3(2):213–251, 1992.
- J. J. Atick and A. N. Redlich. Towards a theory of early visual processing. *Neural Computation*, 2(3):308–320, 1990.
- J. J. Atick and A. N. Redlich. What does the retina know about natural scenes? *Neural computation*, 4(2):196–210, 1992.
- D. Attwell and S. B. Laughlin. An energy budget for signaling in the grey matter of the brain. *Journal of Cerebral Blood Flow & Metabolism*, 21(10):1133–1145, 2001.
- M. Avermann, C. Tamm, C. Mateo, W. Gerstner, and C. Petersen. Microcircuits of excitatory and inhibitory neurons in layer 2/3 of mouse barrel cortex. *J. Neurophysiol.*, 107(11):3116–3134, 2012.
- O. Avila-Akerberg and M. J. Chacron. Nonrenewal spike train statistics: causes and functional consequences on neural coding. *Exp. Brain Res.*, 210(3-4):353–371, 2011.

Bibliography

- R. Azouz and C. M. Gray. Dynamic spike threshold reveals a mechanism for coincidence detection in cortical neurons in vivo. *Proc. National Academy of Sciences USA*, 97:8110–8115, 2000.
- R. Azouz and C. M. Gray. Adaptive coincidence detection and dynamic gain control in visual cortical neurons in vivo. *Neuron*, 37(3):513–523, 2003.
- S. A. Baccus and M. Meister. Fast and slow contrast adaptation in retinal circuitry. *Neuron*, 36(5):909–919, 2002.
- L. Badel, S. Lefort, R. Brette, C. C. Petersen, W. Gerstner, and M. J. Richardson. Dynamic iv curves are reliable predictors of naturalistic pyramidal-neuron voltage traces. *Journal of Neurophysiology*, 99(2):656–666, 2008.
- F. Baldissera, B. Gustafsson, and F. Parmiggiani. A model for refractoriness accumulation and secondary range firing in spinal motoneurons. *Biological Cybernetics*, 24(2):61–65, 1976.
- H. B. Barlow. Possible principles underlying the transformation of sensory messages. *Sensory communication*, pages 217–234, 1961.
- J. Benda and A. V. Herz. A universal model for spike-frequency adaptation. *Neural computation*, 15(11):2523–2564, 2003.
- A. Bhattacharjee and L. K. Kaczmarek. For K^+ channels, Na^+ is the new Ca^{2+} . *Trends in neurosciences*, 28(8):422–428, 2005.
- G.-q. Bi and M.-m. Poo. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *The Journal of neuroscience*, 18(24):10464–10472, 1998.
- S. M. Bohte and J. O. Rombouts. Fractionally Predictive Spiking Neurons. In *NIPS*, pages 253–261, 2010.
- C. Boucsein, M. Nawrot, S. Rotter, A. Aertsen, and D. Heck. Controlling synaptic input patterns in vitro by dynamic photo stimulation. *Journal of neurophysiology*, 94(4):2948–2958, 2005.
- T. Branco, B. A. Clark, and M. Häusser. Dendritic discrimination of temporal input sequences in cortical neurons. *Science*, 329(5999):1671–1675, 2010.
- N. Brenner, W. Bialek, and R. de Ruyter van Steveninck. Adaptive rescaling maximizes information transmission. *Neuron*, 26(3):695–702, 2000.
- R. Brette. Spiking models for level-invariant encoding. *Frontiers in computational neuroscience*, 5:63, 2011.
- R. Brette. Dynamics of one-dimensional spiking neuron models. *Journal of mathematical biology*, 48(1):38–56, 2004.

- R. Brette and W. Gerstner. Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *Journal of neurophysiology*, 94(5):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.
- D. R. Brillinger. Maximum likelihood analysis of spike trains of interacting nerve cells. *Biological Cybernetics*, 59(3):189–200, 1988.
- D. R. Brillinger and J. P. Segundo. Empirical examination of the threshold model of neuron firing. *Biological Cybernetics*, 35(4):213–220, 1979.
- D. R. Brillinger, H. L. Bryant Jr, and J. P. Segundo. Identification of synaptic interactions. *Biological cybernetics*, 22(4):213–228, 1976.
- D. Brown and P. Adams. Muscarinic suppression of a novel voltage-sensitive k⁺ current in a vertebrate neurone. *Nature*, 283(5748):673–676, 1980.
- N. Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of computational neuroscience*, 8(3):183–208, 2000.
- N. Brunel and V. Hakim. Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural computation*, 11(7):1621–1671, 1999.
- N. Brunel and M. C. Van Rossum. Lapicque’s 1907 paper: from frogs to integrate-and-fire. *Biological cybernetics*, 97(5-6):337–339, 2007.
- N. Brunel and X.-J. Wang. What determines the frequency of fast network oscillations with irregular neural discharges? i. synaptic dynamics and excitation-inhibition balance. *Journal of neurophysiology*, 90(1):415–430, 2003.
- R. M. Bruno. Synchrony in sensation. *Current opinion in neurobiology*, 21(5):701–708, 2011.
- A. Calabrese, J. W. Schumacher, D. M. Schneider, L. Paninski, and S. M. Woolley. A generalized linear model for estimating spectrotemporal receptive fields from responses to natural sounds. *PLoS One*, 6(1):e16104, 2011.
- M. Carandini and D. Ferster. Membrane potential and firing rate in cat primary visual cortex. *The Journal of Neuroscience*, 20(1):470–484, 2000.
- J. A. Cardin, R. D. Kumbhani, D. Contreras, and L. A. Palmer. Cellular mechanisms of temporal sensitivity in visual cortex neurons. *The Journal of neuroscience*, 30(10):3652–3662, 2010.
- N. T. Carnevale and M. L. Hines. *The NEURON book*. Cambridge University Press, 2006.
- W. A. Catterall, I. M. Raman, H. P. Robinson, T. J. Sejnowski, and O. Paulsen. The hodgkin-huxley heritage: from channels to circuits. *The Journal of Neuroscience*, 32(41):14064–14073, 2012.

Bibliography

- M. Chacron, B. Lindner, and A. Longtin. Noise Shaping by Interval Correlations Increases Information Transfer. *Phys. Rev. Lett.*, 92(8), 2004.
- M. J. Chacron, B. Lindner, and A. Longtin. Threshold fatigue and information transfer. *Journal of computational neuroscience*, 23(3):301–311, 2007.
- F. S. Chance, L. Abbott, and A. D. Reyes. Gain modulation from background synaptic input. *Neuron*, 35(4):773–782, 2002.
- C. Clopath, R. Jolivet, A. Rauch, H.-R. Lüscher, and W. Gerstner. Predicting neuronal activity with simple models of the threshold type: Adaptive exponential integrate-and-fire model with two compartments. *Neurocomputing*, 70(10):1668–1673, 2007.
- J. Connor and C. Stevens. Prediction of repetitive firing behaviour from voltage clamp data on an isolated neurone soma. *The Journal of Physiology*, 213(1):31, 1971.
- T. M. Cover and J. A. Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- S. Crochet, J. F. A. Poulet, Y. Kremer, and C. C. H. Petersen. Synaptic mechanisms underlying sparse coding of active touch. *Neuron*, 69(6):1160–1175, 2011.
- Y. Dan, J. Atick, and R. C. Reid. Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *J. Neurosci.*, 16(10):3351–3362, 1996.
- P. Dayan and L. F. Abbott. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. The MIT Press, 2001.
- E. De Schutter and J. M. Bower. An active membrane model of the cerebellar purkinje cell. i. simulation of current clamps in slice. *J. Neurophysiol*, 71, 1994.
- T. Deemyad, J. Kroeger, and M. J. Chacron. Sub-and suprathreshold adaptation currents have opposite effects on frequency tuning. *The Journal of physiology*, 590(19):4839–4858, 2012.
- M. Deger, T. Schwalger, R. Naud, and W. Gerstner. Dynamics of interacting finite-sized networks of spiking neurons with adaptation. *arXiv preprint arXiv:1311.4206*, 2013.
- A. Destexhe and T. Bal. *Dynamic-Clamp*. Springer, 2009.
- A. Destexhe and D. Paré. Impact of network activity on the integrative properties of neocortical pyramidal neurons in vivo. *Journal of Neurophysiology*, 81(4):1531–1547, 1999.
- A. Destexhe, M. Rudolph, J.-M. Fellous, and T. J. Sejnowski. Fluctuating synaptic conductances recreate in vivo-like activity in neocortical neurons. *Neuroscience*, 107(1):13–24, 2001.
- A. Destexhe, M. Rudolph, and D. Paré. The high-conductance state of neocortical neurons in vivo. *Nature reviews neuroscience*, 4(9):739–751, 2003.
- B. Doiron, A. Longtin, R. W. Turner, and L. Maler. Model of gamma frequency burst discharge generated by conditional backpropagation. *Journal of neurophysiology*, 86(4):1523–1545, 2001.

- B. Doiron, A.-M. M. Oswald, and L. Maler. Interval coding. ii. dendrite-dependent mechanisms. *Journal of neurophysiology*, 97(4):2744–2757, 2007.
- D. W. Dong and J. J. Atick. Temporal decorrelation: a theory of lagged and nonlagged responses in the lateral geniculate nucleus. *Network: Computation in Neural Systems*, 6(2):159–178, 1995a.
- D. W. Dong and J. J. Atick. Statistics of natural time-varying images. *Network: Computation in Neural Systems*, 6(3):345–358, 1995b.
- Y. Dong, S. Mihalas, A. Russell, R. Etienne-Cummings, and E. Niebur. Estimating parameters of generalized integrate-and-fire neurons from the maximum likelihood of spike trains. *Neural computation*, 23(11):2833–2867, 2011.
- Y. Dong, S. Mihalas, S. S. Kim, T. Yoshioka, S. Bensmaia, and E. Niebur. A simple model of mechanotransduction in primate glabrous skin. *Journal of Neurophysiology*, 109(5):1350–1359, 2013.
- P. J. Drew and L. F. Abbott. Models and properties of power-law adaptation in neural systems. *J. Neurophysiol.*, 96(2):826–833, 2006.
- S. Druckmann, Y. Banitt, A. Gidon, F. Schürmann, H. Markram, and I. Segev. A novel multiple objective optimization framework for constraining conductance-based neuron models by experimental data. *Frontiers in Neuroscience*, 1(1):1, 2007.
- J. Dunlop, M. Bowlby, R. Peri, D. Vasilyev, and R. Arias. High-throughput electrophysiology: an emerging paradigm for ion-channel screening and physiology. *Nature Reviews Drug Discovery*, 7(4):358–368, 2008.
- F. A. Dunn and F. Rieke. The impact of photoreceptor noise on retinal gain controls. *Current opinion in neurobiology*, 16(4):363–370, 2006.
- S. E. Eikenberry and V. Z. Marmarelis. A nonlinear autoregressive volterra model of the hodgkin–huxley equations. *Journal of computational neuroscience*, 34(1):163–183, 2013.
- A. L. Fairhall, G. D. Lewen, W. Bialek, and R. de Ruyter van Steveninck. Multiple timescales of adaptation in a neural code. *Advances in neural information processing systems*, 2001a.
- A. L. Fairhall, G. D. Lewen, W. Bialek, and R. R. d. R. van Steveninck. Efficiency and ambiguity in an adaptive neural code. *Nature*, 412(6849):787–792, 2001b.
- A. Fairhall and W. Bialek. Adaptive spike coding. *The Handbook of Brain Theory and Neural Networks 2nd edn.* (ed. Arbib, MA), 2002.
- J. Feng. Is the integrate-and-fire model good enough?—a review. *Neural Networks*, 14(6):955–975, 2001.

Bibliography

- F. R. Fernandez, T. Broicher, A. Truong, and J. A. White. Membrane voltage fluctuations reduce spike frequency adaptation and preserve output gain in cal pyramidal neurons in a high-conductance state. *The Journal of Neuroscience*, 31(10):3880–3893, 2011.
- M. J. Ferragamo and D. Oertel. Octopus cells of the mammalian ventral cochlear nucleus sense the rate of depolarization. *Journal of Neurophysiology*, 87(5):2262–2270, 2002.
- D. J. Field. Relations between the statistics of natural images and the response properties of cortical cells. *JOSA A*, 4(12):2379–2394, 1987.
- A. Finkel, A. Wittel, N. Yang, S. Handran, J. Hughes, and J. Costantin. Population patch clamp improves data consistency and success rates in the measurement of ionic currents. *Journal of Biomolecular Screening*, 11(5):488–496, 2006.
- I. A. Fleidervish, A. Friedman, and M. Gutnick. Slow inactivation of Na⁺ current and slow cumulative spike adaptation in mouse and guinea-pig neocortical neurones in slices. *The Journal of Physiology*, 493(Pt 1):83–97, 1996.
- B. Fontaine, K. M. MacLeod, S. T. Lubejko, L. J. Steinberg, C. Köppl, and J. L. Pena. Emergence of band-pass filtering through adaptive spiking in the owl's cochlear nucleus. *Journal of Neurophysiology*, 2014a.
- B. Fontaine, J. L. Peña, and R. Brette. Spike-threshold adaptation predicted by membrane potential dynamics in vivo. *PLoS computational biology*, 10(4):e1003560, 2014b.
- N. Fourcaud-Trocmé, D. Hansel, C. Van Vreeswijk, and N. Brunel. How spike generation mechanisms determine the neuronal response to fluctuating inputs. *The Journal of neuroscience*, 23(37):11628–11640, 2003.
- A. S. French and P. H. Torkkeli. The power law of sensory adaptation: simulation by a model of excitability in spider mechanoreceptor neurons. *Annals of biomedical engineering*, 36(1):153–161, 2008.
- M. Fuortes and F. Mantegazzini. Interpretation of the repetitive firing of nerve cells. *The Journal of general physiology*, 45(6):1163–1179, 1962.
- C. D. Geisler and J. M. Goldberg. A stochastic model of the repetitive activity of neurons. *Biophysical journal*, 6(1):53–69, 1966.
- F. Gerhard, T. Kispersky, G. J. Gutierrez, E. Marder, M. Kramer, and U. Eden. Successful reconstruction of a physiological circuit with known connectivity from spiking activity alone. *PLoS computational biology*, 9(7):e1003138, 2013.
- W. Gerstner and W. M. Kistler. *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge university press, 2002.
- W. Gerstner and R. Naud. How good are neuron models? *Science*, 326(5951):379–380, 2009.

- W. Gerstner and J. L. van Hemmen. Associative memory in a network of spiking neurons. *Network: Computation in Neural Systems*, 3(2):139–164, 1992.
- W. Gerstner, W. Kistler, R. Naud, and L. Paninski. *From single neurons to networks and models of cognition*. Cambridge University Press, 2014.
- M.-O. Gewaltig and M. Diesmann. Nest (neural simulation tool). *Scholarpedia*, 2(4):1430, 2007.
- G. Gilboa, R. Chen, and N. Brenner. History-dependent multiple-time-scale dynamics in a single-neuron model. *The Journal of neuroscience*, 25(28):6479–6489, 2005.
- M. Giugliano, P. Darbon, M. Arsiero, H.-R. Lüscher, and J. Streit. Single-neuron discharge properties and network activity in dissociated cultures of neocortex. *Journal of neurophysiology*, 92(2):977–996, 2004.
- B. Gluss. A model for neuron firing with exponential decay of potential resulting in diffusion equations for probability density. *The Bulletin of mathematical biophysics*, 29(2):233–243, 1967.
- J.-M. Goaillard, A. L. Taylor, D. J. Schulz, and E. Marder. Functional consequences of animal-to-animal variation in circuit parameters. *Nature neuroscience*, 12(11):1424–1430, 2009.
- C. Gold, D. A. Henze, C. Koch, and G. Buzsáki. On the origin of the extracellular action potential waveform: a modeling study. *Journal of Neurophysiology*, 95(5):3113–3128, 2006.
- E. M. Goldberg, B. D. Clark, E. Zagha, M. Nahmani, A. Erisir, and B. Rudy. K-channels at the axon initial segment dampen near-threshold excitability of neocortical fast-spiking gabaergic interneurons. *Neuron*, 58(3):387–400, 2008.
- M. Häusser, N. Spruston, and G. J. Stuart. Diversity and dynamics of dendritic signaling. *Science*, 290(5492):739–744, 2000.
- E. Hay, S. Hill, F. Schürmann, H. Markram, and I. Segev. Models of neocortical layer 5b pyramidal cells capturing a wide range of dendritic and perisomatic active properties. *PLoS computational biology*, 7(7):e1002107, 2011.
- D. Henze and G. Buzsaki. Action potential threshold of hippocampal pyramidal cells in vivo is increased by recent spiking activity. *Neuroscience*, 105(1):121–130, 2001.
- L. Hertäg, J. Hass, T. Golovko, and D. Durstewitz. An approximation to the adaptive exponential integrate-and-fire neuron model allows fast and predictive fitting to physiological data. *Frontiers in computational neuroscience*, 6, 2012.
- A. V. Herz, T. Gollisch, C. K. Machens, and D. Jaeger. Modeling single-neuron dynamics and computations: a balance of detail and abstraction. *Science*, 314(5796):80–85, 2006.

Bibliography

- M. H. Higgs and W. J. Spain. Kv1 channels control spike threshold dynamics and spike timing in cortical pyramidal neurones. *The Journal of physiology*, 589(21):5125–5142, 2011.
- M. H. Higgs, S. J. Slee, and W. J. Spain. Diversity of gain modulation by noise in neocortical neurons: regulation by the slow afterhyperpolarization conductance. *The Journal of neuroscience*, 26(34):8787–8799, 2006.
- A. Hill. Excitation and accommodation in nerve. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, 119(814):305–355, 1936.
- B. Hille. *Ionic channels of excitable membranes*. Sinauer, Sunderland, 1992.
- A. L. Hodgkin and A. F. Huxley. Currents carried by sodium and potassium ions through the membrane of the giant axon of loligo. *The Journal of physiology*, 116(4):449, 1952.
- F. W. Howell, J. Dyhrfeld-Johnsen, R. Maex, N. Goddard, and E. D. Schutter. A large-scale model of the cerebellar cortex using pgenesis. *Neurocomputing*, 32:1041–1046, 2000.
- Q. J. Huys, M. B. Ahrens, and L. Paninski. Efficient estimation of detailed single-neuron models. *Journal of neurophysiology*, 96(2):872, 2006.
- V. Ilin, A. Malyshev, F. Wolf, and M. Volgushev. Fast computations in cortical ensembles require rapid initiation of action potentials. *The Journal of Neuroscience*, 33(6):2281–2292, 2013.
- E. M. Izhikevich. Resonate-and-fire neurons. *Neural networks*, 14(6):883–894, 2001.
- E. M. Izhikevich. Which model to use for cortical spiking neurons? *IEEE transactions on neural networks*, 15(5):1063–1070, 2004.
- E. M. Izhikevich. *Dynamical systems in neuroscience : the geometry of excitability and bursting*. MIT Press, Cambridge, Mass., 2007.
- E. M. Izhikevich and G. M. Edelman. Large-scale model of mammalian thalamocortical systems. *Proceedings of the National Academy of Sciences*, 105(9):3593–3598, 2008.
- E. M. Izhikevich et al. Simple model of spiking neurons. *IEEE Transactions on neural networks*, 14(6):1569–1572, 2003.
- J. J. Jack, D. Noble, and R. W. Tsien. *Electric current flow in excitable cells*. 1975.
- C. E. Jahr and C. F. Stevens. Voltage dependence of nmda-activated macroscopic conductances predicted by single-channel kinetics. *The Journal of Neuroscience*, 10(9):3178–3182, 1990.
- R. Jolivet, T. J. Lewis, and W. 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(2), 2004.
- R. Jolivet, A. Rauch, H. Lüscher, and W. Gerstner. Predicting spike timing of neocortical pyramidal neurons by simple threshold models. *J Comput Neurosci*, 21(1):35–49, 2006a.

- R. Jolivet, A. Rauch, H.-R. Lüscher, and W. Gerstner. Predicting spike timing of neocortical pyramidal neurons by simple threshold models. *Journal of computational neuroscience*, 21(1):35–49, 2006b.
- R. Jolivet, R. Kobayashi, A. Rauch, R. Naud, S. Shinomoto, and W. Gerstner. A benchmark test for a quantitative assessment of simple neuron models. *Journal of Neuroscience Methods*, 169(2):417–424, 2008a.
- R. Jolivet, F. Schürmann, T. K. Berger, R. Naud, W. Gerstner, and A. Roth. The quantitative single-neuron modeling competition. *Biological cybernetics*, 99(4-5):417–426, 2008b.
- E. R. Kandel, J. H. Schwartz, T. M. Jessell, et al. *Principles of neural science*. McGraw-Hill New York, 2000.
- J. Keat, P. Reinagel, R. C. Reid, and M. Meister. Predicting every spike: a model for the responses of visual neurons. *Neuron*, 30(3):803–817, 2001.
- W. Kistler, W. Gerstner, and J. Hemmen. Reduction of the hodgkin-huxley equations to a single-variable threshold model. *Neural Computation*, 9(5):1015–1045, 1997.
- B. W. Knight. Dynamics of encoding in a population of neurons. *The Journal of general physiology*, 59(6):734–766, 1972.
- R. Kobayashi, Y. Tsubo, and S. Shinomoto. Made-to-order spiking neuron model equipped with a multi-timescale adaptive threshold. *Frontiers in Computational Neuroscience*, 3:9, 2009.
- C. Koch. *Biophysics of computation : information processing in single neurons*. Oxford University Press, New York, 1999.
- C. Koch and I. Segev. The role of single neurons in information processing. *Nature Neuroscience*, 3:1171–1177, 2000.
- S. B. Kodandaramaiah, G. T. Franzesi, B. Y. Chow, E. S. Boyden, and C. R. Forest. Automated whole-cell patch-clamp electrophysiology of neurons in vivo. *Nature methods*, 9(6):585–587, 2012.
- H. Köndgen, C. Geisler, S. Fusi, X.-J. Wang, H.-R. Lüscher, and M. Giugliano. The dynamical response properties of neocortical neurons to temporally modulated noisy inputs in vitro. *Cereb. Cortex*, 18(9):2086–2097, 2008.
- P. König, A. K. Engel, and W. Singer. Integrator or coincidence detector? the role of the cortical neuron revisited. *Trends in neurosciences*, 19(4):130–137, 1996.
- G. La Camera, A. Rauch, H.-R. Lüscher, W. Senn, and S. Fusi. Minimal models of adapted neuronal response to in vivo-like input currents. *Neural computation*, 16(10):2101–2124, 2004.

Bibliography

- G. La Camera, A. Rauch, D. Thurbon, H.-R. Lüscher, W. Senn, and S. Fusi. Multiple time scales of temporal response in pyramidal and fast spiking cortical neurons. *Journal of neurophysiology*, 96(6):3448–3464, 2006.
- W. Landgren. On the excitation mechanism of the carotid baroreceptors. *Acta physiologica Scandinavica*, 26(1):1–34, 1952.
- S. Lang, V. J. Dercksen, B. Sakmann, and M. Oberlaender. Simulation of signal flow in 3d reconstructions of an anatomically realistic neural network in rat vibrissal cortex. *Neural Networks*, 24(9):998–1011, 2011.
- L. Lapicque. Recherches quantitatives sur l’excitation électrique des nerfs traitée comme une polarisation. *J. Physiol. Pathol. Gen.*, 49:620–635, 1907.
- M. E. Larkum, T. Nevian, M. Sandler, A. Polsky, and J. Schiller. Synaptic integration in tuft dendrites of layer 5 pyramidal neurons: a new unifying principle. *Science*, 325(5941):756–760, 2009.
- P. Latham, B. Richmond, P. Nelson, and S. Nirenberg. Intrinsic dynamics in neuronal networks. i. theory. *Journal of Neurophysiology*, 83(2):808–827, 2000.
- S. B. Laughlin. Energy as a constraint on the coding and processing of sensory information. *Curr. Opin. Neurobiol.*, 11(4):475–480, 2001.
- S. B. Laughlin et al. A simple coding procedure enhances a neuron’s information capacity. *Z. Naturforsch*, 36(910-912):51, 1981.
- S. Lefort, C. Tómm, J. C. F. Sarria, and C. C. H. Petersen. The Excitatory Neuronal Network of the C2 Barrel Column in Mouse Primary Somatosensory Cortex. *Neuron*, 61(2):301–316, 2009.
- M. London and M. Häusser. Dendritic computation. *Annual Review of Neuroscience*, 28: 503–532, 2005.
- S. B. Lowen, L. S. Liebovitch, and J. A. White. Fractal ion-channel behavior generates fractal firing patterns in neuronal models. *Phys. Rev. E Stat. Nonlin. Soft.*, 59(5 Pt B):5970–5980, 1999.
- B. N. Lundstrom, M. H. Higgs, W. J. Spain, and A. L. Fairhall. Fractional differentiation by neocortical pyramidal neurons. *Nature neuroscience*, 11(11):1335–1342, 2008.
- B. N. Lundstrom, A. L. Fairhall, and M. Maravall. Multiple timescale encoding of slowly varying whisker stimulus envelope in cortical and thalamic neurons in vivo. *The Journal of Neuroscience*, 30(14):5071–5077, 2010.
- R. MacGregor and S. Sharpless. Repetitive discharge rate of a simple neuron model with accumulation of after-hyperpolarization conductance. *Brain research*, 64:387–390, 1973.

- D. Madison and R. Nicoll. Control of the repetitive discharge of rat ca 1 pyramidal neurones in vitro. *The Journal of Physiology*, 354(1):319–331, 1984.
- Z. F. Mainen and T. J. Sejnowski. Reliability of spike timing in neocortical neurons. *Science*, 268(5216):1503–1506, 1995.
- Z. F. Mainen and T. J. Sejnowski. Influence of dendritic structure on firing pattern in model neocortical neurons. *Nature*, 382(6589):363–366, 1996.
- D. J. Mar, C. C. Chow, W. Gerstner, R. W. Adams, and J. J. Collins. Noise shaping in populations of coupled model neurons. *Proc. Natl. Acad. Sci. U.S.A.*, 96(18):10450–10455, 1999.
- M. Maravall, R. S. Petersen, A. L. Fairhall, E. Arabzadeh, and M. E. Diamond. Shifts in coding properties and maintenance of information transmission during adaptation in barrel cortex. *PLoS biology*, 5(2):e19, 2007.
- M. Maravall, A. Alenda, M. R. Bale, and R. S. Petersen. Transformation of adaptation and gain rescaling along the whisker sensory pathway. *PloS one*, 8(12):e82418, 2013.
- H. Markram. The blue brain project. *Nature Reviews Neuroscience*, 7(2):153–160, 2006.
- H. Markram, J. Lübke, M. Frotscher, and B. Sakmann. Regulation of synaptic efficacy by coincidence of postsynaptic eps and epsps. *Science*, 275(5297):213–215, 1997.
- 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, 2004.
- V. Z. Marmarelis, D. C. Shin, D. Song, R. E. Hampson, S. A. Deadwyler, and T. W. Berger. Nonlinear modeling of dynamic interactions within neuronal ensembles using principal dynamic modes. *Journal of computational neuroscience*, 34(1):73–87, 2013.
- D. A. McCormick and J. R. Huguenard. A model of the electrophysiological properties of thalamocortical relay neurons. *Journal of Neurophysiology*, 68:1384–1384, 1992.
- J. M. McFarland, Y. Cui, and D. A. Butts. Inferring nonlinear neuronal computation based on physiologically plausible inputs. *PLoS computational biology*, 9(7):e1003143, 2013.
- R. A. Mease, S. Lee, A. T. Moritz, R. K. Powers, M. D. Binder, and A. L. Fairhall. Context-dependent coding in single neurons. *Journal of computational neuroscience*, pages 1–22, 2014.
- I. V. Melnick, S. F. A. Santos, and B. V. Safronov. Mechanism of spike frequency adaptation in substantia gelatinosa neurones of rat. *J. Physiol.*, 559(Pt 2):383–395, 2004.
- S. Mensi, R. Naud, C. Pozzorini, M. Avermann, C. Petersen, and W. Gerstner. Parameter Extraction and Classification of Three Cortical Neuron Types Reveals Two Distinct Adaptation Mechanisms. *J. Neurophysiol.*, 2011.

Bibliography

- T. Mickus, H.-y. Jung, and N. Spruston. Properties of slow, cumulative sodium channel inactivation in rat hippocampal ca1 pyramidal neurons. *Biophysical journal*, 76(2):846–860, 1999.
- Ş. Mihalas and E. Niebur. A generalized linear integrate-and-fire neural model produces diverse spiking behaviors. *Neural Computation*, 21(3):704–718, 2009.
- K. S. Miller. Derivatives of noninteger order. *Mathematics magazine*, pages 183–192, 1995.
- M. N. Miller, B. W. Okaty, and S. B. Nelson. Region-specific spike-frequency acceleration in layer 5 pyramidal neurons mediated by kv1 subunits. *The Journal of Neuroscience*, 28(51):13716–13726, 2008.
- C. Morris and H. Lecar. Voltage oscillations in the barnacle giant muscle fiber. *Biophysical journal*, 35(1):193–213, 1981.
- R. Naud and W. Gerstner. Coding and decoding with adapting neurons: A population approach to the peri-stimulus time histogram. *PLoS computational biology*, 8(10):e1002711, 2012.
- R. Naud and W. Gerstner. *Can we predict every spike?* In Spike Timing: Mechanisms and Function, P.M. Diloranzo and J.D. Victor (Editors) CRC Press, 2013.
- R. Naud, N. Marcille, C. Clopath, and W. Gerstner. Firing patterns in the adaptive exponential integrate-and-fire model. *Biological Cybernetics*, 99(4-5):335–347, 2008.
- R. Naud, F. Gerhard, S. Mensi, and W. Gerstner. Improved similarity measures for small sets of spike trains. *Neural. Comput.*, 23(12):3016–3069, 2011.
- R. Naud, B. Bathellier, and W. Gerstner. Spike-timing prediction in cortical neurons with active dendrites. *Frontiers in Computational Neuroscience*, 8(90), 2014.
- B. Naundorf, F. Wolf, and M. Volgushev. Unique features of action potential initiation in cortical neurons. *Nature*, 440(7087):1060–1063, 2006.
- M. Okatan, M. A. Wilson, and E. N. Brown. Analyzing functional connectivity using a network likelihood model of ensemble neural spiking activity. *Neural computation*, 17(9):1927–1961, 2005.
- K. Padmanabhan and N. N. Urban. Intrinsic biophysical diversity decorrelates neuronal firing while increasing information content. *Nature neuroscience*, 13(10):1276–1282, 2010.
- A. R. Paiva, I. Park, and J. C. Príncipe. A reproducing kernel hilbert space framework for spike train signal processing. *Neural Computation*, 21(2):424–449, 2009.
- L. Paninski. Maximum likelihood estimation of cascade point-process neural encoding models. *Network: Computation in Neural Systems*, 15(4):243–262, 2004.
- L. Paninski, B. Lau, and A. Reyes. Noise-driven adaptation: in vitro and mathematical analysis. *Neurocomputing*, 52:877–883, 2003.

- L. Paninski, J. W. Pillow, and E. P. Simoncelli. Maximum likelihood estimation of a stochastic integrate-and-fire neural encoding model. *Neural computation*, 16(12):2533–2561, 2004.
- L. Paninski, J. Pillow, and E. Simoncelli. Comparing integrate-and-fire models estimated using intracellular and extracellular data. *Neurocomputing*, 65:379–385, 2005.
- L. Paninski, J. Pillow, and J. Lewi. Statistical models for neural encoding, decoding, and optimal stimulus design. *Progress in brain research*, 165:493–507, 2007.
- J. Pillow. Likelihood-based approaches to modeling the neural code. *Bayesian brain: Probabilistic approaches to neural coding*, pages 53–70, 2007.
- J. Pillow, L. Paninski, V. J. Uzzell, E. P. Simoncelli, and E. Chichilnisky. Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model. *The Journal of Neuroscience*, 25(47):11003–11013, 2005.
- J. Pillow, J. Shlens, L. Paninski, A. Sher, A. M. Litke, E. Chichilnisky, and E. P. Simoncelli. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*, 454(7207):995–999, 2008.
- X. Pitkow and M. Meister. Decorrelation and efficient coding by retinal ganglion cells. *Nature neuroscience*, 15(4):628–635, 2012.
- J. Platkiewicz and R. Brette. A threshold equation for action potential initiation. *PLoS computational biology*, 6(7):e1000850, 2010.
- J. Platkiewicz and R. Brette. Impact of fast sodium channel inactivation on spike threshold dynamics and synaptic integration. *PLoS computational biology*, 7(5):e1001129, 2011.
- H. E. Plesser and W. Gerstner. Noise in integrate-and-fire neurons: From stochastic input to escape rates. *Neural Computation*, 12(2):367–384, 2000.
- C. Pozzorini, T. Coletta, and D. Bianchi. Analyse historique du mouvement creationniste aux etats-unis. 2008.
- C. Pozzorini, R. Naud, S. Mensi, and W. Gerstner. Temporal whitening by power-law adaptation in neocortical neurons. *Nature Neuroscience*, 16(7):942–948, 2013.
- S. A. Prescott, S. Ratté, Y. De Koninck, and T. J. Sejnowski. Nonlinear interaction between shunting and adaptation controls a switch between integration and coincidence detection in pyramidal neurons. *The Journal of neuroscience*, 26(36):9084–9097, 2006.
- A. A. Prinz, C. P. Billimoria, and E. Marder. Alternative to hand-tuning conductance-based models: construction and analysis of databases of model neurons. *Journal of Neurophysiology*, 90(6):3998–4015, 2003.
- W. Rall. *Core conductor theory and cable properties of neurons*. In Kandel, E.R., J.M. Brookhardt, and V.M. Mountcastle eds. *Handbook of physiology, cellular biology of neurons*. Bethesda, MD: American Physiological Society, 1977.

Bibliography

- S. Ratté, S. Hong, E. De Schutter, and S. A. Prescott. Impact of neuronal properties on network coding: roles of spike initiation dynamics and robust synchrony transfer. *Neuron*, 78(5): 758–772, 2013.
- A. Rauch, G. La Camera, H.-R. Lüscher, W. Senn, and S. Fusi. Neocortical pyramidal cells respond as integrate-and-fire neurons to in vivo-like input currents. *Journal of neurophysiology*, 90(3):1598–1612, 2003.
- L. Ricciardi. *Diffusion Processes and Related Topics in Biology*. Springer, Berlin, 1977.
- M. J. Richardson. Effects of synaptic conductance on the voltage distribution and firing rate of spiking neurons. *Physical Review E*, 69(5):051918, 2004.
- M. J. Richardson, N. Brunel, and V. Hakim. From subthreshold to firing-rate resonance. *Journal of neurophysiology*, 89(5):2538–2554, 2003.
- F. Rieke, D. Warland, R. de Ruyter van Steveninck, and W. Bialek. *Spikes: Exploring the Neural Code*. MIT Press, Cambridge, MA, USA, 1997.
- C. Rossant, D. F. Goodman, J. Platkiewicz, and R. Brette. Automatic fitting of spiking neuron models to electrophysiological recordings. *Frontiers in neuroinformatics*, 4, 2010.
- C. Rossant, D. F. Goodman, B. Fontaine, J. Platkiewicz, A. K. Magnusson, and R. Brette. Fitting neuron models to spike trains. *Frontiers in neuroscience*, 5:9, 2011a.
- C. Rossant, S. Leijon, A. K. Magnusson, and R. Brette. Sensitivity of noisy neurons to coincident inputs. *The Journal of Neuroscience*, 31(47):17193–17206, 2011b.
- M. Rossum. A novel spike distance. *Neural Computation*, 13(4):751–763, 2001.
- W. A. Rushton. The ferrier lecture, 1962: visual adaptation. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, pages 20–46, 1965.
- H. M. Sakai, J. Wang, and K. Naka. Contrast gain control in the lower vertebrate retinas. *The Journal of general physiology*, 105(6):815–835, 1995.
- M. V. Sanchez-Vives, L. G. Nowak, and D. A. McCormick. Cellular mechanisms of long-lasting adaptation in visual cortical neurons in vitro. *J. Neurosci.*, 20(11):4286–4299, 2000a.
- M. V. Sanchez-Vives, L. G. Nowak, and D. A. McCormick. Cellular mechanisms of long-lasting adaptation in visual cortical neurons in vitro. *The Journal of Neuroscience*, 20(11):4286–4299, 2000b.
- O. Schwartz, J. W. Pillow, N. C. Rust, and E. P. Simoncelli. Spike-triggered neural characterization. *Journal of Vision*, 6(4):13, 2006.
- P. C. Schwindt, W. J. Spain, and W. E. Crill. Long-lasting reduction of excitability by a sodium-dependent potassium current in cat neocortical neurons. *J. Neurophysiol.*, 61(2):233–244, 1989.

- M. Sekerli, C. A. Del Negro, R. H. Lee, and R. J. Butera. Estimating action potential thresholds from neuronal time-series: new metrics and evaluation of methodologies. *Biomedical Engineering, IEEE Transactions on*, 51(9):1665–1672, 2004.
- B. Sengupta, S. B. Laughlin, and J. E. Niven. Consequences of converting graded to action potentials upon neural information coding and energy efficiency. *PLoS computational biology*, 10(1):e1003439, 2014.
- M. N. Shadlen and W. T. Newsome. Noise, neural codes and cortical organization. *Current opinion in neurobiology*, 4(4):569–579, 1994.
- C. E. Shannon. Communication in the Presence of Noise. *Proc. IRE*, 37(1):10–21, 1949.
- J. Shin. Adaptation in spiking neurons based on the noise shaping neural coding hypothesis. *Neural Netw.*, 14(6-7):907–919, 2001.
- S. Shinomoto, H. Kim, T. Shimokawa, N. Matsuno, S. Funahashi, K. Shima, I. Fujita, H. Tamura, T. Doi, K. Kawano, et al. Relating neuronal firing patterns to functional differentiation of cerebral cortex. *PLoS Computational Biology*, 5(7):e1000433, 2009.
- Y. Shu, A. Hasenstaub, M. Badoual, T. Bal, and D. A. McCormick. Barrages of synaptic activity control the gain and sensitivity of cortical neurons. *The Journal of neuroscience*, 23(32):10388–10401, 2003.
- E. P. Simoncelli and B. A. Olshausen. Natural image statistics and neural representation. *Annual review of neuroscience*, 24(1):1193–1216, 2001.
- S. M. Smirnakis, M. J. Berry, D. K. Warland, W. Bialek, and M. Meister. Adaptation of retinal processing to image contrast and spatial scale. *Nature*, 386(6620):69–73, 1997.
- W. R. Softky and C. Koch. The highly irregular firing of cortical cells is inconsistent with temporal integration of random epsps. *The Journal of Neuroscience*, 13(1):334–350, 1993.
- W. Spain and P. Schwindt. Two transient potassium currents in layer V pyramidal neurones from cat sensorimotor cortex. *J. Physiol.*, 1991.
- 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, 1982.
- G. B. Stanley. Reading and writing the neural code. *Nature neuroscience*, 16(3):259–263, 2013.
- R. B. Stein. A theoretical analysis of neuronal variability. *Biophysical Journal*, 5(2):173–194, 1965.
- G. Stuart and N. Spruston. Determinants of voltage attenuation in neocortical pyramidal neuron dendrites. *The Journal of neuroscience*, 18(10):3501–3510, 1998.

Bibliography

- A. Y. Tan, Y. Chen, B. Scholl, E. Seidemann, and N. J. Priebe. Sensory stimulation shifts visual cortex from synchronous to asynchronous states. *Nature*, 2014.
- T. Tchumatchenko, A. Malyshev, F. Wolf, and M. Volgushev. Ultrafast population encoding by cortical neurons. *The Journal of Neuroscience*, 31(34):12171–12179, 2011.
- J. Thorson and M. Biederman-Thorson. Distributed relaxation processes in sensory adaptation spatial nonuniformity in receptors can explain both the curious dynamics and logarithmic statics of adaptation. *Science*, 183(4121):161–172, 1974.
- K. Thurley, W. Senn, and H.-R. Lüscher. Dopamine increases the gain of the input-output response of rat prefrontal pyramidal neurons. *Journal of neurophysiology*, 99(6):2985–2997, 2008.
- A. Toib, V. Lyakhov, and S. Marom. Interaction between duration of activity and time course of recovery from slow inactivation in mammalian brain Na⁺ channels. *The Journal of neuroscience*, 18(5):1893–1903, 1998.
- S. J. Tripathy, K. Padmanabhan, R. C. Gerkin, and N. N. Urban. Intermediate intrinsic diversity enhances neural population coding. *Proceedings of the National Academy of Sciences*, 110(20):8248–8253, 2013.
- T. W. Troyer and K. D. Miller. Physiological gain leads to high isi variability in a simple model of a cortical regular spiking cell. *Neural Computation*, 9(5):971–983, 1997.
- W. Truccolo, U. T. Eden, M. R. Fellows, J. P. Donoghue, and E. N. Brown. A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *Journal of Neurophysiology*, 93(2):1074–1089, 2005.
- M. V. Tsodyks and H. Markram. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proceedings of the National Academy of Sciences*, 94(2):719–723, 1997.
- H. C. Tuckwell. *Introduction to theoretical neurobiology*. Cambridge University Press, 1988.
- N. Ulanovsky, L. Las, and I. Nelken. Processing of low-probability sounds by cortical neurons. *Nature neuroscience*, 6(4):391–398, 2003.
- N. Ulanovsky, L. Las, D. Farkas, and I. Nelken. Multiple time scales of adaptation in auditory cortex neurons. *The Journal of Neuroscience*, 24(46):10440–10453, 2004.
- W. Van Geit, E. De Schutter, and P. Achard. Automated neuron model optimization techniques: a review. *Biological cybernetics*, 99(4-5):241–251, 2008.
- J. Van Hateren. Real and optimal neural images in early vision. *Nature*, 1992.
- J. Van Hateren and A. Van Der Schaaf. Temporal properties of natural scenes. In *Electronic Imaging: Science & Technology*, pages 139–143. International Society for Optics and Photonics, 1996.

- N. G. Van Kampen. *Stochastic processes in physics and chemistry*, volume 1. Elsevier, 1992.
- R. d. R. Van Steveninck and S. Laughlin. The rate of information transfer at graded-potential synapses. *Nature*, 379(6566):642–645, 1996.
- D. V. Vavoulis, V. A. Straub, J. A. Aston, and J. Feng. A self-organizing state-space-model approach for parameter estimation in hodgkin-huxley-type models of single neurons. *PLoS Computational Biology*, 8(3):e1002401, 2012.
- J. D. Victor and K. P. Purpura. Metric-space analysis of spike trains: theory, algorithms and application. *Network: computation in neural systems*, 8(2):127–164, 1997.
- T. P. Vogels, K. Rajan, and L. Abbott. Neural network dynamics. *Annu. Rev. Neurosci.*, 28: 357–376, 2005.
- M. J. Wainwright. Visual adaptation as optimal information transmission. *Vision Res.*, 39(23): 3960–3974, 1999.
- M. M. Waldrop. Computer modelling: Brain in a box. *Nature*, 482(9):456–458, 2012.
- X.-J. Wang, Y. Liu, M. V. Sanchez-Vives, and D. A. McCormick. Adaptation and temporal decorrelation by single neurons in the primary visual cortex. *Journal of neurophysiology*, 89 (6):3279–3293, 2003.
- B. Wark, B. N. Lundstrom, and A. Fairhall. Sensory adaptation. *Current opinion in neurobiology*, 17(4):423–429, 2007.
- B. Wark, A. Fairhall, and F. Rieke. Timescales of inference in visual adaptation. *Neuron*, 61(5): 750–761, 2009.
- J. C. Wester and D. Contreras. Biophysical mechanism of spike threshold dependence on the rate of rise of the membrane potential by sodium channel inactivation or subthreshold axonal potassium current. *Journal of computational neuroscience*, pages 1–17, 2013.
- W. B. Wilent and D. Contreras. Stimulus-dependent changes in spike threshold enhance feature selectivity in rat barrel cortex neurons. *The Journal of neuroscience*, 25(11):2983–2991, 2005.
- J. Xu, A. Guia, D. Rothwarf, M. Huang, K. Sithiphong, J. Ouang, G. Tao, X. Wang, and L. Wu. A benchmark study with seal chip™ planar patch-clamp technology. *Assay and Drug Development Technologies*, 1(5):675–684, 2003.
- Z. Xu, J. R. Payne, and M. E. Nelson. Logarithmic time course of sensory adaptation afferent nerve fibers in a weakly electric fish. *Journal Neurophysiology*, 1996.
- Y. Yu, B. Potetz, and T. S. Lee. The role of spiking nonlinearity in contrast gain control and information transmission. *Vision research*, 45(5):583–592, 2005.

Christian Pozzorini

Rue du Petit-Beaulieu 7
1004 Lausanne
christian.pozzorini@epfl.ch

EDUCATION

Ph. D. Candidate (Neuroscience) Ecole Polytechnique Fédérale de Lausanne (EPFL) School of Life Sciences (SV), Brain Mind Institute (BMI) Laboratory of Computational Neuroscience (Supervisor: Prof. W. Gerstner)	04/2010 – present
Master of Science Ecole Polytechnique Fédérale de Lausanne (EPFL) School of Life Sciences (SV), Specialization in Bioinformatics Master Thesis: “ <i>The functional shape of spike-triggered adaptation</i> ” (Prof. W. Gerstner)	10/2007 – 07/2009
Bachelor of Science Ecole Polytechnique Fédérale de Lausanne (EPFL) School of Life Sciences (SV) Bachelor Thesis: “ <i>In silico evolution of bistable gene regulatory networks</i> ” (Prof. F. Naef)	10/2004 – 07/2007
Cours de mathématiques spéciales (CMS) Ecole Polytechnique Fédérale de Lausanne (EPFL) Entrance exam at EPFL	10/2003 – 07/2004
Scuola Superiore di Informatica di Gestione di Bellinzona (TI) Diploma di Informatico di Gestione	09/2000 – 07/2002
Suola Professionale Commerciale di Locarno (TI) Maturità Professionale Commerciale	09/1996 – 07/2000

PROFESSIONAL EXPERIENCE

Swiss Center of Scientific Computing (Manno, TI) Internship, Software development (Java, SQL)	07/2002 – 02/2003
---	-------------------

SCIENTIFIC PUBLICATIONS

- C. Pozzorini, R. Naud, S. Mensi and W. Gerstner, *Temporal whitening by power-law adaptation in neocortical neurons*, Nature Neuroscience (2013)
- S. Mensi, R. Naud, C. Pozzorini, M. Avermann, C.C. Petersen and W. Gerstner, *Parameter extraction and classification of three cortical neuron types reveals two distinct adaptation mechanisms*, Journal of neurophysiology (2012)
- J. Bieler, C. Pozzorini, and F. Naef, *Whole-embryo modeling of early segmentation in drosophila identifies robust and fragile expression domains*, Biophysical journal (2011)

PUBLICATIONS IN PREPARATION

- C. Pozzorini*, S. Mensi*, O. Hagens and W. Gerstner, *Enhanced temporal coding by nonlinear threshold dynamics in cortical neurons*
- C. Pozzorini*, S. Mensi*, O. Hagens, R. Naud, C. Koch and W. Gerstner, *Automated high-throughput characterization of single neurons by means of simplified spiking models*

- equal contribution

INVITED TALKS

- University of Oxford, Center for Neural Circuits and Behaviour, *Adaptive coding in single neurons* (Oxford, 2014)
- University of Washington, Computational Neuroscience Seminar, *Adaptive coding in single neurons* (Seattle, 2014)
- Allen Institute for Brain Science, Internal Seminar, *Adaptive coding in single neurons* (Seattle, 2014)
- EPFL, Annual BMI retreat, *Encoding and decoding with spiking neuron models* (Villars, 2013)
- EPFL, BMI Research Days, *Power law adaptation in neocortical pyramidal neurons* (Lausanne, 2012)

SELECTED POSTER PRESENTATIONS

- **C. Pozzorini**, S. Mensi, O. Hagens and W. Gerstner, *Nonlinear threshold dynamics implements multiplicative adaptation in pyramidal neurons*, COSYNE 2014
- **C. Pozzorini**, R. Naud, S. Mensi and W. Gerstner, *Temporal decorrelation by power law adaptation in pyramidal neurons*, COSYNE 2013
- **C. Pozzorini**, S. Mensi, O. Hagens and W. Gerstner, *Evidence for a nonlinear coupling between firing threshold and subthreshold membrane potential*, COSYNE 2013
- S. Mensi, R. Naud, **C. Pozzorini**, M. Avermann, C.C. Petersen and W. Gerstner, *Automatic characterization of three cortical neuron types reveals two Distinct adaptation mechanisms*, CNS 2011
- **C. Pozzorini**, R. Naud, S. Mensi and W. Gerstner, *Power law adaptation in spiking neuron models*, COSYNE 2011
- J. Bieler, **C. Pozzorini**, and F. Naef, *Reconstructing the Drosophila segmentation network in real space*, IPCAT 2009

AWARDS AND FELLOWSHIPS

- Swiss Society for Neuroscience, Best Paper Award 2013
- Best Student Poster Award at CNS 2011 (Stockholm)
- Mention of Excellence for Master Studies (EPFL, 2009)
- EPFL Excellence Scholarship for Master Studies (EPFL, 2008)

COLLABORATIONS

- C. C. Petersen and S. Naskar, Laboratory of Sensory Processing, EPFL
Electrophysiological recordings to study spike-frequency adaptation in neocortical neurons
- H. Markram and O. Hagens, Laboratory of Neural Microcircuitry, EPFL
Electrophysiological recordings to study threshold dynamics in neocortical neurons
- C. Koch and S. Mihalas, Allen Institute for Brain Science, Seattle
Pipeline for automatic high-throughput modeling of spiking neuron models

TEACHING ACTIVITIES

Teaching assistant for several Bachelor and Master courses at EPFL:

- Unsupervised and reinforcement learning in neural networks (Prof. W. Gerstner, 2013)
- Machine learning (Prof. M. Seeger, 2012)
- C++ programming (Dr. J. Sam, 2012)
- Supervisor of two Bachelor projects (2011-2012)
- Neural networks and biological modeling (Prof. W. Gerstner, 2011)
- Signals and systems (Prof. M. Unser, 2010)
- Biomathematical programming (Prof. F. Naef, 2007)

ACADEMIC SERVICES

- Tutor at the Summer Research Program in Life Sciences and Technology (EPFL, 2013)
- Served as a referee for the academic journals: PLOS One, PLOS Computational Biology

Languages: Italian (mother tongue), French (bilingual), English (fluent written/spoken), German (elementary proficiency).