Oscillations and gain control in sensory systems

Alexandre Payeur

Thesis submitted to the
Faculty of Graduate and Postdoctoral Studies
in partial fulfillment of the requirements
for the Doctorate in Philosophy degree in Physics

Ottawa-Carleton Institute for Physics
Department of Physics
Faculty of Science
University of Ottawa

À ma famille.

Abstract

Sensory neurons assemble to form networks that process inputs coming from the senses. Through synaptic connections neurons interact and create complex dynamical states in response to these inputs. Networks with different connectivity patterns are thought to display different states and therefore subserve different computational goals.

In this thesis, we mainly study brain rhythms, a dynamical state that occurs in various neural structures. Rhythms are emergent oscillations that typically occur in homogeneous recurrent networks, whose neurons have identical properties and are densely interconnected. Many sensory systems comprise neurons with opposite ON and OFF responses to inputs. We show that homogenous recurrent networks fail to sustain rhythms when ON and OFF neurons are present in equal proportions. This happens even when the network is subjected to spatially correlated inputs, which are known to promote synchronized oscillations. In this context, we adapted the so-called linear response theory to include networks containing ON and OFF neurons with different intrinsic properties. In this asymmetric case, oscillations can be recovered. A simpler approach is to segregate the ON and OFF populations, thus producing two oscillating subnetworks.

The dynamics of purely feedforward networks are studied next. These networks are composed of two or more populations. The populations are connected in a serial fashion, but neurons are unconnected within the populations. This connectivity scheme is drastically different from the fully recurrent network. Yet, this network is shown to display oscillatorylike properties when subjected to spatially correlated stimulation under certain conditions. We also find that this network can implement various types of gain control, depending on the noise in the system and the strength of synaptic interactions. These results establish some unexpected links between feedforward and recurrent networks.

Along the way, we apply our results and conclusions to a well-characterized sensory network, the electrosensory system of weakly electric fish.

Sommaire

Les neurones sensoriels s'assemblent pour former des réseaux qui traitent les signaux venant des sens. Les neurones interagissent *via* les synapses et établissent des états dynamiques complexes en réponse à ces signaux. Il est supposé que des réseaux possédant des patrons de connectivité différents sous-tendent des états dynamiques différents et, de ce fait, accomplissent des tâches différentes.

Dans cette thèse, nous étudions principalement les rythmes neuronaux, un état dynamique survenant dans une variété de structures neuronales. Les rythmes neuronaux sont des oscillations émergentes qui apparaissent habituellement dans des réseaux homogènes et récurrents, i.e. des réseaux où les neurones ont des propriétés identiques et sont densément interconnectés. Plusieurs systèmes sensoriels comprennent des neurones avec des réponses opposées, ON et OFF, aux signaux. Nous montrons que des réseaux homogènes récurrents ne présentent pas de rythmes lorsque des neurones ON et OFF sont présents en quantité égale. Cela survient même lorsque le réseau est soumis à des signaux corrélés spatialement, lesquels sont reconnus comme facilitant les oscillations synchronisées. Dans ce contexte, nous avons adapté la théorie de la réponse linéaire aux réseaux contenant des neurones ON et OFF avec des propriétés intrinsèques différentes. Dans ce cas dit asymétrique, les oscillations peuvent être recouvrées. Une approche plus simple consiste à déconnecter les populations ON et OFF et ainsi générer deux sous-réseaux oscillants.

Nous étudions ensuite la dynamique de réseaux purement "feedforwards". Ces réseaux sont composés d'une ou de plusieurs populations. Les populations sont connectées en série, mais les neurones sont déconnectés à l'intérieur des populations. Ce patron de connectivité est drastiquement différent des réseaux complètement connectés. Néanmoins, ce réseau feedforward présente des propriétés quasi-oscillatoires dans certaines conditions lorsque soumis à des signaux spatialement corrélés. Nous trouvons aussi que ce réseau peut implémenter divers types de contrôles du gain, qui dépendent du bruit dans le système et de la force des interactions synaptiques. Ces résultats établissent un lien inattendu entre les

réseaux feedforwards et récurrents.

En cours de route, nous appliquons nos résultats et conclusions à un système sensoriel bien caractérisé, le système électrosensoriel des poissons faiblement électriques.

Acknowledgments

"[...] combien sont plus heureux que nous les Colomb qui se lancent à travers les mers à la recherche d'un nouveau monde; ils n'ont qu'à aller devant eux. Les difficultés qui les arrêtent, ne viennent que d'obstacles matériels qu'un homme hardi franchit toujours; tandis que nous, ballotés sans cesse sur l'océan des incertitudes, entraînés brusquement par une hypothèse comme un navire par l'aquilon, nous rencontrons tout à coup, ainsi qu'un vent contraire, une doctrine opposée, qui nous ramène, sans espoir, au port dont nous étions sortis."

– Le Docteur Héraclius Gloss, Maupassant, Contes et nouvelles

Science requires that knowledge be transferred. My supervisor André Longtin gave me the autonomy and freedom needed for any creative activity, and the necessary guidance when required. The result was a happy collaboration fuelled by mutual respect. Len Maler—whose name likely stands for *L*iving *E*ncyclopedia of *N*euroscience—has been the outstanding reference on the weakly electric fish. His passion after all these years of research is truly wonderful to witness.

Science is a collaborative endeavour. I had the good fortune to be surrounded by smart and knowledgeable individuals during my PhD. I have had the pleasure to work with Jorge Mejias. His general happiness and stressless attitude were a refreshing contrast to my neurotic obsession with details. Richard Naud was a real inspiration, and still is.

Science is a nonlinear process. One little step forward is on average followed by another step backward (see quote above). These ups and downs are less painful when surrounded by friends and nice colleagues. The Longtin lab is filled with this kind of people, that I wish to thank here. To name but a few: Alexandre Melanson, Alexandre René, Grégory Dumont and Louis Jacques.

Science needs funding. During my PhD various organizations supported my research, either directly through scholarships (FRQNT, OGS, OGF and UOttawa), or indirectly through my supervisors' grants (NSERC-CREATE, CIHR). Collectively, they allowed me to focus on

research. Also, Prof. Longtin—through the CREATE grant—provided me with the privilege to spend a month at EPFL (Lausanne).

Science is not everything. Without my family, (1) I would not be here and (2) most of what I am doing would be meaningless. Merci pour votre support durant toutes ces années.

Contents

List of Tables

List of Figures

List of symbols and acronyms

Symbols

Time constants

τ_d	Axonal propagation delay
τ_s	Synaptic time constant
τ_m	Membrane time constant
τ_r	Absolute refractory period

Membrane dynamics

v, V	Membrane potential	v_R and v_T	Reset and threshold
μ	Bias	σ	Internal noise intensity
ξ, η	Gaussian white noise process	N	Number of neurons
$y(t)$	Spike train	$Y(t) = R_N(t) = \frac{1}{N}\sum_{i=1}^{N} y_i(t)$	Population activity

Synapses

α or s	Synaptic function
$\mathbf{W}$	Connectivity (or weight) matrix
W_{ij}	Synaptic weight from neuron j to neuron i
G	Total synaptic strength

Definitions of functions and mathematical operations

$\Theta(t)$	Heaviside function (1 if $t > 0$, 0 otherwise)
$\delta(t)$	Dirac delta function
$\mathcal{D}_a(z)$	Parabolic cylinder function
$\mathrm{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-z^2} dz$	Error function
$\mathrm{erfc}(x) = 1 - \mathrm{erf}(x)$	Complementary error function
$\langle \cdot \rangle$	Ensemble average of $\cdot$
$\tilde{x}(\omega) = \mathscr{F}\{x\}(\omega) \equiv \int_{-\infty}^{\infty} x(t) e^{i\omega t} dt$	Fourier transform of $x(t)$
$x(t) = \mathscr{F}^{-1}\{\tilde{x}\}(t) \equiv \int_{-\infty}^{\infty} \tilde{x}(\omega) e^{-i\omega t} \frac{d\omega}{2\pi}$	Inverse Fourier transform
$\equiv$	Identity, equivalence or definition

Spike-train statistics

r or f	Firing rate
σ_{ISI}	Standard deviation of interspike intervals
C_V (or CV)	Coefficient of variation
Q	Autocovariance function of a spiking process
C_X	Autocovariance function of signal $X(t)$
C_{XY}	Cross-covariance function of $X(t)$ and $Y(t)$
R_X	Autocorrelation function of signal $X(t)$
R_{XY}	Cross-correlation function of $X(t)$ and $Y(t)$
$X_T(t)$	Restriction of signal $X(t)$ to interval $[0, T]$

Spectral quantities

f	Frequency
ω	Angular frequency ($\omega = 2\pi f$)
$\mathbf{S}$	Spectral matrix
S	Power spectral density or power spectrum
$S^{(0)}$ (or S_0)	Unperturbed power spectral density
S_{pop}	Population spectrum or population power spectral density
S_{cross}	Cross-spectrum
$\tilde{A}$ (or simply A)	Susceptibility
A_{eff}	Effective susceptibility

Acronyms

Electrosensory system

General

EOD	Electric organ discharge	ELL	Electrosensory lateral line lobe
EA	Electroreceptor afferent	AM	Amplitude modulation
EGp	*Eminentia granularis posterioris*	Pd	*Nucleus praeementialis dorsalis*

Segments of the ELL

CMS	Centromedial segment
CLS	Centrolateral segment
LS	Lateral segment

Cell types

PC	Pyramidal cell	DP cell	Deep PC
IP cell	Intermediate PC	SP cell	Superficial PC
GC	Granule cell	BP	Bipolar cell
St	Stellate cell	MP	Multipolar cell

Layers of the ELL

DML	Dorsal molecular layer	VML	Ventral molecular layer
StF	Stratum fibrosum	PCL	Pyramidal cell layer
PIL	Plexiform layer	GCL	Granule cell layerl
DNL	Deep neuropil layer	DFL	Deep fiber layer

Miscellaneous

LRT	Linear response theory	PSD	Power spectral sensity or power spectrum
PopSD	Population power spectral density	GWN	Gaussian white noise
CCF	Cross-correlation function	ACF	Autocorrelation or autocovariance function
FBN	Feedback network	FFN	Feedforward network
ING	Interneuron network gamma	PING	Pyramidal-interneuron network gamma

Introduction

All creatures must acquire information about their environment in order to form percepts of varying levels of complexity. The information is physical or chemical in nature; for instance, photons impinging on the retina, air pressure variations giving rise to sounds, mechanical pressures on the skin causing tactile sensations, or molecules producing odor sensations. In sensory systems, these stimuli are transduced by various types of receptors into signals that can be interpreted by assemblies of neurons. An overarching goal of neurophysics is to understand how this "interpretation" is achieved by networks of biological neurons, i.e., to find neural mechanisms and principles underlying the formation of these perceptions. These perceptions ultimately dictate the appropriate behavioral response of the organism.

In general, the aforementioned mechanisms can be studied from two points of view: etiological or functional. The first pertains to the biological implementation of the mechanism, i.e., how a given neural system performs this task. The second tries to establish the purpose or function of the mechanism, i.e., how it subserves a particular computational goal. In this thesis, we study two of these mechanisms/principles, namely neural rhythms and gain control, from the etiological point of view.

In a nutshell, neural rhythms consist of any coherent oscillatory activity occurring in neural networks. Often, these rhythms are detectable by coarse, macroscopic measures like electroencephalograms. They are ubiquitous across species and brain regions, span many frequency ranges and various functions have been ascribed to them. Rhythms are often induced by well-defined, structured stimulations. Here, neural rhythms in response to spatially coherent stimuli will be the main focus. Gain control, on the other hand, concerns the modulation of the input-output relationships of neurons. It is not difficult to envision how important gain control is for neural dynamics. Gain control can be performed by biophysical processes intrinsic to the neuron; in other cases, it can be a direct consequence of the connectivity of the network to which the neuron belongs. It is the latter possibility that will be studied here. These phenomena—rhythms and gain control—will be studied in

neural networks inspired by the electrosensory system of weakly electric fish. Theoretical neurophysics is in constant need of well-characterized biological networks allowing direct comparison with experiment. The results presented herein, however, are not rigidly bound to this sensory system.

This thesis is organized as follows. In chapter 1 I first introduce some important notions pertaining to the dynamics of neurons and neural networks. I then present the concepts of gain control and neural rhythms in some detail. I try to convey the general importance of these two subjects for theoretical neurophysics in general. The chapter concludes with a description of the electrosensory system. In chapter 2, I present and analyze the neuron model that will be used throughout, and I introduce the linear response theory, a theoretical framework allowing the analysis of neural networks subjected to inputs of various types. Chapters 3 to 5 each contain a peer-reviewed paper. In chapter 3 we study the dynamics of a fully recurrent network—i.e. with all-to-all connectivity—whose neurons respond with opposite polarities (ON and OFF) to external inputs. Contrary to a similar network with neurons responding with the same polarity, that network does not oscillate even when the external input is spatially correlated. Introducing asymmetries between the two populations allows to recover the neural rhythms. We conclude that, in this context, a segregation of the two pathways (ON and OFF) is the simplest way to conserve the oscillatory properties of this type of network and may explain the oscillations seen experimentally with spatially coherent stimuli. Chapter 4 is concerned with gain control in feedforward networks. A feedforward network contains two or more neural populations connected in "series", meaning that the neural signal propagates in a sequential manner. This network displays the well-known subtractive and divisive gain controls, together with a more exotic non-monotonic gain control. Switching between these controls is performed by changing the amount of neural noise and by modifying the strength of synaptic connections. Finally, in chapter 5 we prove that this same feedforward network possesses oscillatory-like properties under spatially correlated stimulation when neuronal noise is large and when signal propagation delays are small. Heterogenous delays accentuate this phenomenon.

Statement of originality

This document is a thesis by articles. The first two chapters present the background information necessary to understand chapters 3 to 5, which contain papers that have already been published in peer-reviewed journals. Here, I delineate my contributions to these articles. All work has been done under the guidance of my co-supervisors.

Chapter 3 This chapter is a copy of the article [159]:

A. Payeur, J. Lefebvre, L. Maler, and A. Longtin. Linear response theory for two neural populations applied to gamma oscillation generation. *Phys Rev E*, **87**:032703, 2013.

I performed all numerical simulations and obtained all the analytical results. I am the main author of the text and figures. Corrections and comments were added by the other three authors.

Chapter 4 This chapter is a copy of the article [144]:

J. F. Mejias*, A. Payeur*, E. Selin, L. Maler, and A. Longtin. Subtractive, divisive and non-monotonic gain control in feedforward nets linearized by noise and delays. *Front Comput Neurosci*, **8**(19), 2014. [* = equal contribution]

Jorge Mejias and I performed the theoretical analysis together. I performed most of the numerical simulations and produced almost all figures. Jorge wrote a first draft of the paper, which I then modified and edited. The section "Comparison with experimental data" and the appendix were written by me alone. The feedforward network studied in this paper comprises an input layer and an output layer. Both layers receive an external input and the input-layer neurons project to (makes synaptic contacts with) the output-layer neurons with a delay. Using the linear response formalism, the appendix details the effect of the interaction between the external input and the delayed synaptic current emanating from the input layer on the input-output relationship of the output-layer neurons. The effect is a linearization of the input-output relationship. This idea of "linearization by delay" is Jorge's and I worked out the mathematical theory. The final draft was commented on and corrected by Profs. Maler and Longtin. E. Selin performed some preliminary work on the topic of the paper.

Chapter 5 This chapter is a copy of the article [160]:

A. Payeur, L. Maler, and A. Longtin. Oscillatorylike behavior in feedforward neuronal networks. *Phys Rev E*, **92**:012703, 2015.

I performed all numerical simulations and obtained all analytical results and figures. Comments and corrections by André Longtin mainly.

Chapter 1

Background

In this chapter we present background information about the topics covered in this thesis. After a discussion of neuronal dynamics, we present the concepts of gain control and neural oscillations. We finish with a description of the neural system serving as canvas for all the research contributions herein. For each topic we begin by presenting the ideas in a general setting, after which we cover in more detail some aspects that are either important for the remainder of the thesis or representative cases that help understand the topic.

1.1 Neural dynamics

Neurons are complex biological entities. A central dogma of neuroscience is that assemblies of interconnected neurons are responsible for the inner working of the mind. Anatomically, neurons consist in a soma from which an axon and multiple dendrites branch out (Fig. 1.1B, left). The axon is the output process which connects to other neurons, whereas the dendrites and the soma constitute the receiving end of the neuron. All neurons maintain a potential difference between their interior and the extracellular medium; the intracellular compartment is negative with respect to the exterior. The intra and extracellular media are separated by a lipid bilayer playing the electrical role of a capacitor (Fig. 1.1A). The potential across the membrane (membrane potential) changes with time according to the opening and closing of ion channels that inhabit the lipid bilayer, allowing ions to enter or leave the cell. These channel dynamics are usually voltage-dependent and can depend on the concentration of ions (typically calcium) as well. When the membrane potential reaches a "threshold", nonlinear processes produce an action potential, i.e. a large positive excursion of the membrane potential followed by a steep decrease and a return to a steady state. This

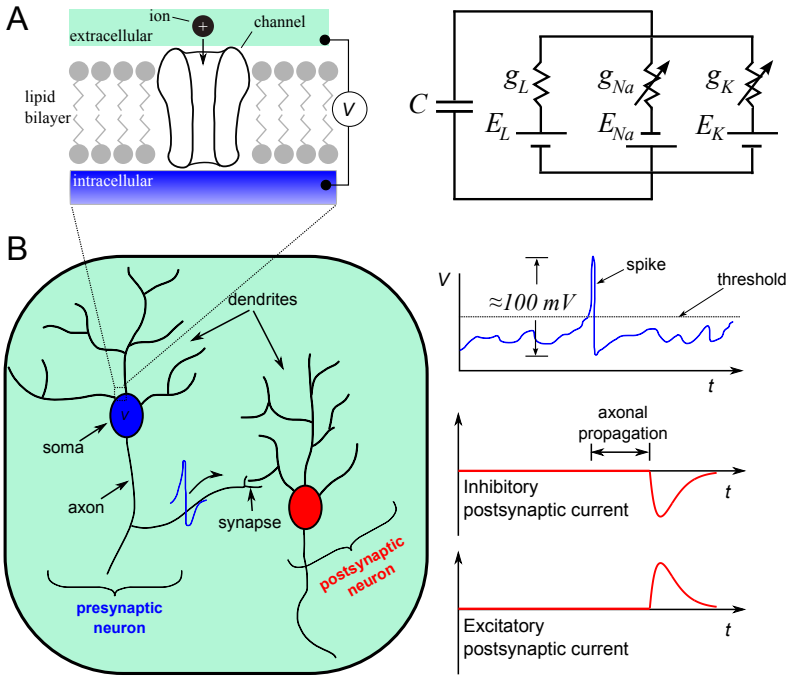

Figure 1.1: Neuronal dynamics and interaction.
(A) Left: Zoom in on a small portion of the blue neuron of panel B. A channel embedded in the lipid bilayer composing the cell membrane is depicted. These channels are selective for given ion types. For the simple preparation of a squid axon (used by Hodgkin and Huxley), only sodium, potassium and "leak" (mainly chloride) selective channels are necessary. A myriad of other channel types exist. Right: Equivalent electrical circuit of a neuron's membrane with leak (L), sodium (Na) and potassium (K) channels. Arrows across conductances mean that they depend on the value of the voltage across the membrane; these voltage-dependent channels determine the occurrence of action potentials.
(B) Schematic interaction between two neurons. When the blue neuron reaches threshold (due to excitatory inputs coming from excitatory neurons not depicted in the figure), it emits an action potential. The latter propagates along the axon which makes synaptic contact with the dendritic tree of the red cell. After a small delay induced by the axonal propagation of the action potential, a synaptic current progressively invades the postsynaptic neuron, leading to an increase (depolarization) or a decrease (hyperpolarization) of the red membrane potential with respect to the extracellular medium.

event—the emission of an action potential—is viewed as the important element in regard to information processing. That is, the shape of the action potential is thought to be irrelevant.

If neurons were isolated, they would remain at their steady state forever (with a given negative membrane potential)[1]. However, neurons are highly interconnected to form networks. An action potential emitted by a neuron travels in a solitonic manner until it reaches the end of its axon (Fig. 1.1B). The axon forms synapses (connections) with many other neurons, which are called *postsynaptic* neurons. Upon reaching the termination points of the axon, the action potential of the so-called *presynaptic* neuron triggers a series of biochemical processes. Ultimately, the membrane potential of the postsynaptic neuron changes as a result of these processes. If the postsynaptic membrane potential decreases (increases), we say that the synapse is inhibitory (excitatory). If the postsynaptic neuron receives a sufficient number of excitatory synaptic inputs in a short time span, it may reach its threshold and an action potential ensues.

Action potentials are brief events ($\sim$ 1 ms) and look like spikes in long recordings of a neuron's activity. Hence, the word *spike* is routinely used to denote an action potential and its timing. The goal of sensory systems is to acquire the physical information from the senses, to transform it into spikes, and to preprocess the resulting sequences of spikes (or *spike trains*) in order to extract low-level information about the stimulus. Further processing and information encoding and decoding may be done by other brain areas. What kind of information is encoded, how is it encoded and decoded, how efficient is the decoding and what are the dynamics of neural networks due to this processing are all questions asked in computational and theoretical neuroscience.

1.1.1 Compartmental models and complexity reduction

Modeling the dynamics of neurons can be an intricate process because of their complicated morphologies, the nonlinear dynamics of ion currents and, when embedded in networks, the high density of interconnections with other neurons. In the compartmental approach, a neuron is spatially subdivided into finite-size segments (compartments) each having its own membrane potential, channel dynamics and synaptic inputs. The voltage dynamics of each compartment are obtained by applying the Kirchhoff law of currents, taking into account the capacitive current ($c\dot{V}$), the membrane current through voltage-dependent channels, the synaptic currents coming from presynaptic inputs and the electrical coupling between adjacent compartments.

[1] Neglecting intrinsic noise sources like channel noise; see section 1.1.2.

If V_i^μ is the membrane potential of compartment μ of neuron i, then the voltage dynamics of this compartment evolve according to [70]

$$c_i^\mu \frac{dV_i^\mu}{dt} = I_{i,\mathrm{m}}^\mu(t) + \sum_{\mu'} g_{\mu,\mu'}[V_i^{\mu'}(t) - V_i^\mu(t)] + I_{i,\mathrm{syn}}^\mu(t). \tag{1.1}$$

The constant c_i^μ is the capacitance per unit area of the membrane[2] ($\sim$ 10 nF/mm^2). The sum in the second term on the right-hand side accounts for the electrical coupling with neighbouring compartments (for a non-branching dendrite, for instance, the sum includes two terms only). $I_{i,\mathrm{m}}^\mu$ is the ion current per unit area that enters the cell due to the opening and closing of voltage-gated ion channels. This current is the summation of all types of local ion-specific membrane currents, $I_{i,\mathrm{m}}^\mu = \sum_\alpha I_{i,\alpha}^\mu$. A standard Ohmic form for each contribution is

$$I_{i,\alpha}^\mu(t) = g_\alpha(t)[E_\alpha - V_i^\mu(t)], \tag{1.2}$$

where g_α is the conductance per unit area associated with channel α. It is proportional to the density of open channels, which itself depends in a complicated way on the membrane voltage. The ion-specific constant E_α is the reversal or Nernst potential; the sign of the current changes according to whether the voltage is above or below E_α.

Following the descriptive account of synaptic interaction given above, the synaptic current per unit area, $I_{i,\mathrm{syn}}^\mu(t)$, depends on the occurrence of action potentials in neurons presynaptic to neuron i. Arrival of an action potential at a synapse triggers the release into the synaptic cleft—extracellular space between the two connected neurons—of vesicles containing neurotransmitters. These chemical compounds bind to ligand-gated receptors on postsynaptic neuron i, which open to let ions flow into and out of the cell. Excitatory neurons typically produce glutamate which binds to NMDA[3] and AMPA[4] receptors, whereas inhibitory neurons mainly produce GABA[5] binding to $\mathrm{GABA_A}$ or $\mathrm{GABA_B}$ receptors. The postsynaptic current produced by a presynaptic spike depends on the neurotransmitter excreted by the presynaptic neuron (i.e., whether it is excitatory or inhibitory), on the type of ligand-gated channels present on the postsynaptic neuron's compartment and on the ion species flowing in or out of the postsynaptic cell due to the channel opening.

The total postsynaptic current (PSC) impinging on compartment μ of neuron i is the sum of the total excitatory postsynaptic current (EPSC) and of the total inhibitory postsynaptic

[2] Capacitance decreases for myelinated axons, hence the dependence upon μ.
[3] N-Methyl-D-aspartate.
[4] α-amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid.
[5] γ-aminobutyric acid.

current (IPSC). A phenomenological and simplified model of EPSC for receptor type β is

$$\sum_j \sum_{t_j^{(f)}} g_\beta(t - t_j^{(f)})[E_\beta - V_i^\mu(t)], \tag{1.3}$$

where β is either AMPA or NMDA. There are two summations: the first is on all the presynaptic neurons connected to that compartment, and the second is on all the spike times of these neurons. Often, g_β is written as the product of a maximal conductance, $\bar{g}_\beta > 0$, and a synaptic function, s_β, giving its time dependence. A general formulation for this function that is consistent with the experimentally observed rise and decay of the synaptic conductances is [83]

$$s_\beta(t) = A_\beta \left(e^{-\frac{t-\tau_d}{\tau_{\text{decay}}}} - e^{-\frac{t-\tau_d}{\tau_{\text{rise}}}} \right) \Theta(t - \tau_d), \tag{1.4}$$

where A_β is a normalization factor[6] ensuring that the maximum of s_β is 1 and $\Theta(\cdot)$ is the Heaviside function implementing causality. The time constants τ_{rise} and τ_{decay} represent respectively the rise and the decay time of the conductance of type β, whereas τ_d takes into account the propagation of the action potential along the presynaptic axon. As above, E_β is the reversal potential and depends on the ions involved. For excitatory synapses, $E_\beta \approx 0$ mV (for AMPA and NMDA). For inhibitory inputs *via* GABA_A receptors, $E_\beta \approx -70$ mV; for GABA_B, E_β is more negative than that of GABA_A. Again, it is the value of E_β that determines the sign of the contribution since $\bar{g}_\beta > 0$. Neurons spend most of their time around their resting potential, which is around -70 mV. Hence, $E_\beta - V$ is on average large and positive for excitatory inputs, almost zero for GABA_A and negative for GABA_B. IPSCs *via* GABA_A receptors are said to mediate *shunting inhibition* because their main effect is a conductance increase.

Looking at Eqs. 1.1-1.3, one can be either fascinated or baffled by the sheer complexity of neural network dynamics. For large-scale neural nets, like the ones that will be used in the present work, simulations of these equations would be extremely intensive computationally, and constraining parameter values would be difficult. Instead, we can consider simplified *point* neurons, for which a single compartment (an "effective" soma) is considered. The archetypal point neuron model is the single-compartment Hodgkin-Huxley model [106], which comprises sodium, potassium and leak currents:

$$c\frac{dV}{dt} = g_{\text{Na}}[E_{\text{Na}} - V(t)] + g_{\text{K}}[E_{\text{K}} - V(t)] + \bar{g}_{\text{L}}[E_{\text{L}} - V(t)] + I(t) \tag{1.5}$$

[6] For NMDA A_β is voltage-dependent because of the magnesium block.

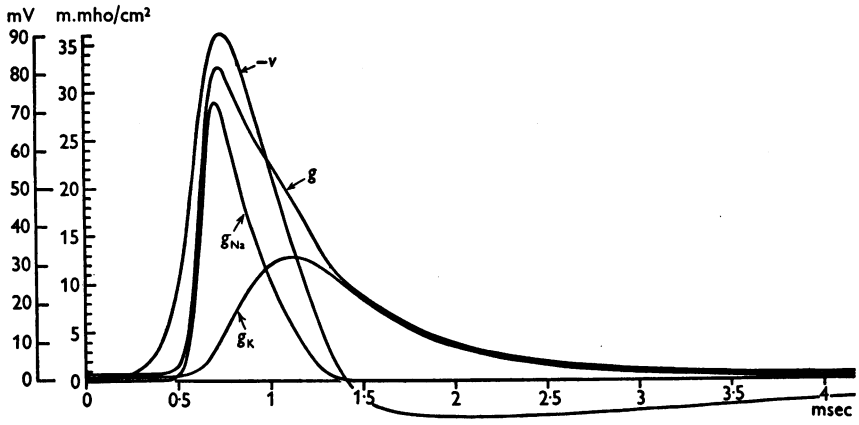

Figure 1.2: Action potential in the Hodgkin-Huxley model. "$-V$" denotes the membrane potential relative to rest. g_{Na} and g_{K} are the sodium and potassium conductances as appearing in Eq. 1.5, whereas g is the total conductance, including the leak conductance. Figure from [106] (with permission from John Wiley and Sons through RightsLink®).

with, for example[7], $E_{\mathrm{Na}} \approx 50$ mV and $E_{\mathrm{K}} \approx -77$ mV. The penultimate term is a leak current that includes contributions with voltage-independent conductances; ion pumps, which maintain the steady-state ion concentrations, may fall into that category ($E_L \approx -70$ mV, say). $I(t)$ comprises all other currents (synaptic or injected through electrodes by skilled experimentalists). The transient sodium current (first term) is responsible for the initial rise of the membrane potential during an action potential; the potassium current (second term) subsequently hyperpolarizes the membrane, and the leak current helps recover the steady state (*cf.* Fig. 1.2). Because conductances g_{Na} and g_{K} increase drastically during an action potential—increasing the total neuronal conductance—the neuron is then mostly unresponsive to the comparatively tiny conductance increases due to synaptic inputs. The result is that the shape of action potentials is fairly constant from one occurrence to the other.

The dynamics of the Hodgkin-Huxley model, although simpler than the cumbersome multi-compartmental model, are rich and somewhat complicated to analyze due to the nonlinear terms responsible for the onset of action potentials. Another simplification is to define an artificial *threshold*, so as to get rid of the action potential dynamics itself—whose invariant shape is, from the point of view of neural coding, unimportant. The leaky integrate-and-fire (LIF) model is the simplest instance of such a *spiking neuron model*. The subthreshold dynamics are given by

$$c\frac{dV}{dt} = -\bar{g}_{\mathrm{L}}[V(t) - E_{\mathrm{L}}] + I(t), \tag{1.6}$$

[7] Nernst potentials depend on temperature and the ion concentrations inside and outside the neuron.

and whenever a threshold V_T is reached, V is reset to a value $V_R \sim E_L$. Then, for a short time interval corresponding to the time course of the action potential, the dynamics are momentarily frozen at $V = V_R$. This model can be used to simulate very large networks due to its simplicity; however, it may not be very realistic biologically (but see [175]). Often, Eq. 1.6 is divided throughout by $\bar{g}_\mathrm{L}$ to make the *membrane time constant* $\tau_m = c/\bar{g}_L$ appear on the left-hand side. This time constant dictates how fast the membrane potential relaxes towards its steady state.

1.1.2 Neuronal noise and variability

The issue of *variability* is important in neuroscience, and in science in general. Variability means that repeating the same experiment under the same conditions leads to different responses. Variability is related to the issue of noise, i.e. unpredictable and usually unwanted signals that perturb the activity of a given system. However, variability is not synonymous with noise; deterministic sources of variability exist, with chaos being the most obvious one. Linking the macroscopic variability in behavioral responses with microscopic sources of variability is a big challenge. From molecular, to synaptic, to cellular, to network level, several spatial and temporal scales are swept and variability at all these levels must be reconciled and connected [84].

For neurons, variability means that action potentials' timings will change from one presentation of a stimulation to the other [128]. This variability suggests the presence of noise in the underlying voltage dynamics. A typical consequence of noise is irregular—i.e., not periodic—spiking activity. Given that spike times are the meaningful quantities for neural information processing [179], uncovering the sources of this variability is crucial.

A first cause of variability[8] is *channel noise*, which is mainly a finite-size effect. A finite number of ion channels resides in the neuronal membrane, and the stochastic openings and closings of these channels are a source of noise for the neuron. For N independent channels, current fluctuations scale like $1/\sqrt{N}$ [222]. Also, the conditions under which experiments are performed determine the possible noise sources. Under *in vivo*[9] conditions, cortical pyramidal neurons are bombarded by thousands of synaptic inputs that drive their membrane potentials up and down, similar to shot noise in electronic circuits. The result is an irregular spiking activity [197]. Reduced variability occurs *in vitro* because the

[8] In this discussion, we leave aside thermal (Johnson) noise—coming from the finite temperature of neurons and typically negligible [138]—and more exotic noise like ephaptic coupling [7]—nonsynaptic coupling between a neuron and the surrounding electromagnetic field.

[9] *In vivo* studies use whole, living organisms. In contrast, *in vitro* experiments are conducted on a portion of the system under study, separated from the organism (for example, a slice of cortical tissue).

number of synaptic inputs severely decreases. *In vitro*, variability also depends on the type of stimulation: rapidly varying suprathreshold inputs produce very reliable spiking, precise to a submillisecond timescale [131]. It must be added that being intensely driven by synaptic inputs does not necessarily lead to irregular spiking [197]. The average value of the total synaptic current (excitatory and inhibitory inputs) plays the role of a drift term that, when positive, drives the membrane potential towards threshold. When the drift is zero on average (balanced condition) or comparatively low with respect to threshold, current fluctuations cause the spiking, hence the spiking variability [197, 185]. Finally, the type of the neuron is important. So-called stuttering and irregular spiking cortical interneurons are irregular *in vitro*, and may themselves induce variability in their postsynaptic targets [202].

Various approaches have been taken to model noise in neurons. One approach is to add a *diffusive* noise term on the right-hand side of, say, Eq. 1.6. When this diffusive term is a Gaussian white noise, Eq. 1.6 corresponds to a Ornstein-Ulhenbeck process [215] with an absorbing boundary (due to the threshold). This model will be analyzed in detail in the next chapter. Gaussian white noise is considered "fast" noise because its power spectrum is flat even for high frequencies. A second approach is to use *escape* noise [91]. In this case, explicit noise is not added to the membrane potential dynamics. Instead, the temporal evolution of V is deterministic, and noise takes the form of a firing probability that depends on $V(t)$. This noise is also considered "fast" because the firing probability changes instantaneously according to the membrane potential dynamics. An example of "slow" noise is to change a parameter (e.g., the reset value V_R) after each spiking event according to a given probability distribution [113]. This noise model may not be very realistic. A more biologically sound slow noise is a Gaussian white noise low-pass filtered by synapses [41]. This can be implemented by synaptically connecting a multitude of noisy Poisson neurons (see chapter 2) to the model neuron; in this case, noise is equivalent to shot noise. Noise can also be self-generated by the network embedding the neuron when the connectivity of the network is low (sparse network) [40].

1.1.3 Firing rate

Among all the possible quantities characterizing a neural response, the firing rate is certainly among the most important. A spike train emitted by a given neuron is a realization of an underlying stochastic *point process* $y(t)$. That is, during an experiment, we record the spike times of a neuron, which are point events on the timeline. Repeating this experiment under the same conditions will yield a different spike train than the first (owing to neural variability); each trial gives a different realization of the process. As with any stochastic

process, various statistics can be computed: expectation values, autocorrelations, etc. With point processes, auxiliary processes are often constructed as well, e.g. counting processes (counting the number of events in interval $[0, t)$), or interarrival processes (the differences of timing between consecutive events). Various definitions of the firing rate can be conceived which are related to these primary and auxiliary processes [70, 91]. In the first definition, we compute the expectation value of the counting process and divide by time; in the second definition, we find the average of the point process at each instant. The third definition is related to the second, but instead of computing an ensemble average, the average is performed on a neural population.[10] Finally, a fourth definition exists that involves computing the firing rate as the inverse of the average interarrival time. To simplify, we assume that the stimulus to which the neuron responds can be considered a deterministic function of time.

First definition This definition assumes that the stimulation to which the neuron responds is constant during the recording. If the stimulus varies with time, then another approach is needed. Let $N(t)$ be the counting process associated with the stochastic spiking process $y(t)$. For every $t > 0$, $N(t)$ is a random variable taking values in $\{0, 1, \ldots\}$, the possible spike counts in interval $[0, t)$. The firing rate for the interval $[0, T)$ is defined by[11]

$$r = \frac{\langle N(T) \rangle}{T}. \tag{1.7}$$

Here and in the following, the brackets denote an ensemble average over the sample functions representing the stochastic process. That is, if N_k is the number of spikes in trial (realization) k and if every trial lasts the same amount of time, then the firing rate is

$$r = \lim_{M \to \infty} \frac{1}{M} \sum_{k=1}^{M} \frac{N_k(T)}{T}. \tag{1.8}$$

In experiments (real or numerical), only a finite number of realizations can be carried out. Thus, we have to rely on a sample average of $N(T)$ to estimate the firing rate:

$$\check{r} = \frac{1}{M} \sum_{k=1}^{M} \frac{N_k(T)}{T}, \tag{1.9}$$

where M is the number of trials and the inverted hat signifies that this is a sample statistic.

Second definition We assume now that the stimulation varies as a function of time. Ex-

[10] Assuming that the neurons have similar properties.

[11] In general, the interval could be $[t_1, t_1 + T)$, where t_1 is any time. Then the firing rate would be $r = \langle N(t_1 + T) - N(t_1) \rangle / T$.

perimentally, we record the spiking activity of a neuron, and repeat the experiment (or simulation) several times. Spike times are monitored with finite time resolution. Therefore, the trial length is divided into equal bins and we determine to which bin each spike time belongs; the result is a binary event train of zeros and ones for each trial if the bin width is small enough. We then compute the average over trials for each bin, and divide the result by the bin width Δt, i.e.,

$$\check{r}(t_i) = \frac{1}{M} \sum_{k=1}^{M} \frac{n_k(t_i)}{\Delta t}, \tag{1.10}$$

where $n_k(t_i) \in \{0, 1\}$ is the number of spikes in bin i for trial k and $t_i = i\Delta t$. In the limit $\Delta t \to 0$ we can represent each sample spike train by a sum of Dirac delta functions centered at the spike times. Denoting the sequence of spike times by the ensemble $\{t_k^{(f)}\}_{f=1,2,\ldots}$, where k indexes the trial and f orders the spikes, we can write for the spike train of trial k

$$y_k(t) = \sum_f \delta(t - t_k^{(f)}). \tag{1.11}$$

This representation of a spike train can be justified by the fact that the derivative of the associated counting process $N(t)$—which varies by a unit step at each spike time—yields Dirac delta functions. With $y(t)$ the parent spiking process of which $y_k(t)$ is a sample function, we could write symbolically

$$y(t) = \frac{dN(t)}{dt}. \tag{1.12}$$

The rate at time t is the expectation value of $y(t)$ (which corresponds to taking the limits $\Delta t \to 0$ and $M \to \infty$ in Eq. 1.10):

$$r(t) = \langle y(t) \rangle. \tag{1.13}$$

This ensemble average of Dirac delta functions is well-defined because

$$\langle y(t) \rangle = \frac{d}{dt} \langle N(t) \rangle,$$

and $\langle N(t) \rangle$ is assumed to be a continuous and differentiable function [157].

Third definition Consider an ensemble of M identical neurons. We record their spike re-

sponses to a time-varying stimulation and then compute the *population activity*, i.e.,

$$\check{R}_M(t_i) = \frac{1}{M}\sum_{j=1}^{M}\frac{n_j(t_i)}{\Delta t},$$

where $n_j(t_i)$ is 1 if *neuron j* fired in bin i and 0 if not. This third definition is formally equivalent to the second: instead of repeating the experiment to yield an estimate of the rate, we perform a single experiment but average over several neurons. $\check{R}_M(t)$ is a sample function of parent process $R_M(t)$; other sample functions are obtained by stimulating again the neural assembly. The parent process can be expressed in terms of the spiking process of each neuron:

$$R_M(t) = \frac{1}{M}\sum_{j=1}^{M} y_j(t). \tag{1.14}$$

The expectation value of this process is trivially $\langle R_M(t)\rangle = r(t)$, recovering the second definition. Alternatively, the average of $R_M(t)$ can be computed by considering the limit when the number of neurons becomes infinite, i.e.,

$$\langle R_M(t)\rangle = \lim_{M\to\infty}\check{R}_M(t). \tag{1.15}$$

Fourth definition At last, in theory, we can compute the rate for a constant stimulation by using the probability density function of interarrivals—called *interspike intervals (ISIs)* in the neural context—i.e., $P_{\text{ISI}}(x)$. An ISI is just the time difference between two successive spikes. The inverse of the average ISI is the rate:

$$r = \left(\int_0^\infty x P_{\text{ISI}}(x)dx\right)^{-1}. \tag{1.16}$$

It is possible to infer the meaning of the firing rate from the definitions above. The average number of spikes in bin $(t-\Delta t/2, t+\Delta t/2)$ is

$$\bar{n}(t) = \sum_{k=1}^{\infty} k\, P[n(t)=k],$$

where $P[n(t)=k]$ is the probability of getting k spikes in that bin. If Δt is small, $P[n(t)=k]$ becomes negligible for $k>1$ and $\bar{n}(t)\approx P[n(t)=1]$. But the average number of spikes in

$(t - \Delta t/2, t + \Delta t/2)$ can also be obtained from

$$\bar{n}(t) = \int_{t-\frac{\Delta t}{2}}^{t+\frac{\Delta t}{2}} r(t')dt' \approx r(t)\Delta t. \tag{1.17}$$

Hence,

$$r(t) \approx \frac{P[n(t) = 1]}{\Delta t}, \tag{1.18}$$

and the firing rate is the "probability density" of firing a spike in a small time interval.[12] In the next section, using the first and second definitions, we will present the concept of gain control, i.e. a transformation of the firing rate versus input curve (f-I curve).

1.2 Gain control

Neurons are sometimes viewed as input-output transformation devices. They integrate postsynaptic currents caused by presynaptic neurons' firing and nonlinearly produce action potentials which, in turn, propagate and affect other neurons. Given that this causal chain of interactions can be traced back to sensory receptors, neurons ultimately respond to stimuli of the physical realm.

As discussed above, a first-order characterization of neural outputs consists in measuring the firing rate. For the "input", one actually has a myriad of possibilities; "input" must here be interpreted as a portmanteau word to mean any parameter on which the neuron's output firing rate depends. Some examples are: a direct intracellular injection of current in an *in vitro* slice preparation, the total synaptic current received by a neuron embedded in its network, or the intensity of the external (physical) stimulus to which the neuron is responsive. Hence, the exact nature of the "input" depends on the context. Computing a neuron's firing rate f as

[12] Note that the firing rate cannot be considered a probability density *per se* since it is not normalized to 1. Hence, the phrase "in a small time interval" is mandatory. In that respect, the firing rate is more related to the conditional intensity or hazard function in the theory of point processes [210, 69]. The hazard function $\rho[t|\mathscr{H}(t), \Theta]$ is conditional on the history $\mathscr{H}(t)$ of the neuron's spiking until time t—i.e., the sequence of spike times occurring before t—and on other covariates Θ; it is defined by

$$\rho[t|\mathscr{H}(t), \Theta] = \lim_{\Delta t \to 0} \frac{P[N(t + \Delta t) - N(t) = 1|\mathscr{H}(t), \Theta]}{\Delta t}, \tag{1.19}$$

where $N(t)$ is here the spike count in interval $[0, t)$. The hazard function fully characterizes the neural point process [69]. For small Δt, $\rho[t|\mathscr{H}(t), \Theta]\Delta t$ is the neuron's firing probability for interval $[t, t + \Delta t)$. The difference between ρ and r lies in the fact that ρ depends on the precise spiking history ($\mathscr{H}$) and stimulation protocol (included in the covariates Θ), whereas r is an average over the spiking histories for the given stimulation protocol [150]. Rigorously, we should write $r(t|\Theta)$, so that $r(t|\Theta) = \langle \rho[t|\mathscr{H}(t), \Theta] \rangle_{\mathscr{H}(t)}$, where $\langle \cdot \rangle_{\mathscr{H}(t)}$ denotes the average over histories.

a function of stimulation intensity I gives the *f-I curve*. Plotting the firing rate as a function of a stimulation parameter (e.g., the orientation of a bar in the visual field) constitutes a *tuning curve*. Both of these curves can be considered input-output transformations. With these curves, we can characterize the responsiveness, selectivity and sensitivity of a neuron, and describe the type of transformation performed on its input.

A neuron is usually responsive to a fraction only of the stimulus properties. For instance, retinal ganglion cells respond to visual stimulation (light) applied to only a fraction of the visual field [70]. This fraction is called the receptive field of the neuron. A neuron can also be selective for a specific value of a parameter defining the stimulus. For example, simple cells in the primary visual cortex are selective for the orientation of a light bar appearing on an otherwise black screen [110]. As a result, the tuning curve rate versus orientation will show a maximum for this preferred orientation. The sensitivity of a neuron, on the other hand, is related to how much it responds to a change in its input. This corresponds to the derivative of the f-I or tuning curve, or *gain*. Often, however, gain means the f-I or tuning curves themselves. *Gain control* (or gain modulation) consists in transformations of f-I and tuning curves by a modulating factor in order to accomplish computational goals. Gain control is thus a mean to integrate two input streams, the input and the modulating factor [186, 184].

Various gain control mechanisms have been described in the literature. Gain control can be performed by neuromodulators [74, 226, 165, 204], input noise [78, 58, 11, 103, 85, 108], attention and other brain states [64, 65, 165], neural feedback [183, 203], synaptic plasticity [1, 182], dendritic mechanisms [177, 169, 158] and adaptation [25], to mention just a few. Also, the specific architecture of a network may itself promotes efficient gain control [203, 144]. This has been shown in particular for the electrosensory system of weakly electric fish, as will be seen in chapters 2 (section 2.8.3) and 4. In this section, we first delineate different types of gain control, and we discuss a few classic mechanisms that underly them. Then, we discuss the importance of gain control for neural network dynamics and information processing.

1.2.1 Types of gain control

From the mathematical point of view, gain control is any functional transformation $\mathcal{T}$ of a f-I curve's derivative. If $f_0(I)$ is the f-I curve before any of the modifications discussed above, gain control can be represented by

$$f'(I) = \mathcal{T}\{f_0'(I)\}, \tag{1.20}$$

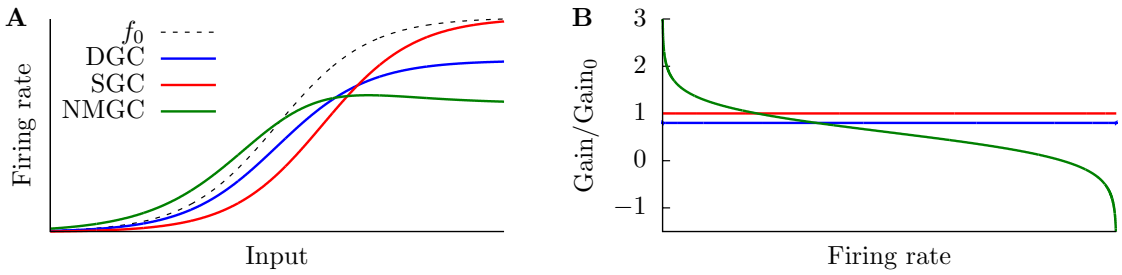

Figure 1.3: Types of gain control. (A) Subtractive (red), divisive (blue) and non-monotonic (green) gain controls. The dotted black curve is the reference f-I curve, $f_0(I)$, here given by a sigmoid. (B) The ratio $f'(I)/f_0'(I)$ as a function of rate. For the subtractive case (red), notice that this ratio is 1 for all rates.

where $f(I)$ is the f-I curve after modification. An analogous definition for tuning curves can be formulated. Two types of gain control are well-known and have been studied extensively: the *subtractive* (SGC) and *divisive* (DGC) gain controls. A third gain transformation, dubbed non-monotonic, will be studied in chapter 4. A brief description of these gain controls follows.

Subtractive gain control

Subtractive (or additive) gain control corresponds to a right (left) shift along the I axis of the f-I curve (red curve in Fig. 1.3A). As per Chance *et al.* [58], this control should not be considered an actual gain control because the slope is merely shifted. However, this shift does change the current-to-gain relationship. For SGC, $f(I) = f_0(I - \text{offset})$, and thus

$$f'(I) = f_0'(I - \text{offset}). \tag{1.21}$$

Perhaps the most well-known mechanism for SGC is shunting inhibition [107]. Shunting inhibition is caused by a synaptic current with a reversal potential that lies around the resting potential of the neuron; therefore, its main effect is to increase the input conductance of the cell. This is readily seen by considering a simple LIF neuron with a constant inhibitory conductance g_I and driven by a constant input I (Eq. 1.6):

$$\tau_m \frac{dV}{dt} = E_L - V(t) + \frac{g_I}{\bar{g}_L}[E_I - V(t)] + \frac{I}{\bar{g}_L}. \tag{1.22}$$

For shunting inhibition $E_I \approx E_L$ so that

$$\frac{\tau_m}{1+\frac{g_I}{\bar{g}_L}}\frac{dV}{dt} = E_L - V(t) + \frac{I}{\bar{g}_L + g_I}. \tag{1.23}$$

Therefore, shunting inhibition decreases the membrane time constant and divisively reduces the input current. A LIF neuron spikes when its membrane potential reaches the threshold V_T and is then reset to V_R for a refractory period τ_r. This refractory period mimics the duration of the action potential in a real neuron. From the equation above the steady state potential without the spiking mechanism is $V_{\text{ss}} = E_L + \frac{I}{\bar{g}_L+g_I}$. If $V_{\text{ss}} < V_T$ then $\lim_{t\to\infty} V(t) = V_{\text{ss}}$ and the neuron never reaches threshold (subthreshold regime); when $V_{\text{ss}} > V_T$ the neuron fires periodically (suprathreshold regime). The firing rate can be obtained by computing the time T_1 that it takes to reach threshold V_T after a spike and after the absolute refractory period is over. Assuming that the last spike occurred at $t = 0$ without loss of generality, we can solve the equation above with $V(\tau_r) = V_R$ to give

$$V(t) = V_{\text{ss}} + (V_R - V_{\text{ss}}) \exp\left(-\frac{t-\tau_r}{\tau'_m}\right), \tag{1.24}$$

with $\tau'_m \equiv \tau_m/(1+\frac{g_I}{\bar{g}_L})$. At $t = 0$, the neuron has just spiked; at $\tau_r + T_1$, the neuron spikes again and we set $V(T_1 + \tau_r) = V_T$. The rate is just the inverse of $\tau_r + T_1$:[13]

$$r(I) = (\tau_r + T_1)^{-1} = \left[\tau_r + \tau'_m \ln\left(\frac{V_{\text{ss}} - V_R}{V_{\text{ss}} - V_T}\right)\right]^{-1} \Theta\left(V_{\text{ss}} - V_T\right). \tag{1.25}$$

We compare this f-I curve with the f-I curve without shunting inhibition, i.e. with $g_I \equiv 0$:

$$r_0(I) = \left[\tau_r + \tau_m \ln\left(\frac{I - \bar{g}_L(V_R - E_L)}{I - \bar{g}_L(V_T - E_L)}\right)\right]^{-1} \Theta\left(I - \bar{g}_L(V_T - E_L)\right). \tag{1.26}$$

With shunting inhibition the neuron starts to fire when $V_{\text{ss}} > V_T$, i.e., when $I > (\bar{g}_L + g_I)(V_T - E_L)$; without inhibition, it fires whenever $I > \bar{g}_L(V_T - E_L)$. Hence, there is a clear shift of the firing onset. For large I, and if the refractory period is small, we can Taylor expand the logarithm to get

$$r_0(I) \approx \frac{I}{\tau_m(V_T - V_R)} - \frac{V_T}{\tau_m(V_T - V_R)} \equiv r_0^{\text{approx}}(I) \tag{1.27}$$

[13] This can be viewed as an application of the fourth definition of the rate, Eq. 1.16, with $P_{\text{ISI}}(s) = \delta(s - \tau_r - T_1)$.

and

$$r(I) \approx r_0^{\text{approx}}(I) - \frac{g_I}{g_L \tau_m} \frac{V_T - E_L}{V_T - V_R}. \tag{1.28}$$

Thus, a subtractive gain control occurs for large I. For intermediate values of I (i.e., not too large and not around the firing onsets), the gain control is not purely subtractive.

Holt and Koch [107] were the first to point out that shunting inhibition does not lead to divisive gain control, as one would have expected from its effect on the membrane potential. Hyperpolarizing inhibition, i.e. inhibition for which the reversal potential lies below the resting potential of the neuron, also produces subtractive gain control [58]. This gain control mechanism will be at play in chapter 4. It will be seen that feedforward inhibition—i.e., a neural population inhibits another population further away in the sensory pathway—, leads to SGC when neuronal noise is low.

Divisive gain control

Divisive (or multiplicative) gain control consists in a division (or multiplication) of the gain by a constant factor; it is a simple scaling of the f-I curve's slope:

$$f'(I) = \alpha f_0'(I); \tag{1.29}$$

where α is a constant (see blue curve in Fig. 1.3A). A classic mechanism for DGC is an alliance between shunting inhibition and noisy synaptic inputs [58].

Neurons are bombarded by synaptic inputs. The resulting conductance fluctuations yield membrane potential fluctuations. In the subthreshold regime, a noiseless neuron cannot reach threshold and thus cannot spike. With noise, impromptu fluctuations may throw the potential towards threshold; large enough upswings may lead to threshold crossings. In a so-called *balanced* state, on average the neuron receives as much inhibition as excitation, i.e. it is neither driven towards threshold, nor away from it. Then fluctuations are solely responsible for spike firing if no external input is applied. Chance and co-workers [58] have shown that a combination of shunting inhibition and noise-induced voltage fluctuations in a balanced state provide the sufficient ingredients to get divisive gain control. They considered a cortical neuron subjected to a barrage of excitatory and inhibitory inputs arriving at given rates, and they measured the neuron's firing rate in response to a constant driving current I. A common multiplicative increase of the rates of excitation and inhibition, causing an augmentation of the noise and a subtractive shift, yields divisive gain control. Figure 1.4 illustrates in a schematic manner how this happens. Experiments with slices of cerebellar granule cells [148] and modeling work with more sophisticated models [149, 169, 78] have

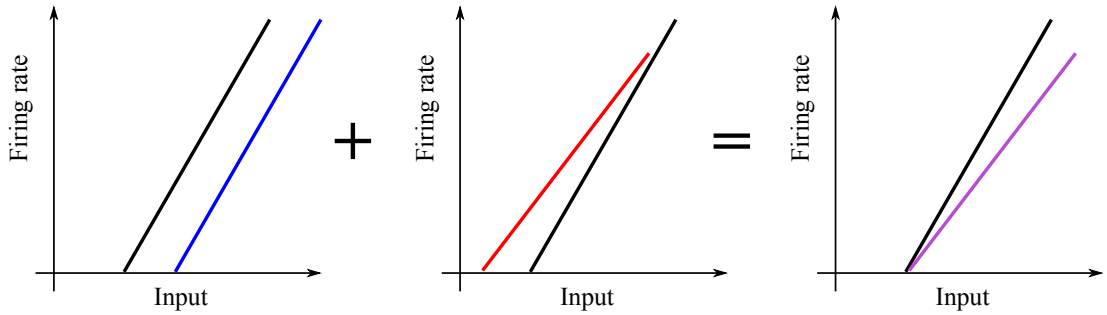

Figure 1.4: Constructing divisive gain control with inhibition and noise in a balanced state. Inhibition (left) induces a subtractive shift of the original f-I curve (black curve). Noise from balanced synaptic inputs increases the firing rate around the onset of firing, i.e. the point along the "Input" axis where the rate becomes nonzero (middle). Put together, these two effects yield a divisive gain control. Schematic figure inspired by [171].

all emphasized shunting inhibition and noise as two important components of divisive gain control. In chapter 4, feedforward inhibition and high neuronal noise will be shown to be an effective way to get DGC in sensory systems.

Non-monotonic

Non-monotonic gain control (NMGC) is related to DGC, but the end product is a f-I curve displaying a local maximum (see Fig. 1.3A, green curve). If f_0 is monotonously increasing with I, NMGC can be represented by

$$f'(I) = \alpha(I) f_0'(I) \tag{1.30}$$

where $\alpha(I)$ changes sign somewhere along the I axis. Mechanisms for NMGC are difficult to find in the literature. Correlated inputs associated with short-term plasticity have been implicated [72]. In chapter 4 we show that strong feedforward inhibition can non-monotonously control the gain of a postsynaptic population.

These gain controls are all depicted in Fig. 1.3A using a sigmoid function as the reference f-I curve. Firing rates often take this sigmoidal shape as a function of the input intensity because they typically plateau as the input gets larger (saturation), and they cannot be negative even for negative inputs (rectification).

A way to decisively determine the type of a given gain control is to plot the ratio between the gain of the transformed f-I curve and that of the original f-I curve as a function of the rate [58]. From Eqs. 1.21-1.30, the resulting curves will give a constant equal to 1 for SGC, a

constant equal to α for DGC and a curve crossing zero for NMGC (see Fig. 1.3B). Pure SGC or DGC may not be possible on the whole dynamic range of a neuron, i.e. the set of all its possible firing rates, as seen above.

1.2.2 The importance of gain control

Gain control is perhaps one the very few computational principles that is general enough so as to apply to all brain regions and sensory modalities [186, 184]. Here, we briefly discuss its functions with emphasis on network stability and information processing.

Stability and network dynamics

In the context of formal neural networks dynamics, f-I curves are called activation functions [137], and only firing rate dynamics—not spiking dynamics—matter. In this framework, the total input received by neuron i within a network of N neurons is given by

$$h_i(t) = \sum_{j=1}^{N} W_{ij} r_j(t) + I_i(t), \tag{1.31}$$

where r_j is the rate of neuron j and $I_i(t)$ is any external input coming from outside the network *per se*. Here, the firing rate comes from the second definition above, Eq. 1.13. The W_{ij}'s are matrix elements of the so-called connectivity matrix $\mathbf{W}$, describing how neurons are interconnected (see chapter 2, Fig. 2.1C). In the simplest case, firing rates change according to the following dynamics:

$$\tau \frac{dr_i}{dt} = -r_i(t) + f[h_i(t)], \tag{1.32}$$

with $f(\cdot)$ the f-I curve; τ is a time constant setting the scale of rate dynamics. Hence, the f-I curve is directly involved in the network dynamics. All the machinery of nonlinear dynamics can then be applied to the above equations to determine the possible dynamical states; in particular, analogous equations will be studied in section 1.3 in the context of neural oscillations. Therefore, factors modifying the shape of f-I curves have a direct impact on neural dynamics.

Information processing

Neural encoding pertains to how neurons and neural networks represent the parameters of an input, e.g. its mean value and standard deviation. f-I and tuning curves provide information about the average response r of a neuron as a function of the parameter. From

trial to trial neural responses deviate from this average. A conditional probability distribution $p(x|s)$ can be constructed such that $r(s) = \int_0^\infty x p(x|s) dx$; $p(x|s)$ is the probability that the rate be x given that the parameter value is s. This probability distribution characterizes the encoding performed by the neuron; in the case of a neural population, x stands for the firing rates of all neurons. From $p(x|s)$ we can produce an estimate s_{est} of the real parameter value s. Various methods exist to achieve this goal, namely the maximum likelihood, the maximum a posteriori inference or the Bayesian inference [70]. Observing rate x on a given trial (or the set $\{x_i\}_{i=1}^N$ for a population of N neurons) gives an estimate for s. When the trial average of these estimates, denoted by $\langle s_{\text{est}} \rangle$, is such that $\langle s_{\text{est}} \rangle = s$, then the estimate is said to be unbiased. The squared estimation error for an unbiased estimate is $\sigma^2_{\text{est}} \equiv \langle (s_{\text{est}} - \langle s_{\text{est}} \rangle)^2 \rangle$. It is limited by the Cramer-Rao bound $\sigma^2_{\text{est}} \geq [I_F(s)]^{-1}$, where $I_F(s)$ is the *Fisher information* of distribution $p(x|s)$, i.e.,

$$I_F(s) = \int_0^\infty dx \; p(x|s) \left(-\frac{\partial^2 \ln p(x|s)}{\partial s^2} \right). \tag{1.33}$$

It is the mean curvature of the logarithm of $1/p(x|s)$; a higher average curvature means a higher Fisher information, which in turn signifies a lower Cramer-Rao bound and a lower squared estimation error.

Consider a population of independent neurons encoding parameter s during time T with conditional probability distributions given by Poisson distributions

$$p(x_i|s) = \frac{[r_i(s)T]^{x_i T}}{(x_i T)!} \exp\left[-r_i(s)T\right], \tag{1.34}$$

where $x_i T$ is the observed spike count for neuron i during time T and $r_i(s)T$ is the average spike count for neuron i when the stimulus has value s. The Fisher information is [70]

$$I_F(s) = T \sum_{i=1}^N \frac{[r_i'(s)]^2}{r_i(s)}. \tag{1.35}$$

Hence, the f-I curves and the gains appear directly in the Fisher information. Each neuron contributes an amount $Tr_i'(s))^2/r_i(s)$ to the total I_F. Interestingly, this information is zero when the gain is zero. This result generalizes to long recordings of renewal spiking processes (see chapter 2, section 2.5.1) for which consecutive ISIs are statistically independent [201]:

$$I_F(s) = \sum_{i=1}^N \frac{[r_i'(s)]^2}{\sigma_i^2(s)}, \tag{1.36}$$

where σ_i^2 is the variance of the rates.

1.3 Rhythms and oscillations

Through synaptic interactions neural networks can display coherent collective dynamical states. Dynamical state here means the set of concomitant spike trains produced by the neural population in response to stimulation.[14] A prominent example is neural oscillations or rhythms, which involve repetitive neuronal discharges and some degree of synchronization among neurons. Various mechanistic explanations have been offered over time to understand their origin, and various putative functions have been ascribed to them. The study of rhythms in health and disease has a rich history and the goal of the present section is definitely not to present a complete picture.[15] The goal is rather to present—sometimes mathematically, other times in a more descriptive manner—a few important ingredients for the emergence of rhythms in neural assemblies. These notions will be useful in chapters 3 and 5, as well as in section 2.8.3 of chapter 2, where oscillations in response to spatially coherent stimulations will be studied.

In this section, we first present a short phenomenological description of neural rhythms, focusing on the so-called gamma rhythms. We then discuss how the interplay between neuronal, synaptic and connectivity properties give rise to oscillations in neural networks.[16] We first use a simple rate model to build intuition about how oscillatory instabilities are generated in homogeneous and fully recurrent networks of excitatory or inhibitory neurons. Then, we discuss the impact of noisy spiking on these single-population rhythms, followed by a discussion of the effects of a sparse connectivity between neurons. In the last subsection, we consider networks of excitatory and inhibitory cells, and finish with a discussion of the ING versus PING mechanisms for gamma-band rhythms generation.

1.3.1 Brief phenomenology

A myriad of electrophysiological techniques exist to probe neural activity at various spatial and temporal resolutions. As early as 1929, Berger [27] used electroencephalography (EEG) to record brain activity at a macroscopic level. EEG measures electric potentials at the surface of the scalp [46]. He discovered, in the occipital lobe region (related to the visual system), two types of oscillatory activity in the voltage trace: alpha waves, with frequency ~ 10 Hz occurring when the brain is in a relaxed state and the patient has his/her eyes closed, and beta rhythms, of higher frequency, when the brain is alert and eyes are open. With time other

[14] Sometimes, the dynamical state can also be spontaneous, i.e., not causally linked to any stimulation.
[15] A complete book by G. Buzsáki [45] and a vast review by X.J. Wang [219] accomplish this task successfully.
[16] We exclude instances where rhythmic activity is caused by pacemaker-like activity.

frequency bands were noticed. The amplitude of the recorded rhythm is a power law of the frequency for frequencies above 80 Hz [146]. Another important frequency band is the theta band, between 4 and 7 Hz. These rhythms predominantly occur in the hippocampus and are involved in position coding by so-called place cells [152]. Delta waves (0.5-4 Hz) are found during non-REM sleep. Slower rhythms can also be found. At the other end of the spectrum, ultra fast $>$ 100 Hz rhythms (ripples) also exist [23]. Causal explanations and behavioral correlates for all these rhythms have been studied extensively, but a coherent picture is still lacking.

Perhaps the most ubiquitous rhythms are those in the gamma-band, between 30-80 Hz, but typically around 40 Hz. The importance of gamma rhythms is due to their putative roles in various brain functionalities, most importantly the *binding by synchrony* hypothesis, which states that different aspects of a percept are put together (bound) by the synchronous oscillations of separate cell assemblies that encode these various aspects [96, 97]. Gamma rhythms have been found in rodents, monkeys, humans but also in low-level creatures like weakly electric fish (see section 2.8.3 of chapter 2 and chapter 3). In mammals, the hippocampus, the olfactory bulb and many cortical areas have been shown to display these oscillations, both *in vitro* [223, 207] and *in vivo* [48, 161, 198, 52], and for a variety of behavioral states (sleep, anesthesia, awakeness). Stimulus-induced gamma oscillations in sensory systems—especially the visual system—have been repeatedly linked to the coherence of the stimulation [96, 117]. Coherence can be loosely defined as the amount of correlation between inputs to the neurons. For instance, gamma rhythms in the cat primary visual cortex were found to decrease in intensity and synchrony when the grating stimulus used in the experiment was polluted by noise [228]. Another aspect of gamma rhythms is their fast temporal decorrelation [68, 117], a characteristic that will be studied in chapter 5. Fast inhibition has been shown to be critical to their generation, as reviewed in [49]. As such, two competing scenarios for gamma oscillation generation exist: the interneuron network gamma (ING) and the pyramidal-interneuron gamma (PING) mechanisms. In the former, a recurrent network of inhibitory neurons is sufficient to generate the rhythm; in the latter, interactions between excitatory pyramidal neurons (having the cortex in mind) and local interneurons generate the oscillations. A third possibility, explored in chapter 5, involves feedforward rather than recurrent nets and produces short-lived oscillatory-like activity.

1.3.2 Networks of inhibitory or excitatory neurons

We begin by studying the occurrence of oscillations in either purely excitatory or purely inhibitory networks. Fully connected (or recurrent) networks are considered first, i.e., networks

whose constituent neurons are connected in an all-to-all fashion. Next, we consider sparse networks, for which the connection probability is low. In all cases we deal with homogenous networks, i.e. networks for which all neuronal and synaptic properties are identical for all neurons and synaptic connections. Our goal here is to illustrate 1) how the *asynchronous state* of these networks loses stability and bifurcates towards an oscillatory state, and 2) how this instability onset depends on noise and synaptic time constants. In the asynchronous state the neurons fire independently and the average population activity is constant. In the next subsection, we use a simple population dynamics *à la* Wilson-Cowan [224] to show how eigenfrequencies are found in delayed network dynamics.

Axonal delay and synaptic time constant

Various neuronal time constants have an effect on the occurrence and properties of oscillations. Among these, the axonal propagation delay of the action potential and the time scale of the synaptic interaction are especially important. To understand their influence, we consider a homogenous network of inhibitory or excitatory neurons.[17] Instead of describing the activity (firing rate) of every neuron, we consider the overall activity of the population. If $R(t)$ is the population activity (*cf.* section 1.1.3, third definition), f the f-I curve of all neurons and $I_s(t)$ the total synaptic current impinging on all neurons, a simplified population dynamics reads

$$\begin{aligned} \tau_s \frac{dI_s}{dt} + I_s(t) &= JR(t-\tau_d) \\ R(t) &= f[I_s(t) + I_{\text{ext}}]. \end{aligned} \tag{1.37}$$

These equations can be compared with Eqs. 1.31-1.32. Here, the time constant setting the scale of the rate dynamics is assumed small compared to the synaptic time constant τ_s: the population's response to a change of current is instantaneous. τ_d is the delay, J is the synaptic strength and I_{ext} is a constant external input (or bias). Changing the time variable to $\tau = t/\tau_s \to t$ gives for the first equation

$$\frac{dI_s}{dt} + I_s = JR(t-\tau_d),$$

where t and τ_d are now measured in units of τ_s. The fixed point I_s^* obeys

$$Jf(I_s^* + I_{\text{ext}}) = I_s^*. \tag{1.38}$$

To solve this equation, a function should be specified for $f(\cdot)$. However, for our purpose, it is only required that $f'(x) > 0$ for all x, which is usual since f-I curves generally have a

[17] The analysis presented in this section is quite standard for delayed dynamical systems.

positive slope. We assume that a solution for I_s^* exists, namely that I_{ext} is chosen in such a way that the curves $Jf(I_s^* + I_{\text{ext}})$ and I_s^* intersect for at least one point. For a sigmoidal-like f-I curve, as happens for noisy LIF neurons, the inhibitory network has a single solution for all I_{ext}. The excitatory net can have either one or three solutions. In the latter case, one fixed point occurs at low firing rate (near the rectified portion of the f-I curve), another at high rate (near the saturation of the f-I curve) and the last point lies in between.

A stability analysis of these fixed points can be performed. We shall show that, for this model: 1) the excitatory network cannot display oscillations for any delay, 2) when three fixed points exist, the lowest and highest fixed points are stable whereas the mid fixed point is unstable, and 3) the inhibitory network can bifurcate towards an oscillatory state for large enough delays. To this end, we write $I_s(t) = I_s^* + I_1 e^{\lambda t}$, with $I_1 \ll I_s^*$ and $\lambda \in \mathbb{C}$, and replace it in equation

$$\frac{dI_s}{dt} + I_s(t) = Jf[I_s(t - \tau_d) + I_{\text{ext}}] \tag{1.39}$$

to give the characteristic equation

$$\lambda + 1 = Jf' e^{-\lambda \tau_d}, \tag{1.40}$$

where f' is the derivative of $f(x)$ evaluated at $I_s^* + I_{\text{ext}}$. This transcendental equation must be solved for λ, resulting in an infinite and countable set of eigenvalues. These eigenvalues will be either purely real or come in complex conjugate pairs. Only two parameters matter: J and τ_d, the latter representing the ratio between the axonal delay and the synaptic time constant. First, note that $\tau_d = 0$ implies that the system is one-dimensional and therefore no oscillation can occur. In this case, fixed points are stable if $Jf' < 1$, and unstable otherwise. Increasing the delay makes the system infinite-dimensional. Purely real solutions only exist for an excitatory network. Writing $\lambda = \mu + i\omega$, real solutions correspond to $\omega = 0$, and Eq. 1.40 becomes $\mu + 1 = Jf' e^{-\mu \tau_d}$. Clearly, solutions exist only if $Jf' > 0$, i.e $J > 0$. If $Jf' > 1$ there exists a positive real eigenvalue and the fixed point is unstable for all τ_d. If three fixed points exist, one of them occurs when $Jf(I_s^* + I_{\text{ext}})$ crosses I_s^* with $Jf' > 1$; thus, this fixed point is unstable. We shall focus on the other fixed points which are stable.

To find the complex eigenvalues for $Jf' < 1$, we replace $\lambda = \mu + i\omega$ in Eq. 1.40 to give two equations to be solved numerically:

$$\begin{aligned} Jf' e^{-\mu \tau_d} \cos(\omega \tau_d) - \mu - 1 &= 0 \\ Jf' e^{-\mu \tau_d} \sin(\omega \tau_d) + \omega &= 0. \end{aligned}$$

Increasing τ_d from zero for J fixed, an oscillatory solution for the network's activity appears

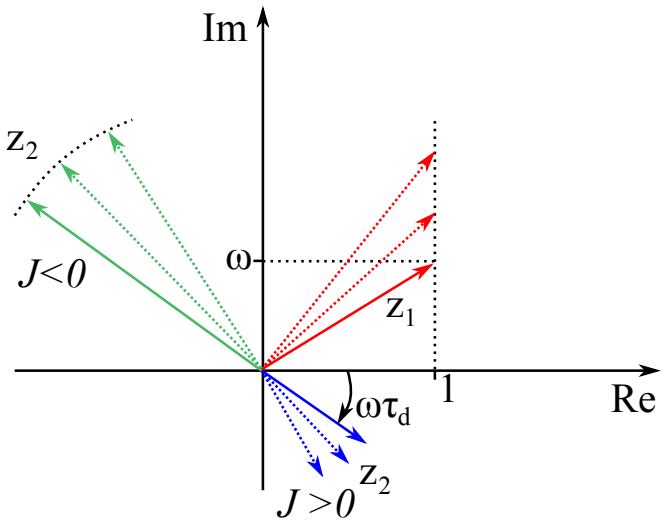

Figure 1.5: Graphical representation of Eq. 1.41. We used $f(x) = x\Theta(x)$, where $\Theta(x)$ is the Heaviside function. We depict the complex numbers $z_1 = 1 + i\omega$ and $z_2 = Jf'e^{-i\omega\tau_d}$ for $J > 0$ (blue arrows) and $J < 0$ (green arrows). For a given value of the delay τ_d, increasing ω moves z_1 along the vertical dotted line at $x = 1$ and z_2 rotates with constant magnitude. The vector for $J < 0$ has a phase advance of π in comparison to the vector for $J > 0$. Solution is reached when vector tips superimpose, which requires $|z_1| = |z_2|$ and $\arg\{z_1\} = \arg\{z_1\}$ mod 2π. For the excitatory network it never happens since $0 < Jf' < 1$ and the modulus of z_1 increases monotonically from 1 with increasing ω.

when a complex conjugate pair with a negative real part crosses the imaginary axis. Setting $\mu = 0$ in Eq. 1.40 yields a condition for the onset of oscillations:

$$1 + i\omega = Jf'e^{-i\omega\tau_d}. \tag{1.41}$$

A graphical interpretation of this equation is given in Fig. 1.5. From this interpretation stems the fact that, for $|Jf'| < 1$, the network cannot bifurcate towards an oscillatory state. We conclude then that the excitatory network cannot oscillate. For $Jf' < -1$, the inhibitory net undergoes a transition from a non-oscillatory to an oscillatory state when the delay is increased from zero (Fig.1.6A). The transition occurs at smaller delay for increasing Jf' in absolute value. Due to our nondimensionalization of the rate dynamics, the delay actually represents the ratio τ_d/τ_s. Hence, from our simple rate model we conclude that oscillations in inhibitory networks are more easily obtainable for small synaptic time constants and long axonal delays. The presence of delays is a necessary condition for the onset of oscillations. For quantitative results, the condition for marginal stability ($\mu = 0$), Eq. 1.41, can be recast

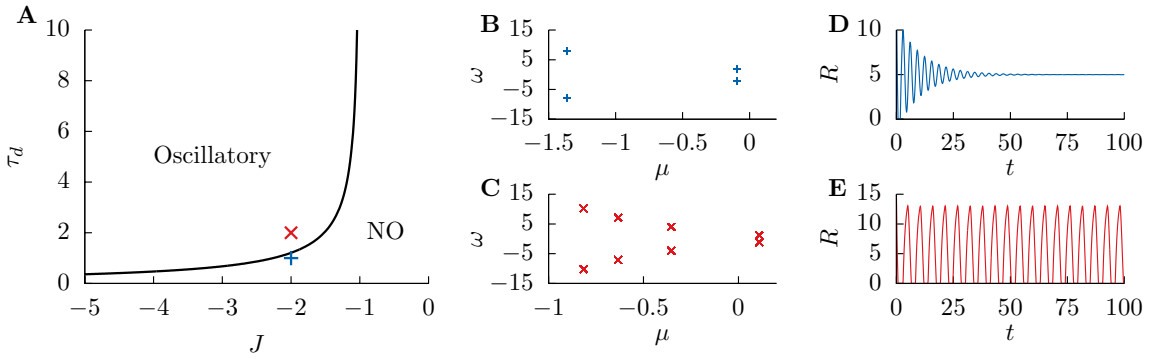

Figure 1.6: Oscillations in the inhibitory network. (A) Phase diagram. The black curve is the transition between the oscillatory and non-oscillatory (NO) states, as determined from Eq. 1.41 with $f(x) = x\Theta(x)$. The red and blue crosses give the values of J and τ_d used in the other panels. The red (respectively, blue) cross lies within the oscillatory (respectively, non-oscillatory) state. (B and D) For the values of (J, τ_d) corresponding to the blue cross of panel A, all eigenvalues have negative real parts (μ). The rate dynamics consists in a transient oscillation, followed by a convergence towards the fixed point (panel D). (C and E) The red cross lies within the oscillatory part of the phase diagram. There exists one eigenvalue pair with real parts greater than zero. Panel E shows the oscillation of the rate.

into two real equations, one for the amplitude of both sides and one for the phase:

$$\begin{aligned} \text{Amplitude}: \sqrt{\omega^2+1} &= |J| f' \\ \text{Phase}: \arctan(\omega) &= \pi - \omega\tau_d + 2\pi k = (2k+1)\pi - \omega\tau_d, \end{aligned}$$

with $k \in \mathbb{Z}_+$ and π comes from $J = |J|e^{i\pi}$ when $J < 0$. For a fixed delay, solving the phase equation for different k gives different eigenfrequencies at marginal stability. When these frequencies are known, they can be used to solve the amplitude equation for $|J|$; it gives the coupling value at which this mode loses stability. The transition line in Fig. 1.6A represents the smallest $|J|$ above which oscillatory solutions exist for a given τ_d.

These results extend to more complicated models. For instance, if the time scale of the rate dynamics is nonzero, instead of Eqs. 1.37 we would have

$$\begin{aligned} \tau_s \frac{dI_s}{dt} + I_s(t) &= J\tau_s R(t-\tau_d) \\ \tau_R \frac{dR}{dt} + R(t) &= f[I_s(t) + I_{\text{ext}}]. \end{aligned}$$

This is a single-population delayed Wilson-Cowan equation [224]. In that case, it is easy to

show that the characteristic equation becomes:

$$(\tau_R\lambda + 1)(\tau_s\lambda + 1) = Jf'e^{-\lambda\tau_d}. \tag{1.42}$$

Again, the excitatory network is unstable for all τ_d when $Jf' > 1$ and, for $0 < Jf' < 1$, the equation for the transition to an oscillatory state, $(i\omega\tau_R + 1)(i\omega\tau_s + 1) = Jf'e^{-i\omega\tau_d}$, has no solution in $\mathbb{R}$.

Emergence of oscillations

In the previous section we used a simple so-called rate model to understand the impact of synaptic and axonal propagation time scales on the occurrence of oscillations in purely excitatory and inhibitory networks. However, this kind of model fails to adequately represent highly synchronous neural activity [192]. Also, it conceals the important role played by the subthreshold dynamics of neurons, the structure of the network connectivity, and sources of noise. In the present section we first discuss, using the phase-resetting theory, the circumstances under which excitatory or inhibitory networks synchronize in absence of delay. We then discuss the impact of noise on rhythms. We postpone the discussion of network connectivity to the next section.

Above, for the excitatory network, it was tacitly assumed that whenever the neurons receive excitation the activity increases. However, this statement is not universally true and depends on the timing of the synaptic coupling. Neurons can be characterized by their phase-response curve (PRC) [82, 81]. Assume that a neuron fires at time t. If, after the absolute refractory period is over, an excitatory input systematically advances the time of the next spike, the neuron is said to possess a *type I PRC*. On the contrary, if for a short time after the waning of the absolute refractory period an excitation leads to a delay of the next spike, we say that the neuron has a *type II PRC*. For fully connected excitatory networks and physiologically sound synaptic coupling strengths, it has been shown that type I excitatory neurons cannot collectively synchronize, but type II neurons can, provided that synaptic excitation is fast enough [100]. Due to the frustration effect occurring in large networks of type I neurons, partial synchrony—where separate clusters of neurons oscillate in synchrony—is nevertheless possible. Van Vreeswijk and collaborators [217] further noticed that mutually coupled inhibitory neurons tend to synchronize when the rise time of the synaptic current is longer than the action potential duration. To summarize, fast synaptic interaction favours the synchronization of excitatory neurons, whereas slow interactions promote in-phase firing of inhibitory neurons. These results were obtained with neurons that are intrinsic oscillators–which can fire repetitively in the absence of periodic input—and

without noise. An extensive literature exists on the synchronization of coupled oscillators and pulse-coupled oscillators inspired by neural dynamics (for instance, the classic [147]).

Network oscillation does not necessarily imply perfect synchrony for the whole population. Noise—whatever its origin—acts as a desynchronizing force, hindering the phase-locking. The typical picture is the following for homogenous (i.e., with identical units) and fully recurrent (i.e., all-to-all connectivity) networks. For high neuronal noise and weak synaptic coupling, the asynchronous state prevails and the population activity is on average equal to the firing rate of the individual neurons. More precisely, using the notation of section 1.1.3, in the asynchronous state $\langle R_M(t) \rangle = r$, where r is the (constant) individual firing rate of any neuron in the network. The individual spiking activity of the neurons is irregular due to the high noise (Fig. 1.7C). If now the synaptic coupling is increased beyond a given value, the asynchronous state loses stability and oscillations develop (Fig. 1.7D). In this oscillatory state at high noise and relatively strong coupling the individual spiking activity of neurons is irregular, even if the population as a whole oscillates. Decreasing the noise gives rise to oscillations involving cluster states, wherein subsets of neurons fire in near-synchrony, with a relative shift between the clusters' spikings (Fig. 1.7B). Note that it is also possible for the network to be in the asynchronous state while the individual spiking activity of the neurons is fairly regular [43], which happens at low noise and very low coupling (Fig. 1.7A). The oscillation frequency depends in a nontrivial manner on the synaptic time constants (delay, decay and rise times; see section 1.3.2 above), the noise level (high or low), the type of noise (fast or slow, see section 1.1.2) and neuronal parameters (e.g., membrane time constant). Section 2.7 of the next chapter will describe more precisely the influence of these factors on rhythms.

Sparse networks

In fully recurrent neural networks all neurons are interconnected. In real neural systems, however, the connectivity is not expected to be so dense. In cortical areas, neurons receive $\sim 10^3$ synaptic contacts [2], but belong to a large network. Hence, to reproduce the irregular cortical spiking measured *in vivo* (*cf.* section 1.1.2), it seems appropriate to consider networks with a sparser connectivity. More generally, network connectivity can be implemented by defining a binary connectivity matrix $\mathbf{W}_b$. Element ij of this matrix is one if neuron j is connected to neuron i and zero otherwise. Fully recurrent nets have all entries equal to one. For a sparse network, only a small fraction of the entries are nonzero.

Brunel and coworkers [42, 44] have considered a network of N sparsely connected inhibitory LIF neurons. Every neuron received C synaptic connections from neurons within

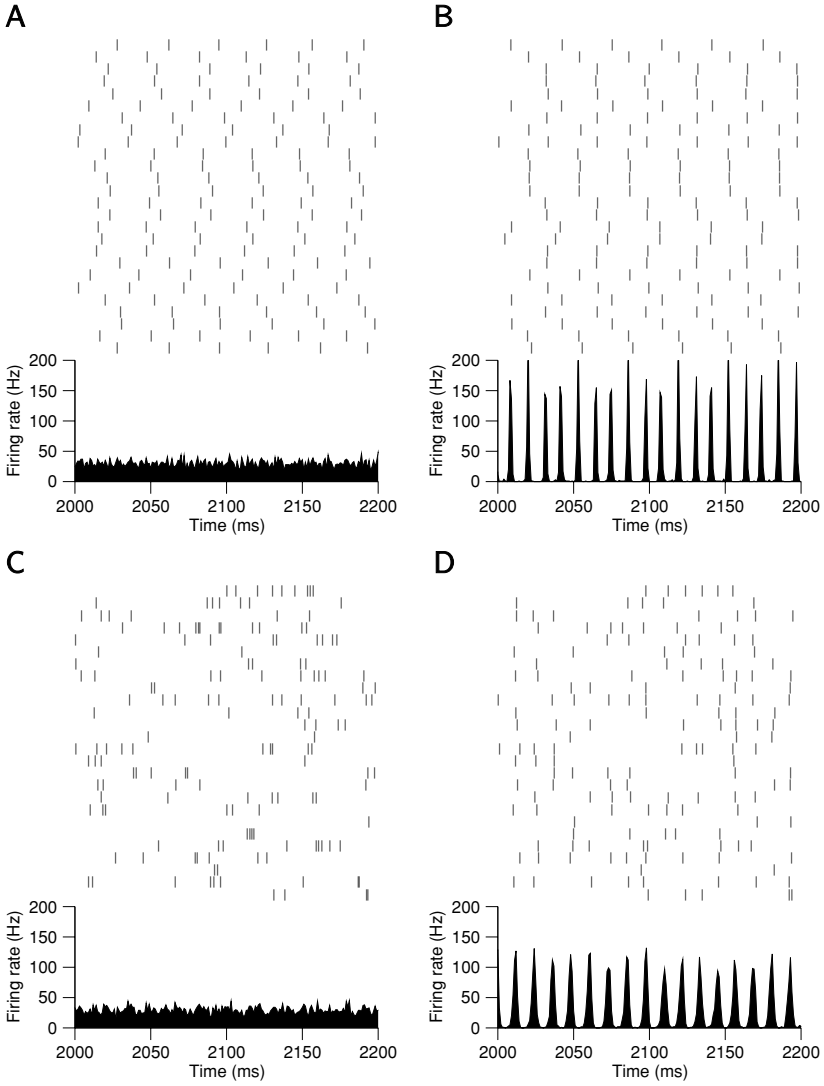

Figure 1.7: Influence of neuronal noise and synaptic coupling strength on the population activity and single-neuron spiking of a network of fully connected inhibitory neurons. The top of each panel represents the spiking activity of a subset of neurons within the population. Each vertical bar is a spike. The bottom of each panel is the population activity. A) The noise level and the synaptic coupling are small, but the network is nevertheless in the asynchronous state. Single-neuron activity is regular. B) Noise is slightly decreased with respect to its level in panel A. The population oscillates in clusters (here, three clusters). C) High noise and strong coupling. The asynchronous state is stable and the single-neuron activity is irregular. D) Noise is lower than in panel C, and a rhythm with irregular neuron spiking occurs. Figure from [43], with permission.

the network, giving a connection probability $p = C/N \ll 1$. All neurons also received noisy excitation that is uncorrelated across the population. The sparsity of the network connectivity and the external noise produce irregular spiking when the average input (external excitation plus recurrent inhibition) is subthreshold: the neurons are noise-driven as in section 1.2.1. In the limit $N \to \infty$, the inhibitory neurons are effectively uncorrelated, i.e., the probability that they share a common input vanishes. A minimal C is required for synchronization to be possible [220, 95]. Analytics have shown that, for fast inhibitory synapses, the critical parameter is the ratio between the mean inhibition and the intensity of the total noise driving the neuron (external and coming from within the network). The asynchronous state becomes unstable when this ratio becomes greater than a given threshold. This is because neurons need to be able to detect the inhibitory current above the ambient noise in order to be modulated by the former, generating oscillations. The resulting rhythm has a higher frequency than the single-neuron firing rate. It corresponds to the situation of the preceding section (section 1.3.2) where the neurons fire in clusters and not at every cycle of the oscillation. The single-neuron activity is irregular.

Hence, sparse inhibitory networks reconcile the notion of irregular low-firing neurons with the occurrence of coherent rhythms at the population level. Pronounced oscillations occur when the axonal delay is large and when the rise and decay time constants are short [44]. For finite-size networks the phase boundary between asynchronous and oscillatory activity is "blurred" by finite-size fluctuations, resulting in damped oscillations in the vicinity of the transition. This constitutes a plausible scenario for the decoherence of gamma rhythms, which will be discussed in chapter 5.

1.3.3 Excitatory-inhibitory networks

In networks with either inhibitory (I) or excitatory (E) neurons the presence of synaptic delays greatly promotes the onset of rhythms. In networks of mutually connected E and I cells, oscillatory instabilities easily appear even with instantaneous interactions. From the dynamical systems point of view the underlying explanation is almost trivial: adding a dynamical variable (the activity of the second population) to a one-dimensional system increases its dimensionality and may allow Hopf bifurcations, i.e. the emergence of oscillations in nonlinear systems from the migration of roots across the imaginary axis. Including delays, even in firing rate models, can produce a very rich phase diagram containing chaotic states [21].

The situation stays interesting and complicated once the spiking activity of the neurons is taken into account. In [40] Nicolas Brunel studies the dynamical states of a sparse E-

I network of LIF neurons. The configuration is similar to that exposed in the preceding section: each neuron receives $C_E = pN_E$ synaptic connections from E neurons, $C_I = pN_I$ connections from I neurons and external noisy excitation (all connections are randomly selected). Synapses are fast, i.e. the rise and decay times are zero, but include a delay for recurrent connections. The important factors are the mean external input, the delay and the balance between the mean recurrent excitation and inhibition. Three types of activity have been delineated [40]. When excitation dominates inhibition, the network is in a synchronous state with neurons firing regularly at high rate; when inhibition dominates excitation, single-neuron firing is irregular in an asynchronous state and synchronous oscillations appear for strong inhibition.

With different neuronal properties for the E and I populations and synaptic delays depending on the connections (I-I, E-I, I-E, E-E), phase lags between the oscillations of the E and I subnetworks may appear. The oscillation frequency depends on the connectivity of the populations: if no recurrent connections exist within the I and E subnetworks, rhythms tend to be slower than the inhibition-only scenario of the preceding section. Introducing I-I recurrent connections (and still no E-E connections) increases the oscillation frequency, whereas E-E connections decrease it. With all four connections, the population frequency decreases with the ratio recurrent excitation/recurrent inhibition (when this ratio is the same for both subpopulations) [44]. Oscillation frequencies tend to lie in the gamma range. Slower excitation favours the I-I subnetwork and produces faster oscillations; fast excitation produces slower rhythms by promoting the E-I connections [44].

With weak or no noise, perfect synchronization ("spike-to-spike synchrony" [44]) is possible in both sparse E-I networks [35] and purely inhibitory networks [95]. Neurons in the network then fire regularly and the dynamics of the network are that of coupled oscillators. Synchronization in sparse I-E networks is favoured by faster decay of inhibition [35]. When the asynchronous firing state becomes unstable at higher noise it gives rise to a "firing-rate synchrony" state [44], as exposed in this section and the preceding one. The underlying neural firing is typically irregular and with a lower rate than the oscillation frequency of the population. For purely inhibitory networks, oscillation frequencies are slower for spike-to-spike synchrony than for firing-rate synchrony [44]. Note that linear stability analyses yield appropriate predictions when oscillation amplitudes are not too large, i.e. not when spike-to-spike synchrony occurs.

1.3.4 ING and PING mechanisms of gamma rhythm generation

Purely inhibitory, purely excitatory and excitatory-inhibitory networks are all capable of displaying collective oscillations under certain circumstances, but they are not all good candidates for gamma-rhythm generation. From the start, "purely inhibitory" is a bit of a misnomer because excitatory inputs are necessary to yield a nonzero fixed-point activity about which the network oscillates. For instance, in subsection 1.3.2 (Emergence of oscillations), extrinsic excitation was provided by a population of randomly firing excitatory neurons. On the other hand, purely excitatory networks tend to produce high firing rates and their propensity to oscillate depends on the level and characteristics of the ambient noise, as seen above. Moreover, in the cortex, a subclass of inhibitory cells—the parvalbumin-positive fast-spiking basket cells—has recently been shown to be mainly responsible for gamma oscillation generation by optogenetic experiments: inhibiting these neurons suppress gamma-band rhythms *in vivo* [198]. Hence, inhibition is required. We are then left with excitatory-inhibitory networks.

Two principal scenarios for gamma-rhythm generation in E-I networks have been described (Fig. 1.8). The interneuron network gamma (ING) requires a synchronization of inhibitory neurons, but not the excitatory cells. In [220], Wang and Buzsáki show that the ING is not robust to the presence of heterogeneities in the cell assembly when the inhibitory neurons fire at each gamma cycle, i.e. with spike-to-spike synchrony. Oscillations are more robust in the presence of noise and heterogeneities when inhibitory neurons skip cycles [205]. However, this scenario may promote rhythms faster than the gamma range. For the pyramidal-interneuron gamma (PING) mechanism, synchronization of both the excitatory and the inhibitory neurons happens. The synchronized excitatory cells send a volley of excitation to inhibitory neurons, making them fire; the inhibitory cells, in turn, inhibit themselves and the excitatory cells [35]. Sometimes, researchers qualify the gamma rhythm as weak—for either the ING or PING mechanism—when the firing rate of the excitatory cells is low compared to the oscillation frequency [34]. In chapter 5, we will propose another putative gamma-generation mechanism that produces weak oscillatorylike activity in a delayed feedforward inhibitory network.

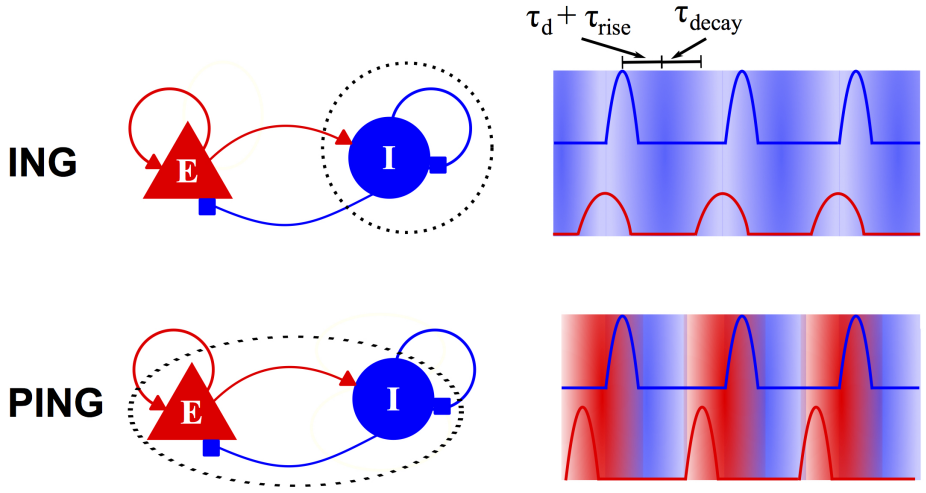

Figure 1.8: Schematics of the ING and PING mechanisms. For the ING mechanism (top), rhythm is fully generated by the inhibitory (I) population. The various synaptic time constants (τ_d = axonal delay; τ_{rise} = rise time of the conductance; τ_{decay} = decay time) set the scale of the oscillation period. The participation of the excitatory (E) population is important only to the extent that it depolarizes the inhibitory cells. E neurons are entrained by the oscillation, but do not synchronize intensely (hence the lower peaks in the E cells' activity; red trace). For the PING mechanism (bottom), the E and I neurons collaborate to establish the rhythm. Contrary to the ING mechanism, the E neurons synchronize and help to synchronize the I cells as well. Gradients in the figure illustrate the rises (light to dark) and falls (dark to light) of excitation (red) and inhibition (blue). All traces represent activity of the E (red) and I (blue) population. Figure redrawn and slightly modified from [206].

1.4 Electrosensory system

The ultimate goal of neuroscience is to understand the inner functioning of the human brain. Its complexity, however, is off-putting. Therefore, neurobiologists typically rely on animal models whose brains, although not simple, do not display the level of sophistication of the human brain. Monkeys, rats and mice are often used due to the presence of cortical structures. Other animals are used for their sheer simplicity, like the worm *C. elegans*, whose neural system contains about 350 neurons and 7000 synapses [66]. The set of behaviors displayed by such a simple animal is, however, limited. Moreover, other animals are studied because they developed a particular sense to its fullest, like echolocation in bats [216] or hearing in barn owls [53]. Deciphering how neural structures in these animals attain optimal performance may help us understand our own senses. In this section, we describe the physiology and anatomy of the neuronal network which inspired the contributions of this thesis: the electrosensory lateral line lobe (ELL) of weakly electric fish. For these fish, behaviorally

relevant stimuli can be artificially generated in laboratory and their neuroanatomy and neurophysiology have been extensively characterized [133].

The account that follows is far from being exhaustive. After a short general description of the electrosensory system, we characterize the nature of the stimuli that these fish receive and describe the ascending and descending pathways that innervate the principal cells of the ELL. A more detailed analysis of neuron and network models associated with this system is presented in chapter 2 in the context of neural rhythms and gain control.

1.4.1 Global description

Weakly electric fish[18] are small ($\sim$ 10 cm) fish living in turbid salt waters. An electric organ located in their tail produces an oscillating dipole-like electric field ($\sim$ 1 kHz) around their body. This field is called the electric organ discharge (EOD). The EOD amplitude is small ($\sim$ 10 mV/cm), so it is not capable of stunning preys—hence the adjective *weakly*. Instead, the EOD is used for object localization and communication with conspecifics. Objects with electrical properties different from the surrounding water and electrocommunication signals cause amplitude modulations of the EOD which are detected by specialized cutaneous electroreceptors. Primary afferents innervate these electroreceptors, and amplitude modulations of the EOD change their firing rates with respect to their baseline activity (i.e., in absence of stimuli in the water).

Electroreceptor afferents send their axons to the electrosensory lateral line lobe (ELL), a large structure in the hindbrain of the fish and the focus of most of this thesis. The principal (pyramidal) cells of the ELL project to two neural assemblies (Pd and EGp; see below) which send feedback signals to various parts of the ELL. ELL cells also project to a structure called the *torus semicircularis*, which then projects to the *optic tectum* and various brain stem regions [24]. Among brain stem regions, the nucleus *electrosensorius* projects to the prepacemaker nucleus which regulates the EOD frequency through the medullary pacemaker nucleus.

As will be detailed below, the ELL has a well-defined laminar and columnar structure. It receives a vast amount of feedback from higher processing centers as well as feedforward inputs from afferents and local inhibitory neurons. Quite strangely, however, recurrent connections between pyramidal cells in the ELL do not seem to exist [133].Therefore, this architecture is a unique setup to study feedforward and feedback contributions to efficient sensory processing.

[18] We focus our attention on the active electroreception of wave-type electric fish and we had in mind the species *Apteronotus Leptorhynchus* while writing this section.

1.4.2 The nature of electrosensory stimuli

The electric organ of weakly electric fish has been modelled as a continuous bimodal distribution of current sources located along the spinal chord of the fish [12]. Using a quasi-static approach [174, 12], the electric potential Φ everywhere in space (inside the fish, inside the skin and in the surrounding water) can be found by solving Poisson's equation [13]. The EOD is a dipole-like electric field (Fig. 1.9A). Since the current distribution is located in the fish's tail, the zero-potential line lies in the tail region; the head and the trunk share the same sign for the potential, whereas the tail has the opposite sign. A large concentration of electroreceptors is located in the head region. The current density oscillates with time, producing the fast oscillating EOD field. The important quantity is the transdermal potential, the potential difference across the fish's skin where electroreceptors are located, especially in the head region.

When the water is free from perturbations—e.g., the fish is immobile in a large water tank—the transdermal potential at a chosen point on the skin varies rapidly at the frequency of the EOD, ~ 1 kHz, and with a given amplitude ~ 1 mV (Fig. 1.9A, top). So-called tuberous electroreceptors sense the transdermal potential and make chemical synaptic contacts with electroreceptor afferents (EAs), whose cell bodies are located in the anterior lateral line nerve ganglion. Usually, one views the compound receptor + synapse + afferent as a single coding unit, with the end product being the EA's spiking activity. Two types of units exist: P- and T-units, differing by the way their spiking activity is modulated by the EOD. We will focus on the P-units here. P-units fire in a phase-locked manner to the EOD, but do not fire at every cycle. Every cycle, the EA has a spiking probability between 0.1 and 0.5 (P in P-units stands for probability). Even without stimulation, P-unit fire rates range between 100 and 500 Hz.

Experimentalists use small metal spheres (a few millimeters in radius) to mimic the presence of small preys on which these fish feed. An object with a conductivity and capacitance different than the surrounding water perturbs the EOD. In a simplified model, a small spherical perfect conductor with radius a centered at position $\mathbf{r}_0$ with respect to a stationary frame of reference yields a dipolar disturbance of the potential at point $\mathbf{r}$ given by [173]

$$\delta\Phi(\mathbf{r}) = -a^3 \nabla\Phi(\mathbf{r})|_{\mathbf{r}_0} \cdot \frac{(\mathbf{r}-\mathbf{r}_0)}{|\mathbf{r}-\mathbf{r}_0|^3} \tag{1.43}$$

where $-\nabla\Phi(\mathbf{r})|_{\mathbf{r}_0}$ is the field evaluated at r_0 *before* putting the sphere. The transdermal potential at a given point $\mathbf{x}$ on the skin is then the difference between $\delta\Phi$ just outside and just inside the skin: $\Delta\Phi_{\mathrm{trans}}(\mathbf{x}) = \delta\Phi_{\mathrm{out}}(\mathbf{x}) - \delta\Phi_{\mathrm{in}}(\mathbf{x})$. The difference of transdermal

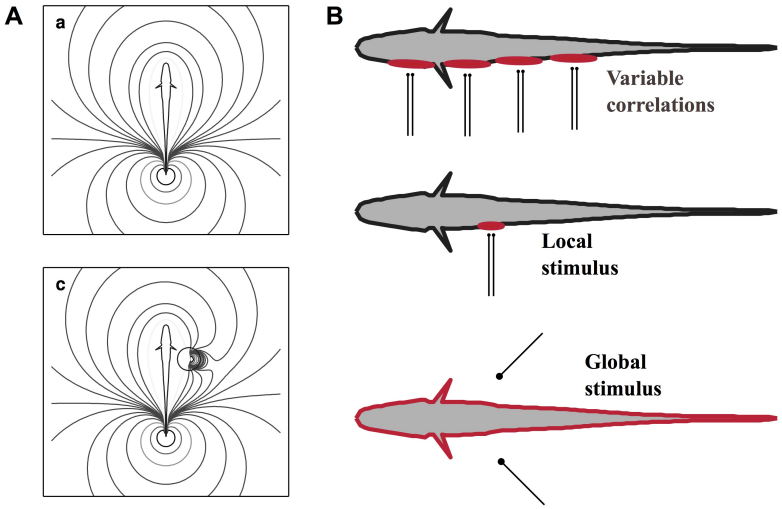

Figure 1.9: Electric organ discharge and experimental stimulus configurations. (A) Equipotential lines of the EOD (top) and perturbation profile (bottom) caused by a sphere in the surroundings. (B) Stimulus configuration used in experiments. Red areas are meant to represent the spatial extent of electric images. Top: Stimulus with spatial correlations changeable by controlling the current delivered to each of the four electrodes [77]. Middle: Local stimulation with a dipole. Bottom: Global spatially correlated stimulation; in reality, stimulating electrodes are farther away. Panel A extracted from [59], figure 2, with permission.

potential with and without the object in the water is called an *electric image*. The object locally modulates the amplitude of the EOD across the skin : an amplitude increase implies an enhanced probability of firing for the EAs (see Fig. 1.9A, bottom). Prey-like signals typically contain low temporal frequencies (< 25 Hz), given the low swimming speed of weakly electric fish [133].

Another important type of stimuli are communication signals with conspecifics. Each fish has its own EOD frequency f. The EODs of two fish in the same neighborhood will add, producing a beat pattern in the EOD oscillations. The beat frequency is the difference between the EOD frequency of both fish. Some weakly electric fish forage in groups of several individuals and the beat pattern can become complicated and aperiodic (due to the sum of incommensurate-frequency sine waves). Efish can respond to beat patterns by transiently increasing their EOD frequency, producing a "chirp" [133]. These chirps are used in courtship and antagonistic encounters.

In experiments, various stimulation configurations can be designed (Fig. 1.9B). A *local*, prey-like, stimulus can be a metal sphere, as explained above, or a small current dipole which mimics the disturbance of a sphere [17]. The dipole is placed close to the fish and its

electric field matches the phase and frequency of the fish's EOD. Amplitude modulations (AMs) can then be applied (steps, sinusoids, random modulations, etc) and be detected and processed by the fish's nervous system. *Global* stimulations intended to activate P-units of the whole body are obtained by using stimulation electrodes far from the body. Using the same protocol as with the stimulating dipole various temporal patterns can be used. For instance, random AMs can be used to mimic a large number of neighboring fish in the water. The global stimulation just described is said to be *spatially correlated* because EAs receive more or less the same stimulation. In [77], the authors used four dipoles located on one side of the fish, near the skin surface. Each dipole electrically "illuminated" a separate portion of the cutaneous surface using a random signal. The stimulation from all dipoles could be *spatially correlated*—when all dipoles produced the same random signal in phase—, or *spatially uncorrelated*—when the random signals were uncorrelated.

Experiments on weakly electric fish can be performed under either *in vitro* (slices of brain material) or *in vivo* (living animal) conditions. The main experimental results inspiring our modeling work in chapters 3-5 were obtained *in vivo* [54, 75, 77, 16, 15]. First, the fish must be anesthetized and a small region of skull is removed [19]. A recording electrode is then inserted and the experimentalist targets the portion of the brain from which he/she wants to record. The fish is subsequently suspended in a water tank and artificially respirated by a flow of water. For some weakly electric fish species—like the one used in the aforementioned studies—the electric organ producing the EOD is of neurogenic origin, meaning that it resists the paralysis induced by muscle-targeting drugs during the experiment [19]. The EOD is thus fully functional and can be recorded with electrodes placed near the animal. Finally, the experimentalist can use any of the stimuli described above and carry out the experiment.

1.4.3 Ascending pathway and architecture of the ELL

Electroreceptor afferents project to the electrosensory lateral line lobe (ELL), in the hindbrain of the fish (Fig. 1.10A-B). This structure is the first-order processing center of electrosensory signals. The projection from the EAs is topographic: two afferents innervating adjacent skin areas will terminate with the same relative position in the ELL [118]. In the ELL, each P-unit trifurcates and ends in three separate segments of the ELL: the centromedial (CMS), centrolateral (CLS) and lateral (LS) segments depending on their position within the lobe (Fig. 1.10B; [193, 194]). Each segment constitutes a topographic map of the fish's body. In a transversal (horizontal) cross-section through the ELL, these topographic maps are along the rostro-caudal axis (see Fig. 1.10A for a definition of rostral and caudal). Each segment

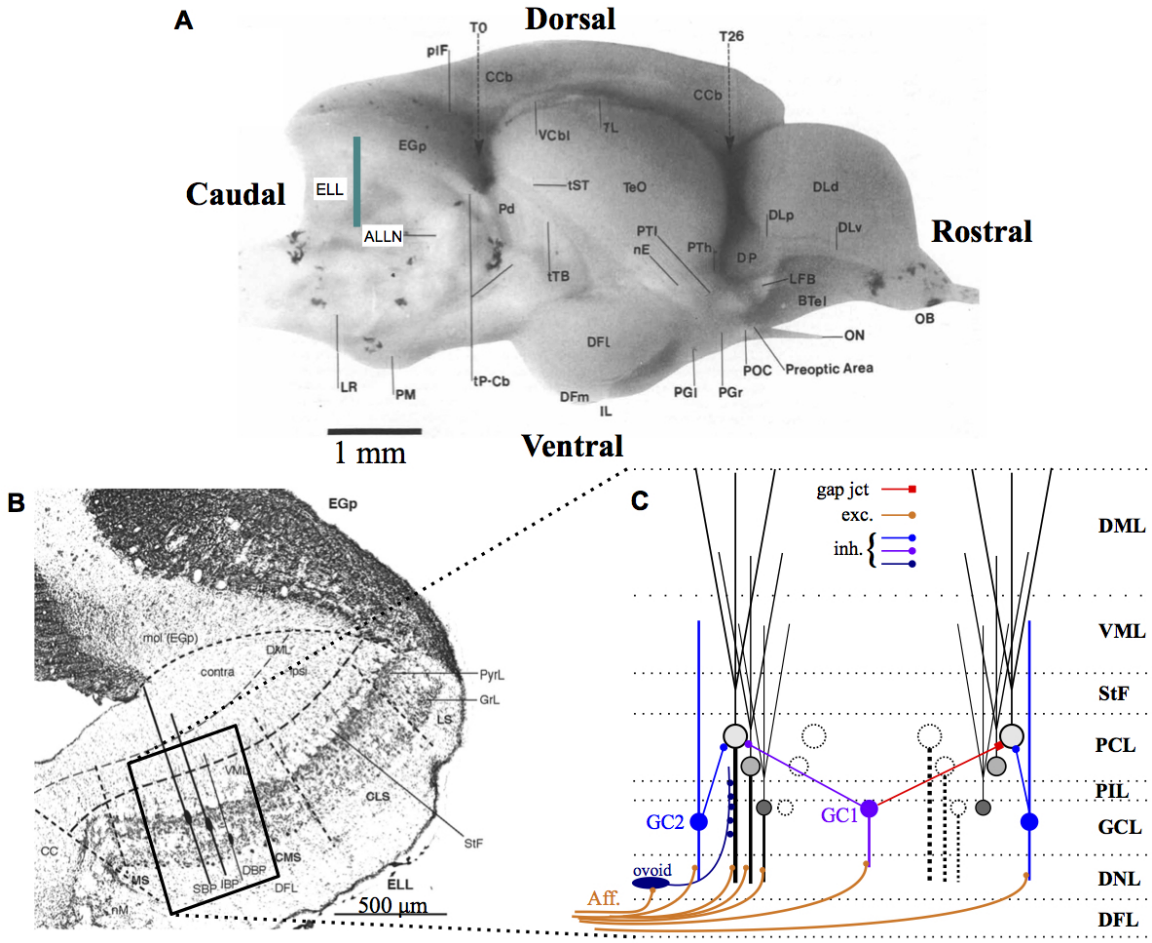

Figure 1.10: The weakly electric fish's brain, its ELL and the ascending pathway.
(A) Lateral view of the weakly electric fish's brain. Blue bar indicates approximately the position of the cut illustrated in panel B. Note the position of the ELL in the caudal part of the brain, the EGp, just above it, the Pd, rostral to the ELL and the ALLN, the nerve formed by electroreceptor afferents.
(B) Coronal section through the ELL, showing the EGp (dark agglomeration at the top), the three segments of the ELL and the real position of the layers schematized in panel C. (Notation: S(I,D)BP = superficial (intermediate, deep) basilar pyramid; PyrL corresponds to PCL in panel C while Grl corresponds to GCL.)
(C) Schematic of the ascending pathway to the ELL. Pyramidal cells' somata are drawn in gray scale (light gray for superficial cells and dark gray for deep cells. Two columns are depicted. For each column, either the basilar (ON; left column) pyramidal cells or the nonbasilar (OFF; right column) pyramidal cells are drawn while the position of the other cells within the columns are suggested by dotted lines. Inhibitory neurons are in cold colors, whereas excitatory neurons are drawn in warm colors. See text for a detailed explanation and table 1.1 for a short description of the layers. Panels A and B slightly modified from [135] and [134], respectively.

Table 1.1: Short description of ELL layers

	Layers of the ELL
Dorsal molecular layer (DML) :	Termination sites of parallel fibers coming from EGp
Ventral molecular layer (VML) :	Termination sites of excitatory fibers from the StF
Stratum fibrosum (StF) :	Axons coming from neurons in the Pd nucleus
Pyramidal cell layer (PCL) :	Contains PCs' somata
Plexiform layer (PIL) :	Contains efferent axons from ELL's PCs
Granule cell layer (GCL) :	Contains granule cells and the somata of DP cells
Deep neuropil layer (DNL) :	Termination sites of EAs' axons on PCs and GCs.
Deep fiber layer (DFL) :	Contains EAs' axons.

has the same global architecture but their pyramidal cells differ by the size of their receptive fields [134], their ability to discharge action potentials in volleys, known as "bursts" [214], and the temporal aspects of their stimulus processing [114, 115].

In a coronal[19] plane through the ELL each map is organized in layers (Fig. 1.10C). The principal neurons of the ELL are (excitatory) pyramidal cells (PCs) which come in two classes [133]. The first class comprises PCs having a thick basilar dendrite whereas the second class has no basilar dendrite [136]. Each class is further subdivided according to their position within the ELL; from ventral to dorsal, we typically delineate the deep, intermediate and superficial PCs (Fig. 1.10C; [19]). EAs' axons enter the ELL through the ventral part. They form excitatory synapses with the basilar PCs and with inhibitory granular interneurons.

Basilar and nonbasilar PCs also differ by their receptive fields, which are a direct consequence of their respective morphologies [17]. Basilar PCs have on-center/off-surround receptive fields, whereas nonbasilar PCs have off-center/on-surround receptive fields [187]. The exception are the deep PCs, which, due to their ventral position in the ELL, do not have a surround. In the following, basilar PCs will be called *ON* cells and nonbasilar PCs *OFF* cells, in analogy to the terms used for similar cells in the retina (historically ON cells are referred to as E cells and OFF as I cells in the electric fish literature [133]). As depicted in Fig. 1.10C, ON cells possess an "on" center because they receive direct excitation (AMPA and maybe NMDA) from the EAs. Researchers have postulated that their "off" surround

[19] A coronal plane is a vertical plane dividing the body—including the brain—into ventral (belly) and dorsal (back) sections.

Table 1.2: Abbreviations for cell types

Cell types	
EA = electroreceptor afferent	GC = granule cell
DP cell = deep PC	BP = bipolar cell
IP cell = intermediate PC	St = stellate cell
SP cell = superficial PC	MP= multipolar cell
PC = pyramidal cell (generic)	

is caused by inhibition coming from remote inhibitory interneurons GC1. Another type of granular interneuron, the GABAergic GC2, may not participate in the surround inhibition, but instead filter the excitation received from the EAs through $GABA_A$ receptors on ON cells. On the other hand, the "on" surround of OFF cells is caused by gap junctions (electrical synapses) with either GC1s or GC2s. The "off" center comes from inhibition from these cells.

ON and OFF PCs are arranged in columns, as depicted in Fig. 1.10 [134]. Each column contains a deep, an intermediate and a superficial PC of each class. Other ELL layers are described in table 1.1. The ON and OFF deep PCs project to the *nucleus praeementialis dorsalis* (Pd), whose stellate and bipolar cells send direct feedback projections to the ELL. The multipolar cells of this nucleus project to the *eminentia granularis posterioris* (EGp), whose granule cells projects back to the dorsal part of the ELL. The next subsection focuses on these two feedback pathways.

1.4.4 Descending pathways: the direct and indirect feedback

Two feedback pathways are defined: the *direct* feedback pathway, ELL $\rightarrow$ Pd $\rightarrow$ ELL and the *indirect* feedback pathway, ELL $\rightarrow$ Pd $\rightarrow$ EGp $\rightarrow$ ELL.

For the direct feedback, axons from bipolar (BP) and stellate (St) cells form the *stratum fibrosum* and enter the ELL just above the pyramidal cell layer (Fig. 1.11). The BP cells inhibit the somatic region of PCs through slow $GABA_B$ synapses (reversal potential ~ -90 mV). Nonlinear interactions with hyperpolarizing I_A-like potassium channels may enhance the resulting inhibitory postsynaptic potentials and modify their temporal dynamics [28]. Inhibition from BP cells is diffuse, i.e., non-topographic. Excitatory St cells, like the ELL PCs, can be either ON or OFF [38]. They do not fire spontaneously, they are unresponsive to

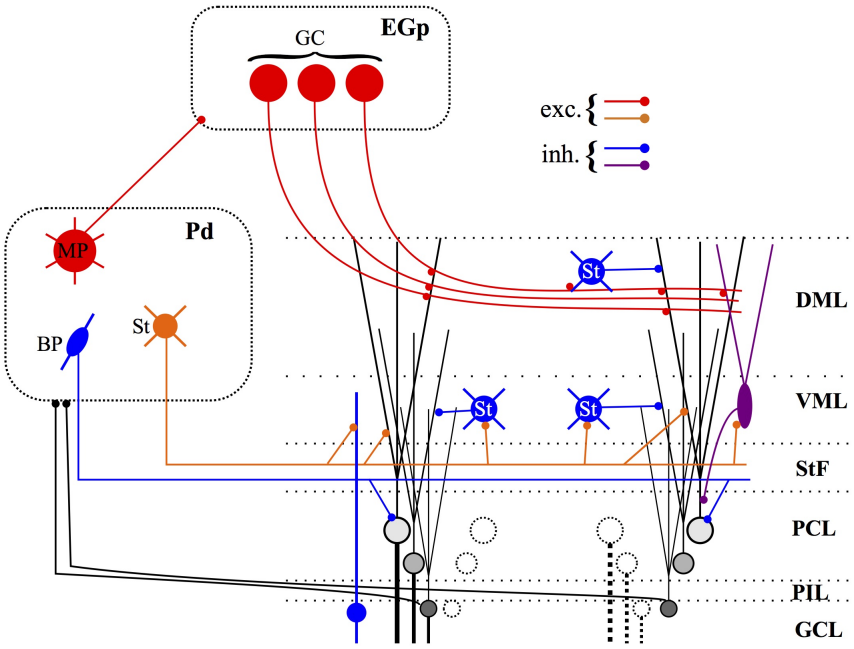

Figure 1.11: Direct and indirect feedback to the ELL. See text for description. Diagram is schematic and obviously not to scale.

amplitude modulations with frequency above 16 Hz and they strongly respond to moving local stimuli [38]. St cells make monosynaptic contacts with pyramidal cells through AMPA and NMDA receptors in a topographic manner; they also disynaptically inhibit pyramidal cells *via* local inhibitory interneurons.

It has been suggested that the direct feedback pathway implements a sensory searchlight [28, 30, 29]. A prey in the surroundings of the electric fish will activate EAs belonging to the receptive field of a subset of pyramidal cells. As a result, some St cells—whose receptive field is larger than those of PCs—will increase their activity and excite the PCs in a topographic manner. At the same time, BP cells diffusively inhibit all pyramidal cells. A PC will only be activated if it receives concomitant excitation from the St cells and the afferents, so as to overshadow the inhibition from the BP cells. The effect is enhanced by various nonlinear processes on the PCs and frequency-dependent synaptic facilitation of the excitatory part of the direct feedback. The searchlight mechanism has been suggested to help the fish detect and focus on preys. On the other hand, and most importantly for this thesis, the direct inhibitory feedback has been proven essential for establishing oscillatory neural states in ELL PCs under spatially correlated global stimulation [75]. Such an oscillatory response is not recorded under local [75] or spatially uncorrelated global stimuli [77]. Models of

this phenomenon have involved delayed inhibitory networks of fully connected neurons [77, 125], as will be explained in chapter 2. The rationale behind this perspective is that the loop PCs $\rightarrow$ BP cells $\rightarrow$ PCs (excitatory neurons to inhibitory cells and back to excitatory neurons) can be viewed, to simplify, as a self-inhibition of ELL pyramidal cells.

The indirect feedback starts in the Pd nucleus. There, multipolar (MP) cells receive excitatory inputs from the deep PCs (Fig. 1.11). MP cells have high spontaneous firing rates and do not respond well to moving objects. MP cells project to the EGp and provide excitation to its granular cells population. The axons of these excitatory neurons constitute the parallel fibers entering the ELL in its dorsal molecular layer. The parallel fibers make synaptic contacts with the distal part of PCs' apical dendrites, and with local interneurons (stellate cells and vml cells). The indirect feedback has two principal roles: 1) controlling the gain of superficial PCs [16, 15], and 2) providing a sophisticated redundancy reduction mechanism through burst-timing dependent plasticity [102, 32, 33]. Redundant inputs are predictable signals (e.g., modulations of the EOD strength on the skin caused by tail bending) that the animal wants to discard in order to detect and highlight behaviourally important stimuli. Concerning point 1, it is the concomitant inhibition from the vml cell—unlabeled purple cell in Fig. 1.11—and the distal inhibition from the stellate cells that control the firing of superficial PCs, as will be studied in a simplified manner in chapter 4.

Chapter 2

Linear response theory

The main objective of this chapter is to introduce the linear response theory or LRT, a framework in which the response of neural networks to various stimulation paradigms can be addressed analytically. We do this in incremental steps.

First, we present in detail the spiking neuron model that will be used throughout this document: the leaky integrate-and-fire neuron with diffusive noise. Second, we analyze the properties of its spiking activity, first in the time domain and then in Fourier space. This is mandatory since the LRT is best exploited in the frequency domain. Third, we characterize the response of the neuron model to external inputs. We then use this information to study oscillations in a fully recurrent network built with these neurons. This will allow us to discuss the limits of so-called mean-field theories of neural dynamics. The LRT will then be introduced as a flexible method superseding the mean-field theories by allowing a description of both the collective response of neural networks to stimulation and the single-neuron spiking activity. To lay the groundwork for the next chapter, we then apply the LRT to study gamma oscillations in a fully recurrent network loosely based on the ELL circuitry and subjected to spatially correlated signals mimicking global amplitude modulations of the EOD.

2.1 Integrate-and-fire models

As mentioned in the previous chapter, one often overcomes the complexity of neuronal dynamics by relying on simplified phenomenological models. One such class of models is the integrate-and-fire model. Various instances of this class exist, all involving an artificial reset mechanism whenever the membrane potential reaches a threshold V_T. The general subthreshold dynamics of this class is [92]

$$\tau_m \dot{V} = F(V) + R(V)I(t). \tag{2.1}$$

The simplest instance is the leaky integrate-and-fire neuron, with $F(V) = E_L - V$ and $R(V) = 1/\bar{g}_L \equiv R$, which has already been introduced in the previous chapter (*cf.* paragraph following Eq. 1.6). Another variant is the exponential integrate-and-fire, with $F(V) = E_L - V + \Delta_T \exp\left(\frac{V-\theta_T}{\Delta_T}\right)$ and $R(V) = R$ [39]. This model includes an exponential nonlinearity mimicking the onset of an action potential: whenever $V \gtrsim \theta_T$, the membrane potential rises abruptly with "sharpness parameter" Δ_T. In this case V_T is merely a computational parameter used to stop the exponential increase and reset the dynamics. Yet another example is the quadratic integrate-and-fire neuron [119], for which $F(V) = a_0(V - E_L)(V - \theta_T)$, with $a_0 > 0$. These two nonlinear integrate-and-fire models reproduce more accurately certain features of real neuronal dynamics, namely the existence of different effective thresholds in response to different types of stimulation [92]. The leaky integrate-and-fire neuron, on the other hand, has the same factitious threshold whether the stimulation $I(t)$ is a short current pulse or a step current. Finally, all these models fail to reproduce another important property of real neurons, adaptation, which has the potential to produce a whole assortment of firing patterns [151].

In the following, we shall only use the leaky integrate-and-fire (LIF) model (Fig. 2.1A). The aforementioned disadvantages are compensated by a higher analytic tractability. Some results of this thesis will nevertheless come out as being fairly independent of the chosen neuron model. Hence, the model that will be used for the remaining part of the thesis is

$$\tau_m \dot{V} = E_L - V + RI(t). \tag{2.2}$$

We express the total current $I(t)$ as

$$I(t) = I^{\text{syn}}(t) + I^{\text{ext}}(t) + I^{\text{noise}}(t), \tag{2.3}$$

where I^{syn} is the total synaptic current coming from other neurons within the network under

consideration. I^{noise} comprises noise sources intended to represent random background synaptic inputs from other neurons not part of the network *per se*, and possibly other noise sources. Finally, I^{ext} is the external input to be encoded by the neural network. It can be viewed as synaptic inputs not included in I^{syn} and I^{noise} that are meant to represent a meaningful sensory input. In the context of the weakly electric fish, it would be, for instance, the synaptic inputs coming from the electroreceptor afferents. The next two sections describe the specific form that the components I^{syn} and I^{noise} take herein.

2.2 Current-based synaptic inputs

In the first chapter we gave a conductance-based model for the synaptic interaction between presynaptic neuron j and postsynaptic neuron i, namely (*cf.* Eq. 1.3)

$$I_{ij}(t) = \sum_f \bar{g} s(t - t_j^{(f)})[E - V_i(t)]. \tag{2.4}$$

where E was the appropriate reversal potential (a characteristic constant for the synapse), $\bar{g}$ the maximal conductance for the synaptic channel and $s(t)$ the specific temporal evolution of the synaptic conductance, whose general formulation was given by Eq. 1.4. For our purpose, it is simpler to consider so-called current-based synapses (Fig. 2.1B). If we assume that V_i does not vary much about its resting state, we can replace the synaptic drive $E - V_i$ by a constant and write

$$I_{ij}(t) = W_{ij} \sum_f \alpha(t - t_j^{(f)}), \tag{2.5}$$

where W_{ij} is the *synaptic weight* between presynaptic neuron j and postsynaptic neuron i. In what follows the synaptic function will typically be

$$\alpha(t) = \left(\frac{t - \tau_d}{\tau_s^2}\right) \exp\left(-\frac{t - \tau_d}{\tau_s}\right) \Theta(t - \tau_d), \tag{2.6}$$

which is a delayed α function; the delay τ_d takes into account the axonal propagation time of action potentials. The expression above corresponds, modulo a constant, to the limit when $\tau_{\text{rise}} \to \tau_{\text{decay}} \equiv \tau_s$ in Eq. 1.4. To be able to compare synaptic functions with different time constant τ_s, the α function is normalized in such a way that its integral over time is 1: $\int_{-\infty}^{\infty} \alpha(t)dt = 1$. Hence, W_{ij} is the total charge transferred to neuron i during a synaptic event, whatever the value of τ_s. The sign of W_{ij} depends on the type of synapse: for excitatory synapses, $\bar{g}(E - V_i)$ is typically greater than zero so that $W_{ij} > 0$; for inhibitory synapses, $\bar{g}(E - V_i)$ is typically smaller than zero so that $W_{ij} < 0$. Finally, using the definition of a

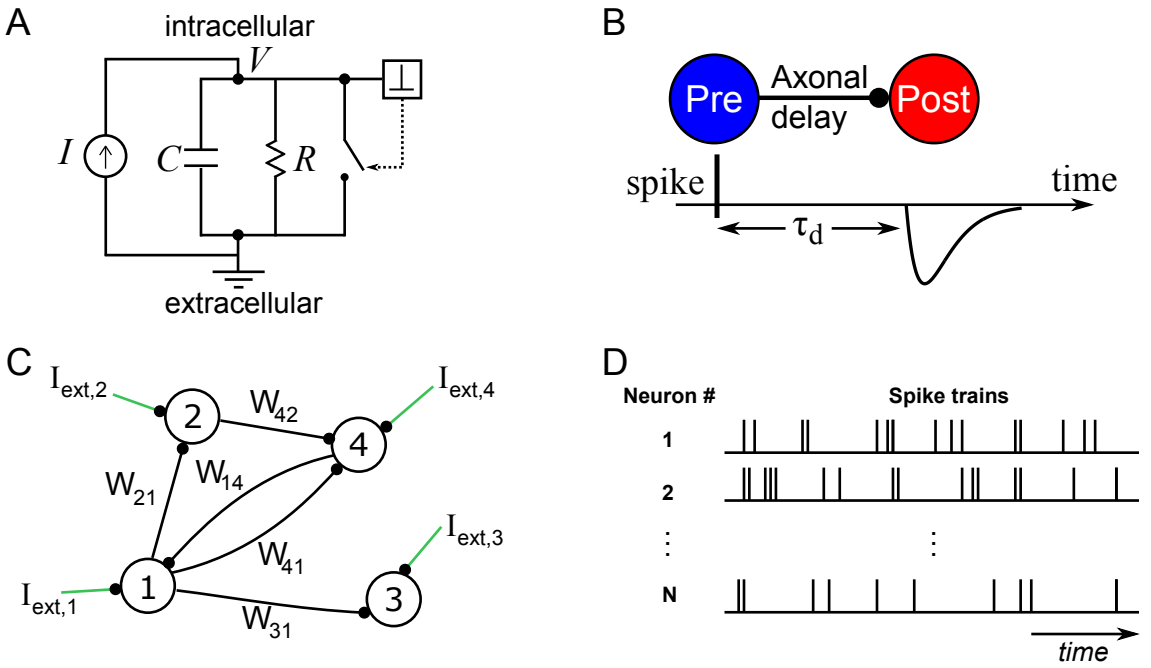

Figure 2.1: Network of LIF neurons.
(A) Dynamics of a leaky integrate-and-fire (LIF) neuron (*cf.* Eq. 2.2). The neuron receives a total current I which is the sum of an external input I^{ext}, a noise current I^{noise} and a synaptic current I^{syn}. The mathematical form of I^{ext} depends on the specific case. The subthreshold dynamics is that of a parallel RC circuit. At every moment the value of the membrane potential V is compared to a threshold (represented by the inverted T). When the threshold is reached, a spike is fired. When this occurs, the switch is closed for a small time interval, known as the absolute refractory period. This period accounts for the action potential generation during which the cell is unresponsive to stimulation (hence the short-circuit).
(B) Current-based synaptic dynamics (Eq. 2.5). Whenever a spike is fired by a presynaptic neuron, all its postsynaptic targets receive a synaptic current (here, inhibitory) after an axonal propagation delay τ_d.
(C) Neural network. Only synapses with nonzero synaptic weights are represented and the noise signal is implicit. Each neuron receives an external input I_{ext} coming from neurons not part of the network. Each neuron is also subjected to noise.
(D) Activity of the network. The output of every neuron of the network is a spike train giving the times at which it spiked.

spike train, $y_j = \sum_f \delta(t - t_j^{(f)})$, we can write the total synaptic input impinging on neuron i as

$$I_i^{\text{syn}}(t) = \sum_j W_{ij}(\alpha * y_j)(t), \tag{2.7}$$

where $*$ denotes a convolution. The W_{ij}'s are viewed as elements of a *connectivity matrix* $\mathbf{W}$. This connectivity matrix is related to the binary connectivity matrix briefly mentioned in section 1.3.2, except that it includes synaptic weights between pairs of neurons; absence of a connection implies a vanishing synaptic weight.

2.3 Noise

The simplest way to add noise in the LIF dynamics is to add a Gaussian white noise (GWN) current:

$$I^{\text{noise}}(t) \equiv \frac{\sigma}{R}\sqrt{\tau_m}\xi(t). \tag{2.8}$$

The factor $\sqrt{\tau_m}/R$ in I^{noise} gives σ units of voltage. The average value of this noise current is zero, $\langle I^{\text{noise}}(t)\rangle \equiv 0$, where the brackets mean an ensemble average over realizations of the noise. The autocorrelation function of I^{noise} is

$$\langle I^{\text{noise}}(t) I^{\text{noise}}(t+\tau)\rangle = \frac{\sigma^2}{R^2}\tau_m\langle\xi(t)\xi(t+\tau)\rangle = \frac{\sigma^2}{R^2}\tau_m\delta(\tau). \tag{2.9}$$

We further assume that the noise is uncorrelated from one neuron to the other. This means that for $i \neq j$, $\langle I_i^{\text{noise}}(t) I_j^{\text{noise}}(t+\tau)\rangle \equiv 0$ for all t and τ.

In the following, $I^{\text{noise}}(t)$ will be interpreted as an *intrinsic* noise (for instance, fast channel noise). However, other origins are possible for this noise source. A typical extrinsic cause for this noise is a synaptic bombardment by excitatory and inhibitory neurons external to the network—i.e., not connected through $\mathbf{W}$—that are using very fast synaptic dynamics; this was the type of noise appearing in subsections 1.3.2 and 1.3.3 [42, 87].

To summarize, the following neural network dynamics will be used throughout this thesis:

$$\begin{aligned} &\tau_m\frac{dV_i}{dt} = E_L - V_i(t) + \sigma\sqrt{\tau_m}\xi_i(t) + R\sum_{j=1}^{N} W_{ij}(\alpha * y_j)(t) + RI_i^{\text{ext}}(t), \\ &V_i(t^-) = V_T \Rightarrow t_i^{(f)} = t \text{ and } V_i(t) = V_R,\ \forall t \in (t^+, t^+ + \tau_r), \end{aligned} \tag{2.10}$$

where the last line spells out the reset mechanism underlying spike emission. In particular, the last part means that the membrane potential is clamped at V_R for a time τ_r after a spike,

where τ_r is the absolute refractory period. See Fig. 2.1 for a graphical summary of the network dynamics.

2.3.1 Nondimensionalization

To simplify the analysis of the neural dynamics, we will use nondimensional quantities. We seek a transformation of V such that the new potential v has the property that $v_T = 1$ and $v_R = 0$. That is, we require $v(V_T) = 1$ and $v(V_R) = 0$ through a linear transformation $v(V) = aV + b$. We readily find

$$v_i = \frac{V_i - V_R}{V_T - V_R}. \tag{2.11}$$

With this transformation and the substitution $\hat{t} = t/\tau_m$, Eq. 2.10 becomes

$$\begin{aligned}\frac{dv_i}{d\hat{t}}(\tau_m\hat{t}) = \frac{E_L - V_R}{V_T - V_R} - v_i(\tau_m\hat{t}) + \frac{\sigma}{V_T - V_R}\sqrt{\tau_m}\xi_i(\tau_m\hat{t}) \\ + \frac{R}{V_T - V_R}\left[\sum_{j=1}^{N} W_{ij}(\alpha * y_j)(\tau_m\hat{t}) + I_i^{\text{ext}}(\tau_m\hat{t})\right].\end{aligned}$$

We define:

$$\begin{aligned}\mu &\equiv \frac{E_L - V_R}{V_T - V_R}, \qquad & \hat{v}(\hat{t}) &\equiv v(\tau_m\hat{t}) \\ \hat{\sigma} &\equiv \frac{\sigma}{V_T - V_R}, \qquad & \hat{I}^{\text{ext}}(\hat{t}) &\equiv \frac{R}{V_T - V_R} I^{\text{ext}}(\tau_m\hat{t}).\end{aligned}$$

We also have for the synaptic current term, focusing on one term of the sum:

$$\begin{aligned}W_{ij}(\alpha * y_j)(\tau_m\hat{t}) &= W_{ij}\int_{-\infty}^{\infty}\alpha(\tau)y_j(\tau_m\hat{t} - \tau)d\tau \\ &= W_{ij}\sum_f \int_{-\infty}^{\infty}\alpha(\tau)\delta\left(\hat{t} - \frac{\tau}{\tau_m} - \frac{t_j^{(f)}}{\tau_m}\right)d\left(\frac{\tau}{\tau_m}\right).\end{aligned}$$

Changing the integration variable to $\hat{\tau} = \tau/\tau_m$ and defining $\hat{t}_j^{(f)} \equiv t_j^{(f)}/\tau_m$ yields

$$W_{ij}(\alpha * y_j)(\tau_m\hat{t}) = W_{ij}\sum_f \int_{-\infty}^{\infty}\alpha(\tau_m\hat{\tau})\delta\left(\hat{t} - \hat{\tau} - \hat{t}_j^{(f)}\right)d\hat{\tau}.$$

Since our normalization of α is such that its integral is equal to unity, we have

$$\int_{-\infty}^{\infty}\tau_m\alpha(\tau_m\hat{\tau})d\hat{\tau} = 1. \tag{2.12}$$

We thus define $\tau_m \alpha(\tau_m \hat{\tau}) \equiv \hat{\alpha}(\hat{\tau})$ and

$$\hat{W}_{ij} \equiv \frac{R}{V_T - V_R} \frac{W_{ij}}{\tau_m}. \tag{2.13}$$

Also, since

$$\langle \xi_i(t) \xi_i(t') \rangle = \langle \xi_i(\tau_m \hat{t}) \xi_i(\tau_m \hat{t}') \rangle = \frac{1}{\tau_m} \delta(\hat{t} - \hat{t}'),$$

we define

$$\hat{\xi}_i(\hat{t}) \equiv \sqrt{\tau_m} \xi_i(\tau_m \hat{t}). \tag{2.14}$$

Hence, removing all hats, we finally get:

$$\boxed{\frac{dv_i}{dt} = \mu - v_i(t) + \sigma \xi_i(t) + \sum_{j=1}^{N} W_{ij} (\alpha * y_j)(t) + I_i^{\text{ext}}(t).} \tag{2.15}$$

with $v_R = 0$, $v_T = 1$ and

$$\alpha(t) = \left(\frac{t - \tau_d}{\tau_s^2} \right) \exp\left(-\frac{t - \tau_d}{\tau_s} \right) \Theta(t - \tau_d). \tag{2.16}$$

All quantities with units of time are measured in multiples of the membrane time constant τ_m. μ will be called a *bias*. Sometimes we will put $\sigma = \sqrt{2D}$ for a reason that will be made clear in the next section. Because of this nondimensionalization, only three parameters determine a LIF neuron's response to synaptic and external inputs, namely τ_r, μ, σ. In the next section we study the LIF dynamics in absence of synaptic and external inputs.

2.4 The LIF neuron and its properties

We study the spiking dynamics of a single noisy LIF neuron subjected to constant stimulation I_0, i.e.,

$$\frac{dv}{dt} = \mu - v(t) + \sigma \xi(t) \tag{2.17}$$

where the constant input has been included in μ. In particular, an analytical expression for its firing rate will be given.

Equation 2.17 is a stochastic differential equation called the Langevin equation. It can be integrated—in the Stratonovich sense [153]—with initial condition $v(\tau_r) = 0$ to give:

$$v(t) = \mu \left[1 - e^{-(t - \tau_r)} \right] + \int_{\tau_r}^{t} \sigma \xi(t') e^{-(t - t')} dt'. \tag{2.18}$$

Without the spiking mechanism, the average value of $v(t)$ follows $\langle v(t)\rangle = \mu[1 - e^{-(t-\tau_r)}]$, whereas its variance is

$$\sigma_v^2(t) = \left\langle [v - \langle v(t)\rangle]^2 \right\rangle = \frac{\sigma^2}{2}[1 - e^{-2(t-\tau_r)}], \tag{2.19}$$

using $\langle \xi(t)\xi(t')\rangle = \delta(t - t')$. Note that $\langle v(t)\rangle \to \mu$ and $\sigma_v^2(t) \to \sigma^2/2$ when $t \to \infty$. Because the fluctuations of $v(t)$ with respect to its mean are given by the integration of a Gaussian process $\xi(t)$, $v(t)$ is normally distributed with mean and variance given by the expressions just obtained. Equivalently, the probability density that the potential takes the value v at time t given that it started at $v = 0$ at $t = \tau_r$, denoted $p(v, t)$, solves the partial differential equation

$$\frac{\partial p}{\partial t} = \frac{\partial}{\partial v}[(\mu - v)p] + \frac{\sigma^2}{2}\frac{\partial^2 p}{\partial v^2} \tag{2.20}$$

with initial condition $p(v, \tau_r) = \delta(v)$. This is the Fokker-Planck equation associated with Eq. 2.17 [180]. The prefactor of the second derivative is the diffusion coefficient $D = \sigma^2/2$, hence the notation $\sigma = \sqrt{2D}$ mentioned previously.

With the spiking mechanism switched on, the potential is reset to $v_R = 0$ after the threshold has been reached. To find the firing rate of the LIF neuron, one needs to find the average time that it takes for the potential to reach v_T, starting at v_R: a first-passage-time problem. If $\langle T_1\rangle$ is the average first-passage time, then the firing is $r = (\tau_r + \langle T_1\rangle)^{-1}$ (compare with section 1.2.1). Numerous authors have solved (and re-solved) this problem over time, albeit Siegert [195] and Ricciardi [50, 178] are usually given credit for the solution. We give the solution appearing in [126]:

$$r = \nu(\mu, \sigma; \tau_r) = \left(\tau_r + \sqrt{\pi} \int_{\frac{\mu - v_T}{\sigma}}^{\frac{\mu - v_R}{\sigma}} e^{z^2} \mathrm{erfc}(z) dz \right)^{-1} \tag{2.21}$$

with

$$\mathrm{erfc}(z) = 1 - \mathrm{erf}(z) = \frac{2}{\sqrt{\pi}} \int_x^\infty e^{-t^2} dt$$

the complementary error function. Figure 2.2 illustrates the effect of the noise intensity on the firing rate. When $\sigma = 0$ (blue curve), the neuron must be in the suprathreshold regime $\mu > v_T$ to fire, as previously noted in section 1.2.1. Voltage fluctuations due to noise allow a neuron to fire in the subthreshold regime $\mu < v_T$. Noise tends to linearize the f-I curve around threshold. Far from the threshold in the suprathreshold regime, noise has little effect on the firing rate because it is the deterministic part of the dynamics that matters.

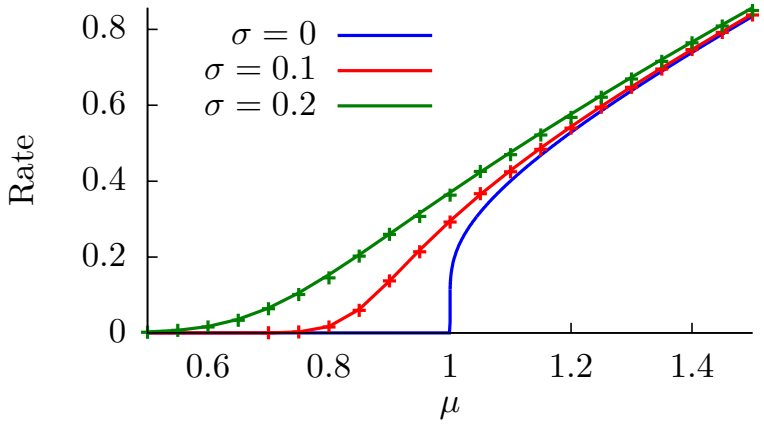

Figure 2.2: Effect of noise on f-I curves. Due to noise, the rate increases in the vicinity of the threshold at $v_T = 1$. The impact of noise on firing rate is small in the suprathreshold regime $\mu > v_T$. Solid lines are results from analytical expressions. Symbols were obtained by simulating the LIF dynamics—as explained in section 2.4.1—and evaluating the rate following the first definition in section 1.1.3.

2.4.1 Spike-train analysis in the time domain

The firing rate is only one statistic that can be utilized to understand the spiking process. The probability distribution of first-passage time—and therefore, the probability density function of ISIs, P_{ISI}—is not known in closed form, but an expression for its characteristic function has been obtained (see below). It is quite simple, however, to do without analytics and simulate the LIF dynamics to obtain an approximation of P_{ISI}, as depicted in Fig. 2.3A. To simulate the noisy LIF dynamics, we can use the Euler-Maruyama scheme:

$$v(t + dt) = \mu dt + (1 - dt)v(t) + \sigma\sqrt{dt}\mathcal{N},$$

where $\mathcal{N}$ stands for a normally distributed random variable with mean 0 and variance 1; at each time step a different sample is drawn (see appendix A.1 for details). When $v(t) < v_T$ and $v(t + dt) > v_T$, we say that a spike has been fired. We then clamp $v(t)$ to v_R for τ_r/dt time steps, implementing the absolute refractory period. Figure 2.4 illustrates these dynamics for five realizations. Every such realization yields a different spike train due to the presence of noise.

The average firing rate does not say anything about how variable the spiking process is. If the rate is, say, 0.2 (spikes per unit time), then we do not know whether the firing is periodic with period 5 or totally random. A useful measure to quantify spiking variability is

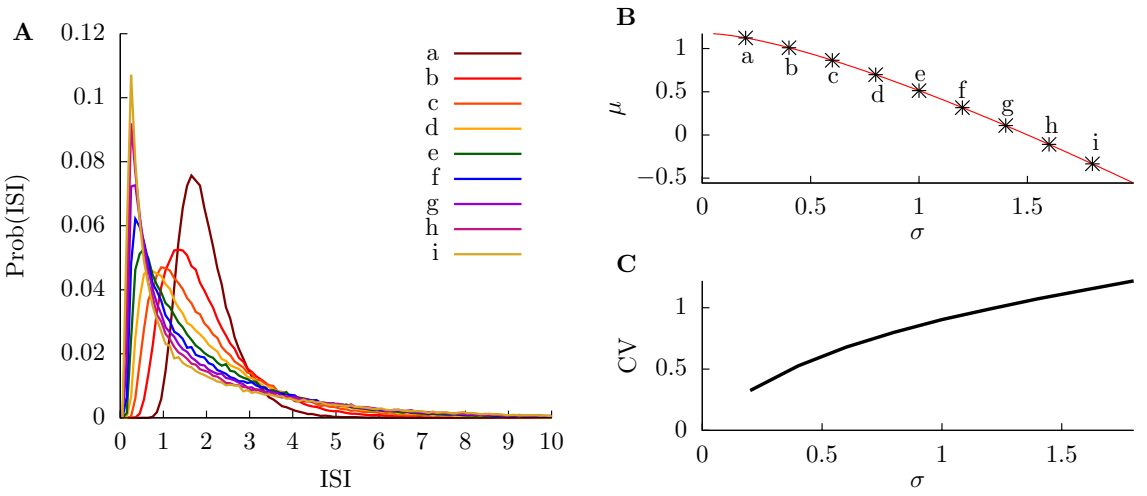

Figure 2.3: (A) Probability density function (pdf) of ISIs for the LIF neuron. The rate has been fixed to 0.5 and the refractory period τ_r to 0.1. We then found the subset of the parameter space (σ, μ) for which $r = \nu(\mu, \sigma; \tau_r) = 0.5$, generating the curve in panel B. We then picked a few points along this line (indicated by asterisks) and computed the corresponding ISI pdf (panel A) and coefficient of variation (panel C). As σ increases (and μ decreases) the ISI pdfs become more and more "Poisson-like". For large noise, the C_V actually becomes larger than that of the Poisson neuron with the same rate. In the more deterministic case ('a'), the ISI pdf has its maximum around $2 = 1/r$; hence, in this case the most probable ISI is the one corresponding to the inverse of the firing rate.

the *coefficient of variation* (C_V):

$$C_V = \frac{\sigma_{\text{ISI}}}{\langle \mathcal{T} \rangle}, \tag{2.22}$$

where $\sigma_{\text{ISI}} = \sqrt{\langle (\mathcal{T} - \langle \mathcal{T} \rangle)^2 \rangle}$ is the standard deviation of ISIs and $\langle \mathcal{T} \rangle = r^{-1}$ is the average ISI. For a perfectly periodic spiking process, $C_V = 0$; for a so-called Poisson spiking process, for which the spikes are uncorrelated, $C_V = 1$. For a Poisson spiking process, the probability that a spike occur in $[t, t + \Delta t)$ is given by $r\Delta t$ and the probability that two spikes occur is $o(\Delta t)$, with $\lim_{\Delta t \to 0} \frac{o(\Delta t)}{\Delta t} = 0$ [111]. With this definition, one can find [70]

$$P_{\text{ISI}}(\tau) = re^{-r\tau}, \tag{2.23}$$

and then $C_V = 1$. A Poisson process is highly variable because with each time increment there is a probability $r\Delta t$ to generate a spike, independently of what happened before. Other measures of variability exist, like the Fano factor [70], but we will only make use of the C_V in this thesis. In Figure 2.3A, we display the probability density functions (pdf) of ISIs of a LIF neuron with rate fixed at $r = 0.5$ and refractory period at $\tau_r = 0.1$, for various values of

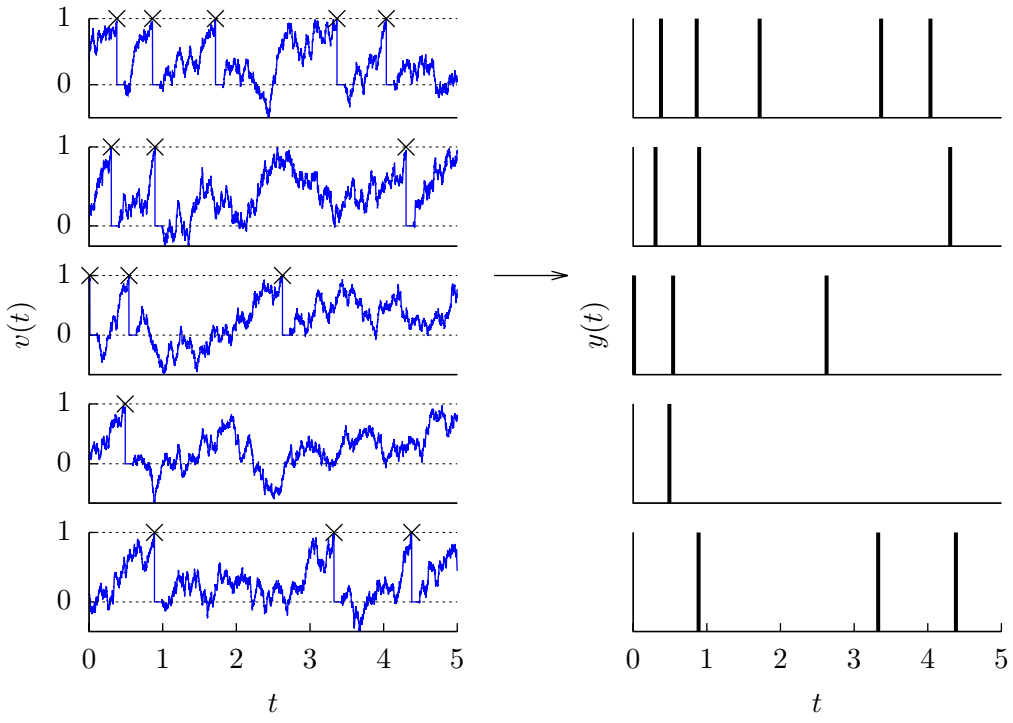

Figure 2.4: Dynamics of a single LIF neuron. The five panels on the left are five different realizations of Eq. 2.17. Blue traces are the voltages and crosses indicate spiking events. When the times of occurrence of these events are known, we can construct the spike trains on the right. Threshold and reset are indicated by dotted lines at $v = 1$ and $v = 0$, respectively. Parameters are: $dt = 10^{-3}$, $\mu = 0.9$, $\sigma = 0.8$ and $\tau_r = 0.1$.

σ and μ. The corresponding coefficients of variation appear in panel C of that figure; they increase monotonically with σ.

By construction, the LIF spiking process with constant input is a *renewal* spiking (point) process. This means that adjacent ISIs are statistically uncorrelated. Hence, if $T^{(i)}$ denotes the random variable associated with the ith ISI of the process, then $\langle T^{(i)} T^{(i+k)} \rangle = \langle T^{(i)} \rangle \langle T^{(i+k)} \rangle$ for all $k \geqslant 1$ and all i. This can be formalized by the correlation coefficient of ISIs:

$$\rho_{\mathrm{ISI}}(k) \equiv \frac{\left\langle \left(T^{(i+k)} - \langle T^{(i+k)} \rangle \right) \left(T^{(i)} - \langle T^{(i)} \rangle \right) \right\rangle}{\left\langle \left(T^{(i)} - \langle T^{(i)} \rangle \right)^2 \right\rangle}, \tag{2.24}$$

where $\langle T^{(i)} \rangle$ is the average ith ISI and $\langle (T^{(i)} - \langle T^{(i)} \rangle)^2 \rangle$ is the variance of the ith ISI. For renewal processes, $\rho_{\mathrm{ISI}}(0) = 1$ and $\rho_{\mathrm{ISI}}(k) = 0$ for all $k > 0$. For a stationary point process, ρ_{ISI} depends only on k, not i, which was implicitly assumed in the definition above. LIF neurons embedded in an oscillating recurrent network and neuron with slow adaptation

are examples of non-renewal spiking processes.

Finally, another useful statistic is the autocovariance function (ACF) of the spike train (sometimes also called autocorrelation function):

$$Q(\tau) \equiv \langle y(t+\tau)y(t)\rangle - r^2, \tag{2.25}$$

where $r = \langle y(t)\rangle$ is the stationary firing rate of the neuron. For a sufficiently small time interval, r^2 is the probability density that a spike occurs at $t+\tau$ and at t by chance (i.e., solely due to the firing rate). Hence, the ACF represents the probability density with respect to chance that a spike occurs at $t+\tau$ given that a spike occurred at t. For a purely periodic spike train, $Q(\tau) \propto \sum_n \delta(\tau - \frac{n}{r})$; for Poisson spiking, $Q(\tau) = r\delta(\tau)$, because spikes are uncorrelated.

2.5 Fourier analysis of spike trains

Spike trains and their statistics give temporal information about a neuron's activity. Another approach is to describe this activity and the resulting statistics in the frequency space *via* Fourier analysis. In this section, we first introduce the concept of the power spectrum or power spectral density (PSD) of a stochastic process. It represents the relative strength of the frequency components of the process. We then apply this definition to spike trains.

Let $X(t)$ be a stochastic process defined on $t \geq 0$ (not necessarily a spiking process). Also, let $X_T(t)$ be the restriction of $X(t)$ to the interval $[0, T]$, i.e.,

$$X_T(t) \equiv \begin{cases} X(t) & \text{if } 0 \leq t \leq T \\ 0 & \text{otherwise.} \end{cases} \tag{2.26}$$

We further define the Fourier operator, $\mathscr{F}$, and its inverse, $\mathscr{F}^{-1}$, through

$$\tilde{x}(\omega) \equiv \mathscr{F}\{x\}(\omega) \equiv \int_{-\infty}^{\infty} x(t)e^{i\omega t}dt$$
$$x(t) \equiv \mathscr{F}^{-1}\{\tilde{x}\}(t) \equiv \int_{-\infty}^{\infty} \tilde{x}(\omega)e^{-i\omega t}\frac{d\omega}{2\pi}.$$

The power spectrum of $X(t)$ is defined by

$$S_X(\omega) \equiv \lim_{T\to\infty} \frac{\langle|\tilde{X}_T(\omega)|^2\rangle}{T}, \tag{2.27}$$

with $\tilde{X}_T(\omega) = \mathscr{F}\{X_T\}(\omega)$. A stochastic process $X(t)$ is said to be wide-sense stationary (WWS) if $\langle X(t)\rangle$ is constant for all $t \geq 0$ and if its autocorrelation function $\langle X(t)X(t')\rangle$ only depends on $|t' - t|$ [157]. By the Wiener-Khintchine[1] theorem, the PSD of such a process is the Fourier transform of its autocorrelation function [157]. Defining the autocorrelation by (compare with Eq. 2.25 above)

$$Q_X(\tau) \equiv \langle X(t)X(t+\tau)\rangle, \tag{2.28}$$

the PSD of X is

$$S_X(\omega) = \mathscr{F}\{Q_X\}(\omega). \tag{2.29}$$

Likewise, we can define the cross-spectrum or cross spectral density between two random processes $X(t)$ and $Y(t)$ by

$$S_{XY}(\omega) \equiv \lim_{T\to\infty} \frac{\langle \tilde{X}_T(\omega)\tilde{Y}_T^*(\omega)\rangle}{T}. \tag{2.30}$$

Note that $S_{XY}^*(\omega) = S_{YX}(\omega)$. If these two processes are jointly WSS—i.e., $\langle X(t)Y(t')\rangle$ only depends on $|t' - t|$—then by the Wiener-Khintchine theorem,

$$S_{XY}(\omega) = \mathscr{F}\{Q_{XY}\}(\omega). \tag{2.31}$$

As for spike train processes (Eq. 2.25), we can also define the autocovariance function of $X(t)$, $Q_{X-\bar{X}}(\tau)$, which is related to the autocorrelation function by

$$Q_{X-\bar{X}}(\tau) \equiv \langle [X(t) - \langle X(t)\rangle]\,[X(t+\tau) - \langle X(t+\tau)\rangle]\rangle = Q_X(\tau) - \langle X(t)\rangle^2,$$

where $\langle X(t)\rangle = m_X$ does not depend on t for a WSS process. A process with correlations on a finite time scale will have

$$\lim_{\tau\to\infty} Q_{X-\bar{X}}(\tau) = 0$$

or $\lim_{\tau\to\infty} Q_X(\tau) = \langle X(t)\rangle^2$. The PSDs corresponding to these two expressions are related by

$$S_X(\omega) = 2\pi\langle X(t)\rangle^2\delta(\omega) + S_{X-\bar{X}}(\omega). \tag{2.32}$$

Hence, only the DC component changes from one definition to the other. In the following, the latter does not typically appear in the results. $S_{X-\bar{X}}(\omega)$ is sometimes called the *noise*

[1] Khintchine is sometimes spelled Khinchin.

spectrum because it represents the amplitude of the Fourier components of the fluctuations of the process around its mean.

2.5.1 Spiking processes

The PSD of stationary spiking processes involves the characteristic functions of random variables representing interspike intervals of order σ. That is, for $\sigma = 1$, T_σ is the random variable representing the ISI as defined above; for $\sigma = 2$, T_σ represents the interval between a pair of spikes separated by two ISIs, and so on. If p_σ is the probability density function for T_σ and $\tilde{p}_\sigma$ its characteristic function, the PSD is [63]

$$S(\omega) = r\left[1 + 2\,\mathrm{Re}\left\{\sum_{\sigma=1}^{\infty}\tilde{p}_\sigma(\omega)\right\}\right] = r\left[1 + \sum_{\sigma=1}^{\infty}\{\tilde{p}_\sigma(\omega) + \tilde{p}_\sigma(-\omega)\}\right]. \tag{2.33}$$

On the other hand, the autocorrelation is just the inverse Fourier transform of this expression, namely

$$Q(\tau) = r\left[\delta(\tau) + \sum_{\sigma=1}^{\infty}\{p_\sigma(\tau) + p_\sigma(-\tau)\}\right]. \tag{2.34}$$

We now apply these expressions to renewal processes and several important examples that will help understand the LIF neuron's PSD.

Renewal processes

There exists an obvious recursive relationship between the random variables T_σ: $T_\sigma = T_1 + T_{\sigma-1}$. For a renewal process, all intervals are independent and therefore the random variables T_1 and $T_{\sigma-1}$ are independent. Hence, the probability density function for T_σ can be expressed as a convolution:

$$p_\sigma(\tau) = \int_0^\tau p_1(t)p_{\sigma-1}(\tau - t)dt = (p_1 * p_{\sigma-1})(\tau), \tag{2.35}$$

or, equivalently, the characteristic functions are given by $\tilde{p}_\sigma(\omega) = \tilde{p}_1(\omega)\tilde{p}_{\sigma-1}(\omega) = \tilde{p}_1^\sigma(\omega)$. Since

$$|\tilde{p}_1(\omega)| = \left|\int p_1(t)e^{i\omega t}dt\right| \leq \int \left|p_1(t)e^{i\omega t}\right| dt = \int p_1(t)dt = 1, \tag{2.36}$$

and assuming that $S(\omega)$ diverges only at $\omega = 0$ (where $\tilde{p}_1(\omega) = 1$), the sum in Eq. 2.33 becomes

$$\sum_{\sigma=1}^{\infty}\tilde{p}_1^\sigma(\omega) = \tilde{p}_1(\omega)\sum_{\sigma=0}^{\infty}\tilde{p}_1^\sigma(\omega) = \frac{\tilde{p}_1(\omega)}{1 - \tilde{p}_1(\omega)}, \tag{2.37}$$

using the well-known formula for geometric series. Thus, the PSD is

$$S(\omega) = r\left[1 + 2\,\mathrm{Re}\,\left\{\frac{\tilde{p}_1(\omega)}{1-\tilde{p}_1(\omega)}\right\}\right] = r\,\mathrm{Re}\,\left[\frac{1+\tilde{p}_1(\omega)}{1-\tilde{p}_1(\omega)}\right] = r\left[\frac{1-|\tilde{p}_1(\omega)|^2}{|1-\tilde{p}_1(\omega)|^2}\right]. \tag{2.38}$$

Examples

We present two extreme cases of spiking processes to illustrate the concept of power spectrum.

Poisson process For a Poisson process, $p_1(\tau) = re^{-r\tau}$ so that $\tilde{p}_1(\omega) = \frac{r}{r-i\omega}$. Then,

$$S(\omega) = r\,\mathrm{Re}\,\left[\frac{1+\frac{r}{r-i\omega}}{1-\frac{r}{r-i\omega}}\right] = r.$$

To show the effect of refractoriness on PSDs, we consider a Poisson process X_2 with rate r and refractory period τ_r, and compare it to another Poisson process X_1 with the same rate but without refraction [91]. For X_2, the random variable T_1 can be seen as a sum of a deterministic number representing the refractory period, plus another random variable $\hat{T}_1$ whose average value is $1/r - \tau_r$, and whose pdf is $p_{\hat{T}_1}(\tau) = \nu e^{-\nu\tau}$ (ν is unknown for now). Therefore,

$$p_1(\tau) = (p_{\mathrm{ref}} * p_{\hat{T}_1})(\tau) = \int_{-\infty}^{\tau} \delta(t-\tau_r)p_{\hat{T}_1}(\tau - t)dt = p_{\hat{T}_1}(\tau - \tau_r). \tag{2.39}$$

It means that $p_1(\tau) = \nu e^{-\nu(\tau-\tau_r)}\Theta(\tau-\tau_r)$. The requirement that $1/r = \int \tau p_1(\tau)d\tau$ implies $\frac{1}{r} = \tau_r + \frac{1}{\nu}$, i.e.

$$\nu = \frac{r}{1-\tau_r r}. \tag{2.40}$$

The characteristic function of $p_1(\tau)$ is

$$\tilde{p}_1(\omega) = \int_{\tau_r}^{\infty} \nu e^{-\nu(\tau-\tau_r)}e^{i\omega\tau}d\tau = e^{i\omega\tau_r}\int_0^{\infty} \nu e^{-\nu t}e^{i\omega t}dt = \frac{\nu}{\nu - i\omega}e^{i\omega\tau_r} \tag{2.41}$$

and

$$S(\omega) = \frac{r}{1 + 2\frac{\nu^2}{\omega^2}[1-\cos(\omega\tau_r)] + 2\frac{\nu}{\omega}\sin(\omega\tau_r)}. \tag{2.42}$$

For small frequencies, $\cos(\omega\tau_r) \approx 1 - \omega^2\tau_r^2/2$ and $\sin(\omega\tau_r) \approx \omega\tau_r$, so that

$$S(\omega) = r(1-\tau_r r)^2 + \mathcal{O}(\omega^2). \tag{2.43}$$

Hence, given that $\tau_r r < 1$, refractoriness has the effect of decreasing the PSD for low

frequency with respect to a refractoriless process with the same rate.

Tonic (periodic) spiking We now consider regular spiking with period P, so that $r = 1/P$. Obviously, $p_\sigma(\tau) = \delta(\tau - \sigma P)$ and $\tilde{p}_\sigma(\omega) = e^{i\omega\sigma P}$, so that

$$S(\omega) = r\left[1 + 2\,\mathrm{Re}\left\{\sum_{\sigma=1}^{\infty} e^{i\omega\sigma P}\right\}\right]. \tag{2.44}$$

Since the firing pattern is a so-called Dirac comb, $\sum_n \delta(t - nP)$, whose Fourier transform is also a Dirac comb, we expect the power spectrum to be a Dirac comb as well. Indeed, using the identity,

$$\sum_{n=-\infty}^{\infty} \delta(t - nP) = \frac{1}{P}\sum_{n=-\infty}^{\infty} e^{i2\pi nt/P} = \frac{1}{P} + \frac{2}{P}\,\mathrm{Re}\sum_{n=1}^{\infty} e^{i2\pi nt/P} \tag{2.45}$$

we get

$$S(\omega) = \frac{2\pi}{P^2}\sum_{n=-\infty}^{\infty} \delta\left(\omega - \frac{2\pi n}{P}\right) = 2\pi r^2 \sum_{n=-\infty}^{\infty} \delta\left(\omega - 2\pi r n\right). \tag{2.46}$$

Leaky integrate-and-fire neuron

A third and more important example is the noisy LIF—it deserves a subsection on its own. It has been proven repeatedly that for a noisy LIF obeying Eq. 2.17 with refractory period τ_r, the characteristic function of the first-order ISI distribution is [123]

$$\tilde{p}_1(\omega) = e^{i\omega\tau_r} e^{\Delta} \frac{\mathcal{D}_{i\omega}\left(\frac{\mu - v_R}{\sigma/\sqrt{2}}\right)}{\mathcal{D}_{i\omega}\left(\frac{\mu - v_T}{\sigma/\sqrt{2}}\right)} \tag{2.47}$$

where

$$\Delta = \frac{v_R^2 - v_T^2 + 2\mu(v_T - v_R)}{2\sigma^2} \tag{2.48}$$

and $\mathcal{D}_a(z)$ are parabolic cylinder functions[2]. Since the LIF spiking process is renewal, its

[2]Parabolic cylinder functions $\mathcal{D}_b(z)$ are solutions of the Weber differential equation. This second-order differential equation,

$$\frac{d^2y}{dx^2} - \left(\frac{x^2}{4} + a\right) y = 0,$$

has two independent solutions, $U(a,x)$ and $V(a,x)$. The parabolic cylinder function used above is related to $U(a,x)$ by $U(a,x) = \mathcal{D}_{-a-1/2}(x)$. Mathematica (Wolfram Research, Champaign, Illinois) and Maple™ (Maplesoft, Waterloo, Canada) have useful implementations of these functions. For the interested reader, see [154] or [3].

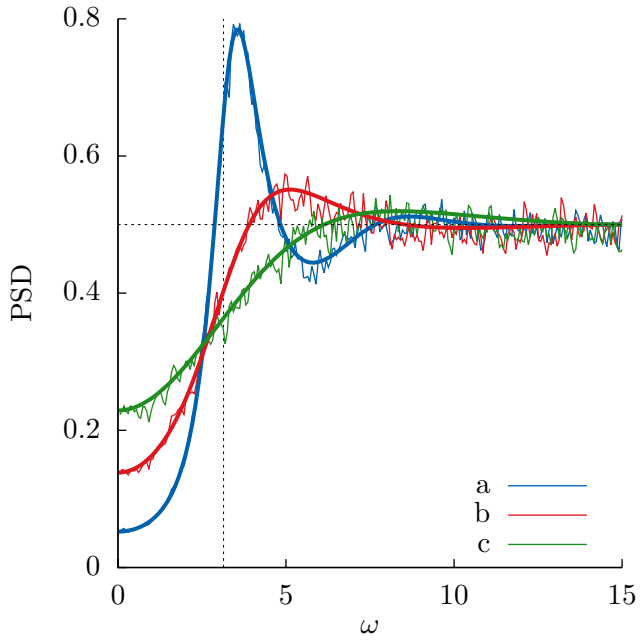

Figure 2.5: Power spectral densities (PSDs) for cases *a*, *b* and *c* in Fig. 2.3. Thick curves were computed from Eq. 2.49 (DC component has been removed); the thin wiggly lines come from simulations of the dynamics of LIF neurons, i.e., spike trains were obtained numerically and PSDs were then computed (see appendix A.2 for details). Notice that the asymptotic value of the PSDs is equal to the common rate $r = 0.5$ (dotted horizontal line). This comes from the zero-lag delta correlation in Eq. 2.34. The vertical dotted line denotes the angular frequency corresponding to the rate ($= 2\pi r$). Note that $S(\omega) = S(-\omega)$.

PSD is then, by Eq. 2.38,

$$S(\omega) = r\frac{\left|\mathcal{D}_{i\omega}\left(\frac{\mu-v_T}{\sigma/\sqrt{2}}\right)\right|^2 - e^{2\Delta}\left|\mathcal{D}_{i\omega}\left(\frac{\mu-v_R}{\sigma/\sqrt{2}}\right)\right|^2}{\left|\mathcal{D}_{i\omega}\left(\frac{\mu-v_T}{\sigma/\sqrt{2}}\right) - e^{\Delta}e^{i\omega\tau_r}\mathcal{D}_{i\omega}\left(\frac{\mu-v_R}{\sigma/\sqrt{2}}\right)\right|^2}. \tag{2.49}$$

The rate r in the expression above is computed from Eq. 2.21. Figure 2.5 shows three examples of PSDs associated with cases *a*, *b*, and *c* of Fig. 2.3. For case *a*, the neuron is in the suprathreshold regime with $\mu = 1.1234$. Hence, even without noise the neuron would spike. A small amount of noise, $\sigma = 0.2$ was added to reach the required rate of 0.5. This somewhat deterministic dynamic results in a PSD displaying a salient peak around the angular frequency corresponding to the rate ($= 2\pi r$), represented by the vertical dotted line. As the bias decreases and the noise intensity increases, the peak rapidly decreases in amplitude and moves toward higher frequencies. For case *c*, μ is slightly lower than the threshold. Noise is thus needed for the neuron to spike. The PSD does not show a clearly

noticeable peak.

In the suprathreshold regime, without noise, tonic spiking occurs and the ISI pdf is $p_1(\tau) = \delta(\tau - \frac{1}{r})$ with (see section 1.2.1)

$$r = \left[\tau_r + \ln\left(\frac{\mu - v_R}{\mu - v_T}\right)\right]^{-1} \Theta(\mu - v_T).$$

For fixed rate, this corresponds to the limit $\sigma \to 0$ in Fig. 2.3B along the curve μ versus σ. Counterintuitively, the opposite limit $\sigma \to \infty$ along this same curve does not yield a Poisson process with rate r and refractory period τ_r [155]. In fact, when μ lies well below reset, short ISIs become more probable than for the corresponding Poisson model with refraction and with the same rate. The spiking activity is then arranged in bursts, i.e. tight sequences of spikes occurring occasionally. In Fig. 2.3C, the coefficient of variation increases monotonically, becoming greater than $(1 - \tau_r r) = 0.95$, the CV of the Poisson model (see Eq. 2.43 and Eq. 2.53 below). We might summarize the situation by saying that, going from $\sigma \to 0$ to $\sigma \to \infty$ morphs the PSD shape from tonic to Poisson-like to burst-like.

2.5.2 Time scales and power spectra

Various time scales are at play in the dynamics of LIF neurons. First, there is the refractory period τ_r, whose effect on power spectra is to reduce the power at low frequency (see section 2.5.1). Intuitively, this is due to the fact that very short ISIs are prohibited by the refractoriness. Therefore, to get the same firing rate for a refractory neuron compared to a non-refractory one, longer intervals must be rarer—the tail of the ISI distribution is smaller—the effect of which is a reduced power at low frequency.

In the suprathreshold regime ($\mu > v_T$), the drift part of the subthreshold LIF dynamics, $\mu - v$, rapidly brings the membrane potential close to threshold after a spike has occurred. Hence, the predominant time scale in this case is the mean ISI, the average time that it takes for v to reach threshold. Noise spreads out the ISI distribution, which is centred around the average ISI (case *a* in Figs. 2.3 and 2.5). In the subthreshold (or noise-activated) regime, the drift part of the subthreshold LIF dynamics brings the membrane potential to $\mu < v_T$ in 1 unit of time, on average. If μ is near threshold, v will quickly cross v_T; the dominant time scale is then the membrane time scale, which is 1 in our nondimensional units. For case *c* in Fig. 2.3, $\mu = 0.8645 < v_T$, and the ISI pdf has its maximum around 1; in the PSD of Fig. 2.5, the maximum is slightly above 2π, as expected. As μ gets smaller with respect to v_T, spikes are caused more and more by noise and the drift dynamic ceases to be relevant.

2.5.3 Coefficient of variation and power spectrum

For renewal point processes, there exists a relationship between the coefficient of variation and PSDs. The power spectrum of a zero-mean renewal point process is given by Eq. 2.38 :

$$S(\omega) = r\,\mathrm{Re}\left[\frac{1+\tilde{p}_1(\omega)}{1-\tilde{p}_1(\omega)}\right] \tag{2.50}$$

An expansion of $\tilde{p}_1(\omega)$ for small ω gives

$$\tilde{p}_1(\omega) \approx \int p_1(\tau)\left(1 + i\omega\tau - \frac{\omega^2\tau^2}{2}\right)d\tau = 1 + i\omega\langle T_1\rangle - \frac{\omega^2\langle T_1^2\rangle}{2}. \tag{2.51}$$

Replacement in the above equation yields

$$S(\omega) = r\,\mathrm{Re}\left[\frac{2 + i\omega\langle T_1\rangle - \frac{\omega^2\langle T_1^2\rangle}{2}}{-i\omega\langle T_1\rangle + \frac{\omega^2\langle T_1^2\rangle}{2}}\right] = r\left[\frac{\omega^2\langle T_1^2\rangle - \omega^2\langle T_1\rangle^2 - \left(\frac{\omega^2\langle T_1^2\rangle}{2}\right)^2}{\omega^2\langle T_1\rangle^2 + \left(\frac{\omega^2\langle T_1^2\rangle}{2}\right)^2}\right].$$

Taking the limit $\omega \to 0$ yields

$$\lim_{\omega\to 0} S(\omega) = r\left[\frac{\langle T_1^2\rangle - \langle T_1\rangle^2}{\langle T_1\rangle^2}\right] = rC_V^2. \tag{2.52}$$

Also, since $\lim_{\omega\to\infty}\tilde{p}_1(\omega) = 0$, $\lim_{\omega\to\infty} S(\omega) = r$ (see Fig. 2.5), hence

$$C_V = \sqrt{\frac{\lim_{\omega\to 0} S(\omega)}{\lim_{\omega\to\infty} S(\omega)}}. \tag{2.53}$$

This fact clearly appears in Fig. 2.5: going from case *a* to *c*, the CV (Fig. 2.3C) increases as $S(0^+)$ increases, with $S(\infty)$ a constant.

2.6 Response of noisy LIF neurons to stimulation

We now consider the response of a LIF neuron to external stimulation. Its subthreshold dynamics is

$$\frac{dv}{dt} = \mu - v(t) + \sigma\xi(t) + I^{\mathrm{ext}}(t). \tag{2.54}$$

$I^{\mathrm{ext}}(t)$ can be a deterministic or stochastic (WSS) external signal. In the latter case, each sample function $I_k^{\mathrm{ext}}(t)$ of the external input will generate an ensemble of spike trains. The

expectation value of the spiking process with respect to this ensemble yields a sample-dependent firing rate:

$$r_k(t) = \langle y(t) \rangle_k, \tag{2.55}$$

where the subscript means that a given sample function for $I^{\text{ext}}(t)$ was chosen. Each presentation of sample function $I_k^{\text{ext}}(t)$ gives a different spike train—because of the "internal" noise $\sigma\xi(t)$. Averaging over all the spike trains generated with this sample function gives the rate above. We can further average over the ensemble of sample functions. Since the stochastic process is WSS, we expect this average to be independent of time:

$$r = \lim_{M\to\infty} \frac{1}{M} \sum_{k=1}^{M} r_k(t) = \langle y(t) \rangle. \tag{2.56}$$

In this case the brackets without subscript signify to take both the average over realizations of the internal noise $\sigma\xi(t)$ and of the external input process.

For deterministic signals, the average over the sample space of the external signal is irrelevant. Here, it helps to have in mind the general shape of a neuron's f-I curve (or r-μ curve for LIF neurons). From chapter 1 and Fig. 2.2, f-I curves typically have a sigmoidal shape. They go to zero from above as the input intensity is lowered, and saturate for large input intensities. A rectification thus occurs for small (possibly negative) inputs and saturation happens for large input values. Assuming that the mean value of $I^{\text{ext}}(t)$ is zero, that the input is weak in amplitude and that the value of μ makes the neuron lie at a point of the f-I curve not too close to the rectification or saturation, we can write the linear approximation

$$r(t) = r_0 + (A * I^{\text{ext}})(t) = r_0 + \int_{-\infty}^{\infty} A(\tau) I^{\text{ext}}(t-\tau) d\tau. \tag{2.57}$$

Here, $r_0 = \nu(\mu, \sigma; \tau_r)$ is the firing rate of the neuron without stimulation and A is the impulse response function or time-dependent transfer function of the neuron. The above expression can be recast in the Fourier representation thus

$$\tilde{r}(\omega) = 2\pi r_0 \delta(\omega) + \tilde{A}(\omega)\tilde{I}^{\text{ext}}(\omega). \tag{2.58}$$

To understand the meaning of $\tilde{A}(\omega)$[3], consider $I^{\text{ext}}(t) = a\cos(2\pi f t)$, a cosine with frequency f and amplitude a. Then,

$$r(t) = r_0 + a|\tilde{A}(2\pi f)| \cos(2\pi f t - \phi), \tag{2.59}$$

[3]Note that in the following we will generally omit the tilde over A. The context will determine whether it is $A(t)$ or $A(\omega)$.

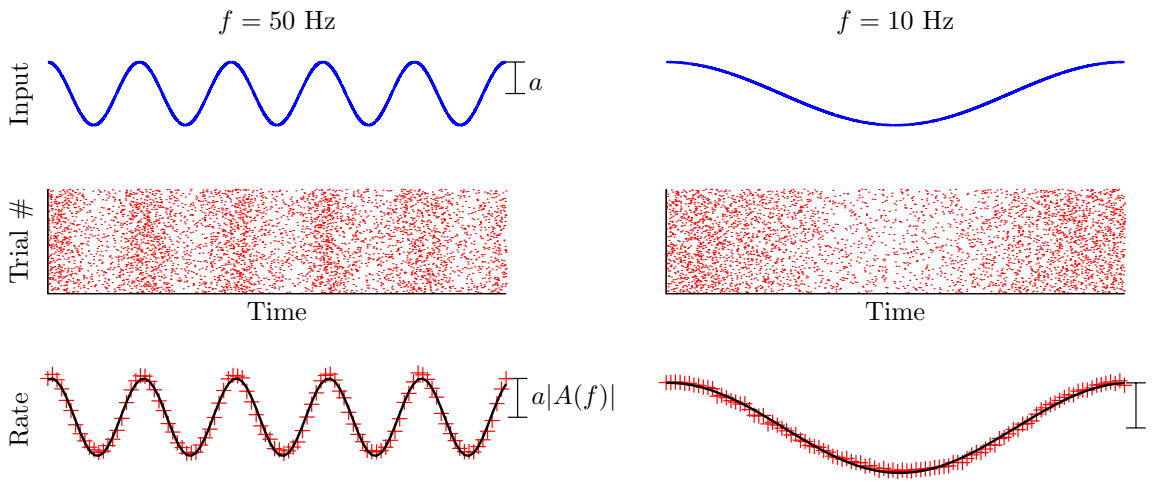

Figure 2.6: The meaning of susceptibility. For a given stimulation frequency (here $f = 50$ Hz and $f = 10$ Hz (in SI units)), we get many different spike trains. We then compute the firing rate $r(t) = \langle y(t) \rangle$. Notice that $|\tilde{A}(2\pi 10)| > |\tilde{A}(2\pi 50)|$. The phase ϕ corresponds to the shift of $r(t)$ with respect to the input.

where $\phi = \arg\{\tilde{A}(2\pi f)\}$. Therefore, $|\tilde{A}(2\pi f)|$ is numerically equal to the ratio between the amplitude of $r(t) - r_0$ and that of $I^{\text{ext}}(t)$, ϕ gives the phase shift of the rate response with respect to the input, and $\tilde{A}(2\pi f) = |\tilde{A}(2\pi f)|e^{i\phi}$.

In principle, we could determine all values of $\tilde{A}(\omega)$ by stimulating the neuron model (or, even, a real neuron) with sinusoidal inputs and compute the aforementioned ratio (see Fig. 2.6). However, this process would be tedious since it requires evaluation of $r(t)$ for numerous input signals. Rather, we can adopt the following approach called the reverse-correlation method [70]. In response to the k-th sample function of process I^{ext}, we compute the output firing rate in the Fourier domain as

$$\tilde{\Delta r}_{k,T}(\omega) = \tilde{A}(\omega)\tilde{I}^{\text{ext}}_{k,T}(\omega)$$

where $\tilde{\Delta r}_{k,T}(\omega) = \mathscr{F}\{\Delta r_{k,T}(t)\}$ and $\Delta r_{k,T}(t) = r_{k,T}(t) - r_0$ (as above, subscript $_T$ means that the function is restricted to $[0, T]$). We then calculate the cross-spectrum between I^{ext} and r, as well as the PSD of I^{ext}, i.e.

$$S_{rI^{\text{ext}}}(\omega) = \lim_{T\to\infty} \frac{\langle \tilde{\Delta r}_T(\omega)\tilde{I}^{\text{ext}}_T(\omega)\rangle}{T} \quad \text{and} \quad S_{\text{ext}}(\omega) \equiv \lim_{T\to\infty} \frac{\langle |\tilde{I}^{\text{ext}}_T(\omega)|^2\rangle}{T}. \tag{2.60}$$

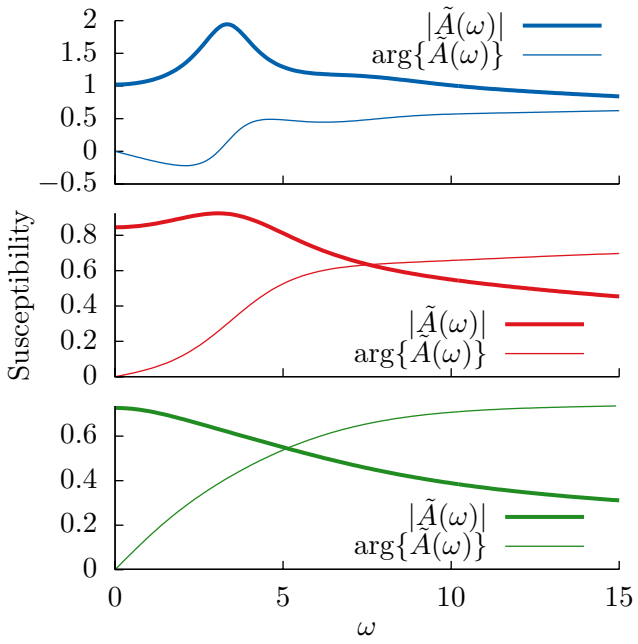

Figure 2.7: Susceptibility of LIF neurons for cases *a*, *b* and *c* in Figs. 2.3 and 2.5. The bias μ decreases from *a* to *c* while the noise intensity σ increases. We represent the amplitude and phase of the susceptibility, with $\tilde{A}(\omega) = |\tilde{A}(\omega)| \exp(i\arg\{\tilde{A}(\omega)\})$. Cases *a* (top) and *b* (middle) show a preferred frequency, but not case *c* (bottom). $\tilde{A}(\omega)$ was evaluated from Eq. 2.62. Note that $|\tilde{A}(\omega)| = |\tilde{A}(-\omega)|$ and $\arg\{\tilde{A}(\omega)\} = -\arg\{\tilde{A}(-\omega)\}$.

The susceptibility is thus given by

$$\tilde{A}(\omega) = \frac{S_{rI^{\text{ext}}}(\omega)}{S_{\text{ext}}(\omega)}. \tag{2.61}$$

This method can be applied to all kinds of model and real neurons, as long as the operating point of the neuron is not too close to nonlinearities such as saturation or rectification.

For LIF neurons, an analytical expression for $\tilde{A}(\omega)$ exists [123, 126, 125]. It is again given in terms of parabolic cylinder functions:

$$\tilde{A}(\omega) = \left(\frac{i\omega r\sqrt{2}}{\sigma(i\omega-1)}\right)\left[\frac{\mathcal{D}_{i\omega-1}(\frac{\mu-v_T}{\sigma/\sqrt{2}}) - e^{\Delta}\mathcal{D}_{i\omega-1}(\frac{\mu-v_R}{\sigma/\sqrt{2}})}{\mathcal{D}_{i\omega}(\frac{\mu-v_T}{\sigma/\sqrt{2}}) - e^{\Delta+i\omega\tau_r}\mathcal{D}_{i\omega}(\frac{\mu-v_R}{\sigma/\sqrt{2}})}\right], \tag{2.62}$$

with Δ as defined above. See Fig. 2.7 for examples.

2.7 Oscillations in a fully recurrent network of LIF neurons

Susceptibility does not only quantify the response of a neuron to external inputs, but also to fluctuations of the synaptic inputs coming from within the network. In this section, we show, following the treatment of Brunel and Hansel [43], how susceptibility dictates the occurrence of oscillations in a fully recurrent network with delayed synaptic interactions. The treatment is analogous to that of section 1.3.2.

We start with a fully recurrent inhibitory network containing N identical neurons. Following [43], we assume that every neuron in the network is connected to every other neuron and itself. Furthermore, each neuron receives excitatory white noise that is uncorrelated across neurons in the network. The subthreshold membrane dynamics of neuron i is:

$$\dot{v}_i = \mu - v_i(t) + I^{\text{syn}}(t) + \sigma\xi_i(t), \tag{2.63}$$

where

$$I^{\text{syn}}(t) = -\frac{J}{N}\sum_{k=1}^{N}(s * y_k)(t),$$

with s the synaptic function (see Eq. 2.7). Note that the total inhibitory current, $I^{\text{syn}}(t)$, does not need an index i because it is the same for all neurons. The synaptic coupling is chosen to be identical for all neural pairs and set equal to $-J/N$ with $J > 0$. Division by N ensures that the synaptic current stays finite in the limit $N \to \infty$.

Brunel and Hansel [43] use the synaptic function

$$s(t) = A\left[\exp\left(-\frac{t-\tau_d}{\tau_s}\right) - \exp\left(-\frac{t-\tau_d}{\tau_r}\right)\right]\Theta(t-\tau_d) \tag{2.64}$$

with the constant $A = 1/(\tau_s - \tau_r)$ such that $\int s(t)dt = 1$ (in this section τ_r is the synaptic rise time, not the refractory period) . It is easy to show (*cf.* section A) that $s(t)$ is a solution of the system

$$\tau_s\dot{s}(t) + s(t) = x(t)$$
$$\tau_r\dot{x}(t) + x(t) = \delta(t-\tau_d)$$

with initial condition $x(0) = s(0) = 0$; x is an auxiliary variable. By extension, $I^{\text{syn}}(t)$ solves

$$\begin{aligned}\tau_s \dot{I}^{\text{syn}}(t) + I^{\text{syn}}(t) &= x(t) \\ \tau_r \dot{x}(t) + x(t) &= -\frac{J}{N}\sum_{k=1}^{N} y_k(t-\tau_d).\end{aligned}$$

In the limit $N \to \infty$, the inhibitory recurrent input $I^{\text{syn}}(t)$ does not fluctuate due to finite-size effects and is given by $I^{\text{syn}}(t) = -J(s * R)(t)$, with $R(t)$ the population activity when $N \to \infty$ (section 1.1.3, third definition). Then the system of equations above becomes

$$\begin{aligned}\tau_s \dot{I}^{\text{syn}}(t) + I^{\text{syn}}(t) &= x(t) \\ \tau_r \dot{x}(t) + x(t) &= -JR(t-\tau_d).\end{aligned} \tag{2.65}$$

Therefore, we have an equation (rather, a system) linking I^{syn} to the past activity. In the asynchronous state, $I_0^{\text{syn}} = -J(s * R_0)(t) = -JR_0 \int s(t)dt = -JR_0$, with $R_0 = \nu(\mu + I_0^{\text{syn}}, \sigma)$, using the LIF firing rate function in Eq. 2.21. What is needed now is a functional equation giving $R(t)$ as a function of $I^{\text{syn}}(t)$. If we assume that the time-dependent population activity does not depart much from the asynchronous state $[R(t) - R_0] \ll R_0$ and that at its minimum value $R(t)$ does not become negative (i.e., no rectification is needed), we can write

$$\Delta R(t) = (A * \Delta I^{\text{syn}})(t) = \int_{-\infty}^{\infty} A(\tau)\Delta I^{\text{syn}}(t-\tau)d\tau, \tag{2.66}$$

where A is the transfer function defined above; it must be evaluated at $\mu + I_0^{\text{syn}}$. Hence, A is a kernel linking the deviation $\Delta I^{\text{syn}}(t)$ of the total synaptic input from its value in the asynchronous state to the deviation of the population activity from its constant mean in the asynchronous state, $\Delta R(t) = R(t) - R_0$.

To find the eigenfrequencies of the network we follow the same process as in section 1.3.2, i.e., we write $\Delta R(t) = R_1 e^{\lambda t}$ and $\Delta I^{\text{syn}}(t) = I_1 e^{\lambda t}$, with $\lambda = \mu + i\omega$ (here μ is different from the bias μ in the LIF neuron model). Replacing this in Eq. 2.66 gives an equation for the eigenvalues λ. If for all λ, $\mu < 0$, then the asynchronous state is stable. If an eigenvalue pair crosses the imaginary axis ($\mu = 0$), the asynchronous state loses stability in favor of the oscillatory state. Setting $\mu = 0$, Eq. 2.66 becomes after simplifications:

$$R_1 = I_1 A^*(\omega) = I_1 |A(\omega)| e^{-i\Phi_A(\omega)}, \tag{2.67}$$

where $A(\omega) = |A(\omega)|e^{i\Phi_A(\omega)}$ is the susceptibility and we used the simple fact that $A(t) \in \mathbb{R}$

implies $A(-\omega) = A^*(\omega)$. From Eqs. 2.65, we extract after a few steps a second equation

$$I_1 = \frac{-JR_1 e^{-i\omega\tau_d}}{(1+i\omega\tau_r)(1+i\omega\tau_s)} \equiv JR_1|B(\omega)|\exp\left(i\pi - i\Phi_B\right), \tag{2.68}$$

with

$$|B(\omega)| = \frac{1}{\sqrt{(1+\omega^2\tau_r^2)(1+\omega^2\tau_s^2)}}$$

and

$$\Phi_B(\omega) = \omega\tau_d + \arctan(\omega\tau_r) + \arctan(\omega\tau_s), \tag{2.69}$$

where we used the arctangent addition formula $\arctan\left(\frac{u+v}{1-uv}\right) = \arctan(u) + \arctan(v)$. Putting Eq. 2.67 and Eq. 2.68 together yields the equation

$$1 = J|A(\omega)||B(\omega)|\exp\left[i\pi - i\Phi_A(\omega) - i\Phi_B(\omega)\right], \tag{2.70}$$

or, equivalently,

$$\begin{aligned} J|A(\omega)||B(\omega)| &= 1 \\ \Phi_A(\omega) + \Phi_B(\omega) &= (2k+1)\pi, \qquad k = 0, 1, \ldots. \end{aligned} \tag{2.71}$$

Brunel and Hansel have called Φ_A the neuronal phase lag because it only depends on properties of the neuron, whereas Φ_B is the synaptic phase lag because it involves synaptic parameters only. Examples for $\Phi_A = \arg\{A(\omega)\}$ are provided in Fig. 2.7; negative neuronal phase lags are possible when noise is low. Note that $\Phi_A \to \pi/4$ when $\omega \to \infty$ and $|A(\omega)|$ behaves like $1/\sqrt{\omega}$ for large frequencies [41].

From here, the procedure is as described in section 1.3.2. The phase equation gives the available eigenfrequencies. For $J \approx 0$ the asynchronous state is stable, hence increasing J above some critical value—or decreasing the negative coupling $-J/N$ below a critical value—will destabilize the asynchronous state, leading to oscillations; the amplitude equation yields this critical value once eigenfrequencies are replaced in it. When noise is very low—lower than the noise intensities depicted in Fig. 2.7—the neuronal phase lag has a sawtooth profile as a function of ω [43]. Combined with Φ_B, this profile leads to cluster states as described in section 1.3.2; families of cluster states exist for each $k \geqslant 0$. For higher noise, a single instability occurs for each $k \geqslant 0$. When the delay τ_d is large compared to the other synaptic time constants, the linear function $\omega\tau_d$ dominates the phase equation and the lowest eigenfrequency (i.e., when $k = 0$) is $\omega = \pi/\tau_d$. For an excitatory network, with coupling J/N and $J > 0$, the process is the same, except that the right-hand side of the phase equation in Eq. 2.71 is $2\pi k$. Thus, at large delay, an oscillatory instability occurs

with eigenfrequency $2\pi/\tau_d$, twice as large as for the inhibitory network. This is readily understandable. For the inhibitory network, if the population activity has a momentarily positive deviation from its mean, then after a time τ_d the whole network is depressed. This depressed state decreases inhibition, which will yield an increased activity after another time delay. Hence, the oscillation period is $2\tau_d$, or $\omega = \pi/\tau_d$. For the excitatory network, increased activity excites the network above its average level after a time τ_d, which is the period of the oscillation. The synaptic kinetics must be fast enough for the network to go back to its average activity before the next excitation kicks in, i.e., large delays are required [141].

2.8 Linear response theory

The formalism exploited in the previous section—and, in a slightly different way, that of section 1.3.2—can be considered a mean-field approach to neural dynamics [176]. Mean-field theories have a long and glorious history in the physics of interacting particles—see any book on condensed matter theory. In a nutshell, mean-field theory circumvents a detailed description of particle-particle interactions by instead considering that every particle interacts with the "mean field" generated by surrounding particles. Short-scale fluctuations—either temporal or spatial—are then neglected. The dynamics of the particle embedded in its mean field are solvable from first principles. From there, macroscopic properties (e.g., magnetization) can be obtained by solving a self-consistent equation: the particle interacts with its mean field, but also participates in the mean field of other particles. This is what we did in the last section—hiding a difficult part of the process: finding $A(\omega)$. The role of the particles is played by the neurons, interactions are *via* synaptic processes, the mean field is the average synaptic input and the macroscopic variable of interest is the population rate $R(t)$. The neurons are effectively uncorrelated even though they are recurrently connected [141]. Once the mean-field approximation has been made—when $N \to \infty$ and $R_N \to R$, removing stochastic fluctuations—neurons interact with the mean field with total disregard for the specific dynamics of other neurons.

From this discussion, we can predict when this mean-field theory fails in the context of neural dynamics. When noise is low ("low temperature"), neurons can synchronize in a spike-to-spike fashion (*cf.* section 1.3.3), thus generating fluctuations and correlation structures beyond the mean-field approximation. For the same reason, this theory is not applicable when the external input is a spatially correlated stochastic signal or, possibly, a deterministic signal with fast temporal variations. By *spatially correlated* we mean that the external input impinging on any neuron pair has non-negligible pairwise correlations, i.e.

$\langle I_i^{\text{ext}}(t) I_j^{\text{ext}}(t)\rangle \neq \langle I_i^{\text{ext}}(t)\rangle\langle I_j^{\text{ext}}(t)\rangle$. Moreover, the asynchronous state must not lie near the f-I curve saturation or rectification when we study oscillatory instabilities. Yet another drawback is that single-neuron spiking properties—which are routinely monitored in neurophysiology experiments—are not accessible analytically.

In this section, we present the linear response theory [125], an analytic method that has been designed to deal with some of these issues. First, we describe the fundamental *ansatz* of the theory in the context of the single neuron, and then we apply it to neural networks. Interestingly, this theory has been first applied to the study of gamma rhythms in the ELL of weakly electric fish [77]. Therefore, it seems rather appropriate to treat this case in some detail. Along the way, the issue of gain control by inhibitory feedback will appear.

2.8.1 Single neuron

The spiking process of a LIF neuron with the dynamics of Eq. 2.17 has a PSD given exactly by Eq. 2.49. An external input $I^{\text{ext}}(t)$ will modify this PSD. For deterministic inputs, the rate is time-dependent and therefore the underlying spiking process is non-stationary (section 2.6). Hence, PSDs are not defined in the usual sense. It is possible to circumvent this problem when the input is periodic, e.g. $I^{\text{ext}}(t) = a\cos(2\pi f t + \phi)$. For it is then possible to average over the initial phase ϕ of the input and thus produce a stationary process. When the external input is *weak*, the PSD of the neuron perturbed by this input is given by [123]

$$S(\omega) \approx S_0(\omega) + |A(\omega)|^2 \left\{ \frac{\pi}{2} a^2 [\delta(\omega - 2\pi f) + \delta(\omega + 2\pi f)] \right\}. \tag{2.72}$$

The background or *unperturbed* spectrum, $S_0(\omega)$, is obtained with $I^{\text{ext}}(t) \equiv 0$. The expression in curly brackets is the external input's PSD. Thus, the *perturbed* spectrum is the sum of an unperturbed spectrum plus the product of the input power and $|A|^2$. Delta peaks appear at the driving frequency $2\pi f$. This is a linear approximation since we assume that the response only occurs at the input frequency.

A similar linear approximation can be given for zero-mean stationary stochastic inputs with finite variances:

$$S(\omega) \approx S_0(\omega) + |A(\omega)|^2 S_{\text{ext}}(\omega), \tag{2.73}$$

with $S_{\text{ext}}(\omega)$ the PSD of the external input. A heuristic way to understand this expression is

to write the *ansatz*[4] [125] (compare with Eq. 2.57):

$$y(t) \approx y_0(t) + (A * I^{\text{ext}})(t), \tag{2.76}$$

where y_0 is the spike train of the neuron when no external input is applied (the *unperturbed spike train*), and y is the spike train of the neuron when subjected to the input (the *perturbed spike train*). This expression is not correct mathematically since it involves a mixture of point processes (y and y_0) and a continuous process ($A * I^{\text{ext}}$). If, however, we compute the PSD of both sides of this expression, we get for the left-hand side:

$$\lim_{T\to\infty} \frac{\langle |\tilde{y}_T(\omega)|^2 \rangle}{T} = S(\omega). \tag{2.77}$$

For the right-hand side,

$$\lim_{T\to\infty} \frac{\langle |\tilde{y}_{0,T}(\omega) + A(\omega)\tilde{I}^{\text{ext}}_T(\omega)|^2 \rangle}{T} = S_0(\omega) + 2\,\text{Re}\,\{A(\omega)S_{y_0 I^{\text{ext}}}(\omega)\} + |A(\omega)|^2 S_{\text{ext}}(\omega),$$

where $S_{y_0 I^{\text{ext}}}(\omega)$ is the cross-spectrum between the unperturbed spike train and the external signal. This cross-spectrum must be zero since, logically, the external signal cannot modulate the unperturbed spike train. Therefore, the two processes are uncorrelated and the cross-spectrum is zero because the input is zero-mean. We thus recover Eq. 2.73. Although not rigorous from the mathematical standpoint, Eq. 2.76 is very useful for computing expressions for perturbed PSDs of neurons embedded in a network, as will be seen in the next section.

The approximation of Eq. 2.73 is good as long as the response of the neuron is sufficiently linear and the variance of the external input is finite. Below, we will often use Gaussian white noise (GWN) as an external input, whose variance is infinite, i.e. $\langle \xi^2(t) \rangle = \delta(0)$. If we write it as $I^{\text{ext}}(t) = \sigma_{\text{ext}}\eta(t)$, with $\langle \eta(t) \rangle = 0$ and $\langle \eta(t)\eta(t+\tau) \rangle = \delta(\tau)$, then the total

[4] Here is another way to see the approximation in Eq. 2.73. In first approximation, we can write for the ACF of the spike train $Q(\tau) \approx Q_0(\tau) + \gamma(\tau)$, where $Q(\tau)$ (resp. $Q_0(\tau)$) is the perturbed (resp. unperturbed) ACF. $\gamma(\tau)$ is a functional of $I^{\text{ext}}(t)$ and $A(t)$. Since I^{ext} is a stochastic function and $Q(\tau)$ is a deterministic function, $\gamma(\tau)$ can only depend on statistics of I^{ext}. Since the latter is WSS with zero mean, the remaining low-order statistics that can be used is its ACF, $R_{\text{ext}}(\tau)$. To linear order in $Q_{\text{ext}}(\tau)$, we have

$$Q(\tau) \approx Q_0(\tau) + \int d\tau_1\, \kappa(\tau_1) Q_{\text{ext}}(\tau - \tau_1) \tag{2.74}$$

where $\kappa(\cdot)$ is the first-order kernel of the expansion; it depends on $A(\tau)$. Taking the Fourier transform of the expression above gives

$$S(\omega) \approx S_0(\omega) + \tilde{\kappa}(\omega) S_{\text{ext}}(\omega). \tag{2.75}$$

Since $S(\omega)$ and $S_{\text{ext}}(\omega)$ are real functions of ω, $\tilde{\kappa}(\omega)$ must be a real function as well. In the theory of linear, translation invariant systems, the PSD of the output signal of a system is related to the PSD of its input signal by $S_{\text{out}}(\omega) = |H(\omega)|^2 S_{\text{in}}(\omega)$, with $H(\omega)$ the Fourier transform of the impulse response function. Therefore, we use $\tilde{\kappa}(\omega) = |A(\omega)|^2$ in the linear approximation (instead of, say, $\text{Re}\{A(\omega)\}$).

stochastic input impinging on the neuron is $\sigma\xi(t) + \sigma_{\text{ext}}\eta(t)$. This signal is equivalent to a zero-mean GWN with intensity $\sigma^2_{\text{eff}} = \sigma^2 + \sigma^2_{\text{ext}}$, because it is the sum of two uncorrelated zero-mean GWN signals. Then, for a GWN external input, the perturbed spectrum is exactly

$$S(\omega) = S_0(\omega)|_{\sigma_{\text{eff}}}, \tag{2.78}$$

where $S_0(\omega)|_{\sigma_{\text{eff}}}$ is the unperturbed PSD of a neuron with effective noise intensity σ_{eff}. Moreover, if the average of I^{ext} is nonzero then the LRT *ansatz* is rather

$$y(t) \approx y_0(t) + (A * [I^{\text{ext}} - \langle I^{\text{ext}} \rangle])(t), \tag{2.79}$$

and the unperturbed PSD must be evaluated with bias $\mu + \langle I^{\text{ext}} \rangle$ in Eq. 2.49 [125].

2.8.2 Network of spiking neurons

We now connect noisy LIF neurons to form a network of N neurons. The connectivity matrix is $\mathbf{W}$ and we denote the spike train of the nth neuron by y_n. Its membrane potential is denoted by v_n and the subthreshold dynamic is given by (*cf.* Eq. 2.15)

$$\frac{dv_n}{dt} = \mu_n - v_n(t) + \sigma_n\xi_n(t) + I^{\text{syn}}_n(t) + I^{\text{ext}}_n(t).$$

With this formulation, all neurons may receive different synaptic and external inputs at any time; for instance, $I^{\text{syn}}_n(t)$ is the total synaptic input received by neuron n at time t, and so on. The neurons may have different biases and internal noise intensities. Also, we stress again that "external input" stands for any inputs which do not come from within the network. Synaptic currents are given by Eq. 2.7:

$$I^{\text{syn}}_n(t) = \sum_{k=1}^{N} W_{nk}(\alpha_{nk} * y_k)(t),$$

where α_{nk} is the synaptic function linking presynaptic neuron k to postsynaptic neuron n. Having these indices allows the use of different synaptic functions and time constants for every synaptic interaction. Again, we assume that the applied currents I^{ext}_n are WSS stochastic processes.

In the framework of the LRT presented above, fluctuations of the external input and of the synaptic bombardment will be taken as perturbations of an underlying background

activity. The latter follows the subthreshold dynamics:

$$\dot{v}_n = \mu_n + \langle I_n^{\text{syn}} \rangle - v_n(t) + \sigma_n \xi_n(t)$$

with $\langle I_n^{\text{syn}} \rangle$ the ensemble average of $I_n^{\text{syn}}(t)$ in the asynchronous state. These dynamics generate *unperturbed* spike trains, denoted by $y_n^{(0)}$. As stated above, the LRT assumes that, in the Fourier representation and for nonzero frequencies ($\omega \neq 0$)[5],

$$\tilde{y}_n(\omega) \approx \tilde{y}_n^{(0)}(\omega) + A_n(\omega)[\tilde{I}_n^{\text{syn}}(\omega) + \tilde{I}_n^{\text{ext}}(\omega)], \tag{2.80}$$

where

$$\tilde{I}_n^{\text{syn}}(\omega) = \sum_{k=1}^{N} W_{nk} \tilde{\alpha}_{nk}(\omega) \tilde{y}_k(\omega),$$

and $A_n(\omega)$ is the susceptibility of the nth neuron. These susceptibilities depend on the *effective biases*,

$$\mu_{\text{eff},n} = \mu_n + \langle I_n^{\text{syn}} \rangle, \tag{2.81}$$

and the noise intensities σ_n [125]. These effective biases take into account the average synaptic input coming from other neurons in the network. More explicitly, in the susceptibility of Eq. 2.62—and in the expression for the rate needed for the calculation of the susceptibility—, μ must be replaced by the effective biases above. If it happens that the average of I_n^{ext} is nonzero, it can always be included in the effective bias.

We define the column vector $\tilde{\mathbf{y}}(\omega) \equiv (\tilde{y}_1(\omega) \dots \tilde{y}_N(\omega))^T$, such that

$$\tilde{\mathbf{I}}_{\text{syn}} \equiv (\tilde{I}_1^{\text{syn}} \dots \tilde{I}_N^{\text{syn}})^T = \mathbf{W}_\alpha \tilde{\mathbf{y}}(\omega),$$

where matrix $\mathbf{W}_\alpha$ has elements $(\mathbf{W}_\alpha)_{ij} = W_{ij}\tilde{\alpha}_{ij}(\omega)$, and ^{T} denotes the transpose operation. The argument ω will be implicit from now on. With these definitions, we can write Eq. 2.80 as

$$\tilde{\mathbf{y}} \approx \tilde{\mathbf{y}}_0 + \mathbf{A}[\tilde{\mathbf{I}}_{\text{syn}} + \tilde{\mathbf{I}}_{\text{ext}}] = \tilde{\mathbf{y}}_0 + \mathbf{A}[\mathbf{W}_\alpha \tilde{\mathbf{y}} + \tilde{\mathbf{I}}_{\text{ext}}],$$

[5] For a LIF neuron with dynamics $dv/dt = \mu - v + \sigma\xi + I_{\text{tot}}$, with $I_{\text{tot}}(t) = I_{\text{syn}}(t) + I_{\text{ext}}(t)$, the LRT *ansatz* asserts that (here A = susceptibility versus time and $\tilde{A}$ = susceptibility versus frequency)

$$y(t) \approx y_0(t) + (A * [I_{\text{tot}} - \langle I_{\text{tot}} \rangle])(t) = y_0(t) + (A * I_{\text{tot}})(t) - \tilde{A}(0)\langle I_{\text{tot}} \rangle.$$

Upon computing the PSD, the constant term $\tilde{A}(0)\langle I_{\text{tot}} \rangle$ produces a zero-frequency component that can be avoided by excluding this constant term.

where $\mathbf{A} \equiv \text{diag}(A_1, \ldots, A_N)$, $\tilde{\mathbf{I}}_{\text{ext}} \equiv (\tilde{I}_1^{\text{ext}} \ldots \tilde{I}_N^{\text{ext}})^T$ and $\tilde{\mathbf{y}}_0 \equiv (\tilde{y}_1^{(0)} \ldots \tilde{y}_N^{(0)})^T$. If the matrix $[\mathbf{1} - \mathbf{A}\mathbf{W}_\alpha]$, with $\mathbf{1}$ the identity matrix, is invertible, then

$$\tilde{\mathbf{y}} \approx [\mathbf{1} - \mathbf{A}\mathbf{W}_\alpha]^{-1}(\tilde{\mathbf{y}}_0 + \mathbf{A}\tilde{\mathbf{I}}_{\text{ext}}). \tag{2.82}$$

To get the power and cross-spectra of all neurons, we multiply both sides of the equation by $\tilde{\mathbf{y}}^\dagger$ to the right ($\dagger$ represents conjugate transpose), take the ensemble average, divide by T and let $T \to \infty$. To simplify notation, we will represent these operations by brackets: $\langle \cdot \rangle \equiv \lim_{T\to\infty} \frac{\langle \cdot \rangle}{T}$. The left-hand side is

$$\langle \tilde{\mathbf{y}}\tilde{\mathbf{y}}^\dagger \rangle \equiv \mathbf{S} = \begin{pmatrix} S_1 & S_{12} & \ldots & \ldots & S_{1N} \\ S_{12}^* & & & & \\ \vdots & & & & \\ & & \ddots & & \\ \vdots & & & & \\ S_{1N}^* & & \ldots & & S_N \end{pmatrix}, \tag{2.83}$$

where S_i is the perturbed PSD of neuron i and S_{ij} with $i \neq j$ is the cross spectrum of neurons i and j; $\mathbf{S}$ is called the spectral matrix. The right-hand side is

$$[\mathbf{1} - \mathbf{A}\mathbf{W}_\alpha]^{-1} \left\langle (\tilde{\mathbf{y}}_0 + \mathbf{A}\tilde{\mathbf{I}}_{\text{ext}})(\tilde{\mathbf{y}}_0 + \mathbf{A}\tilde{\mathbf{I}}_{\text{ext}})^\dagger \right\rangle \left\{ [\mathbf{1} - \mathbf{A}\mathbf{W}_\alpha]^{-1} \right\}^\dagger,$$

with

$$\left\langle (\tilde{\mathbf{y}}_0 + \mathbf{A}\tilde{\mathbf{I}}_{\text{ext}})(\tilde{\mathbf{y}}_0 + \mathbf{A}\tilde{\mathbf{I}}_{\text{ext}})^\dagger \right\rangle = \left\langle \tilde{\mathbf{y}}_0\tilde{\mathbf{y}}_0^\dagger \right\rangle + \mathbf{A} \left\langle \tilde{\mathbf{I}}_{\text{ext}}\tilde{\mathbf{I}}_{\text{ext}}^\dagger \right\rangle \mathbf{A}^\dagger.$$

The last equality comes from the fact that the external input is uncorrelated with the unperturbed activity, i.e. $\langle \tilde{\mathbf{y}}_0\tilde{\mathbf{I}}_{\text{ext}}^\dagger \rangle = \mathbf{0}$. Then, we define $\mathbf{S}_0 = \text{diag}(S_1^{(0)}, \ldots, S_N^{(0)}) = \langle \tilde{\mathbf{y}}_0\tilde{\mathbf{y}}_0^\dagger \rangle$, the diagonal matrix containing the unperturbed PSDs, whose elements are to be computed from Eq. 2.49 with the effective biases of Eq. 2.81. Finally, with $\mathbf{S}_{\text{ext}} = \langle \tilde{\mathbf{I}}_{\text{ext}}\tilde{\mathbf{I}}_{\text{ext}}^\dagger \rangle$ the spectral matrix of the external input, we get

$$\mathbf{S} \approx [\mathbf{1} - \mathbf{A}\mathbf{W}_\alpha]^{-1} \left(\mathbf{S}_0 + \mathbf{A}\mathbf{S}_{\text{ext}}\mathbf{A}^\dagger \right) \{ [\mathbf{1} - \mathbf{A}\mathbf{W}_\alpha]^{-1} \}^\dagger. \tag{2.84}$$

This expression—which has been obtained by Trousdale *et al.* [209]—states that the perturbed spectral matrix is obtained from a matrix transformation of the unperturbed spectral matrix

and of the spectral matrix of the external input, after the latter has been processed by the susceptibility matrix.

For LIF neurons, analytic expressions for $\mathbf{A}$ and $\mathbf{S}_0$ are readily available [125], and this is what motivated the use of this particular model in this thesis. However, Eq. 2.84 is quite general—as long as the linear approximation remains valid—so other unperturbed spectra and susceptibilities could be considered as well, such as ones obtained numerically from simulations of more complex models or from experimental estimates, with the help of the method explained in section 2.6.

2.8.3 Oscillations and gain control through delayed inhibitory feedback

We now apply the LRT formalism to a concrete case, the occurrence of gamma rhythms in the ELL of weakly electric fish under spatially correlated stimuli (Fig. 2.8A). This phenomenon was studied in [75, 77, 125] and provides the motivation for chapter 3. Experiments have been carried out using the setup depicted in the top panel of Fig. 1.9B. That is, four electrode dipoles driven by random signals produced random amplitude modulations (AMs) of the EOD. It was possible to tune the degree of spatial correlation between the electrodes: in-phase signals produced global AMs which concurrently activated a large proportion of electroreceptors, whereas out-of-phase signals yielded a global signal that is spatially uncorrelated (or, at least, less correlated than the in-phase signal).

The experimental measurements were performed in the centrolateral and lateral segments and involved SP and/or IP cells [75] since they are more dorsally localized in the ELL and easier to target with electrodes than DP cells. However, the modeling work featured in these papers did not make any distinction between DP, IP or SP cells, and between ON or OFF cells. The main assumptions were that 1) ELL pyramidal cells constitute a fully recurrent/feedback network whose neurons are connected *via* delayed inhibitory synapses and 2) the electroreceptor afferents linearly code for the AMs of the EOD and project to ELL neurons. Based on section 1.4, the first assumption is obviously false. But, as noted there, the direct feedback ELL→Pd→ELL can be reduced to a feedback loop in a first approximation, if we disregard altogether the dorsoventral position of ELL pyramidal cells. To simplify, it is also assumed that the Pd nucleus merely plays the role of a relay station: a spike emitted by an ELL neuron affects all ELL neurons after a propagation delay corresponding to the time required to travel the distance between the ELL and the Pd nucleus (see Fig. 2.8B). From [75], it is the diffuse inhibition coming from the BP cells that is responsible for the oscillatory response to global stimuli. Therefore, all ELL pyramidal cells are effectively interconnected, forming a fully recurrent network, with inhibitory synapses. The second assumption is

known to be approximately valid for moderate amplitude modulations [98]. Hence, the model featured in [75, 77] and thoroughly analyzed in [125], must not be interpreted as a direct attempt at modeling the ELL in all its intricacies. Rather, it should be viewed as a reductionist attempt at understanding gamma rhythms produced by spatially correlated stimuli.

For the external input, Doiron, Lindner and coworkers used a mixture of a spatially correlated random signal, $\eta_c(t)$, and another random signal uncorrelated across ELL neurons, $\eta_i(t)$ [125]. These two inputs represent the extreme cases of totally correlated and totally uncorrelated dipole stimuli, respectively. They are both zero-mean and their correlation functions are given by $\langle \eta_i(t)\eta_j(t')\rangle = \delta_{ij} Q_{\mathrm{st}}(t-t')$, $\langle \eta_i(t)\eta_c(t')\rangle \equiv 0$ and $\langle \eta_c(t)\eta_c(t')\rangle = Q_{\mathrm{st}}(t-t')$; η_c and η_i thus have the same PSD, $S_{\mathrm{st}}(\omega) = \mathscr{F}\{Q_{\mathrm{st}}\}$. Doiron *et al.* [77] constructed a compound signal using these two signals that allows various degrees of correlation:

$$I_i^{\mathrm{ext}}(t) \equiv \sqrt{c}\eta_c(t) + \sqrt{(1-c)}\eta_i(t), \tag{2.85}$$

where c is a "correlation parameter". When the dipole input has a perfect spatial correlation, $c = 1$ and all neurons receive the same input; when it is totally uncorrelated, $c = 0$ and on average any two neurons do not receive the same input. The square-root prefactors ensure that the PSD is $S_{\mathrm{st}}(\omega)$ for all c. The spectral matrix of I_i^{ext} has elements $(\mathbf{S}_{\mathrm{ext}})_{ij}(\omega) = cS_{\mathrm{st}}(\omega) + (1-c)S_{\mathrm{st}}(\omega)\delta_{ij}$.

The network is fully recurrent (all-to-all connectivity) with uniform synaptic weights, $\mathbf{W}_{ij} = G/N$, with $G < 0$, and α synaptic functions. Therefore, the matrix $\mathbf{W}_\alpha = (G/N)\tilde{\alpha}\mathbf{U}$ in Eq. 2.84, where $\mathbf{U}$ is a matrix whose elements are all 1. Also, assuming that the network is homogeneous, $\mathbf{A} = A\mathbf{1}$ and $\mathbf{S}_0 = S_0\mathbf{1}$. Defining $b(\omega) \equiv (G/N)A(\omega)\tilde{\alpha}(\omega)$, we have for $[\mathbf{1} - \mathbf{A}\mathbf{W}_\alpha]^{-1}$:

$$\begin{aligned}
[\mathbf{1} - b\mathbf{U}]^{-1} &= \mathbf{1} + b\mathbf{U} + b^2\mathbf{U}^2 + b^3\mathbf{U}^3 + \dots \\
&= \mathbf{1} + b\mathbf{U} + b^2N\mathbf{U} + b^3N^2\mathbf{U} + \dots \\
&= \mathbf{1} + \frac{b}{1-bN}\mathbf{U}.
\end{aligned}$$

The largest eigenvalue of $b\mathbf{U}$ in absolute value is $|bN|$. Since $f(x) = [1-x]^{-1} = \sum_{n=0}^{\infty} x^n$ converges for $|x| < 1$, the expansion of $[\mathbf{1} - b\mathbf{U}]^{-1}$ will converge if $|bN| < 1$, i.e. if $|GA(\omega)\tilde{\alpha}(\omega)| < 1$. Using Eq. 2.84, we get the perturbed spectral matrix

$$\mathbf{S} \approx \left(\mathbf{1} + \frac{b}{1-bN}\mathbf{U}\right)\left(S_0\mathbf{1} + c|A|^2S_{\mathrm{st}}\mathbf{U} + (1-c)|A|^2S_{\mathrm{st}}\mathbf{1}\right)\left(\mathbf{1} + \frac{b^*}{1-b^*N}\mathbf{U}\right). \tag{2.86}$$

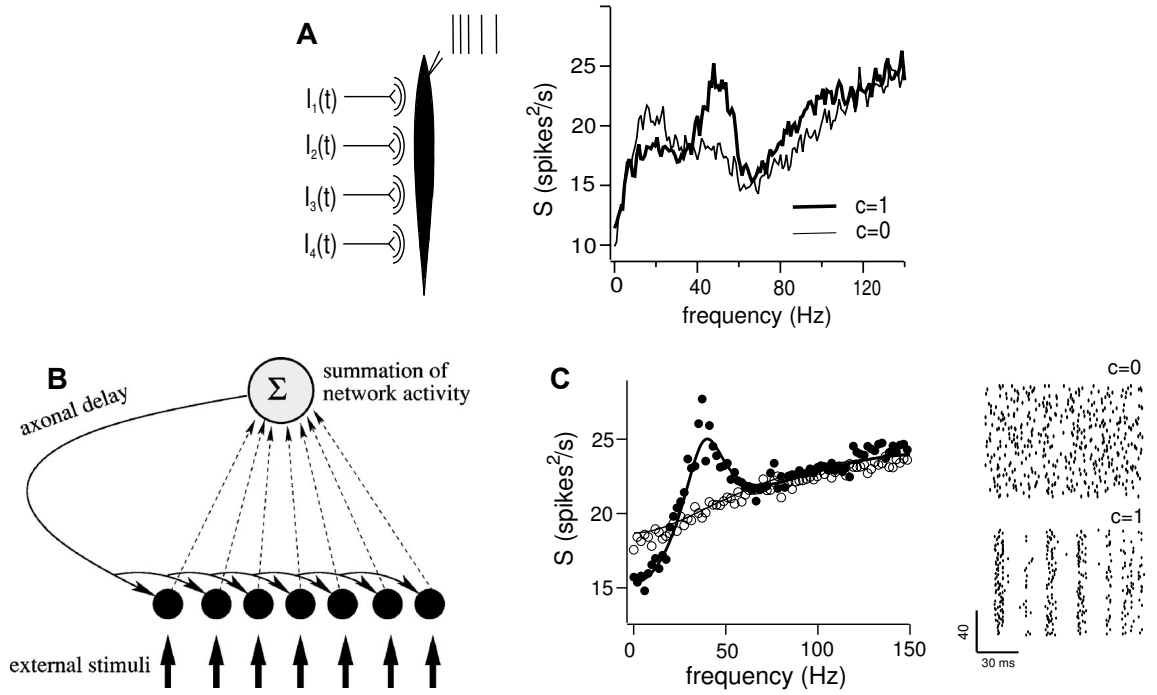

Figure 2.8: Gamma rhythms in the ELL under correlated spatial noise.
A) Experimental setup (left) as described in Fig. 1.9B, and experimental results for perfectly correlated ($c = 1$) and totally uncorrelated ($c = 0$) random AMs. Note the presence of a salient peak around 50 Hz in the PSD when $c = 1$.
B) Schematic of the feedback loop mimicking the loop ELL→Pd→ELL. The processing performed by the Pd nucleus is replaced by a simple summation of incoming spikes from the ELL. Because the diffuse inhibition coming from the BP cells is thought to play a prominent role in the generation of rhythms in the ELL, the spikes arriving at synapses onto ELL neurons after the summation produce inhibitory postsynaptic currents. Effectively, then, the ELL neurons are recurrently connected in an all-to-all fashion with a delay coming from the propagation of spikes in the feedback loop.
C) Numerical results (dots) and analytical results from the LRT (lines) for PSDs in a delayed inhibitory feedback network. On the right-hand side, raster plots are displayed for the two limit cases. The y axis is the neuron number and the x axis is time. Each dot corresponds to a spike. Panels A and C reproduce figures 1 and 2 in [77] and panel B is figure 1 in [125], with permissions.

After performing the matrix multiplication and using the definition of b, we obtain

$$\begin{aligned}\mathbf{S} \approx\ & S_0\mathbf{1} + |A|^2 S_{\text{st}}[c\mathbf{U} + (1-c)\mathbf{1}] \\ & + \left[\frac{S_0}{N} + |A|^2 S_{\text{st}}\left(c + \frac{1-c}{N}\right)\right]\left[\frac{2\text{Re}\{G\tilde{\alpha}A\} - |G\tilde{\alpha}A|^2}{|1-G\tilde{\alpha}A|^2}\right]\mathbf{U},\end{aligned}$$

which represents, in a condensed fashion, the results appearing in [125]. The single-neuron PSDs constitute the diagonal of this matrix, whereas the off-diagonal terms are the cross-spectral densities, which must be all equal to one another given the homogeneity of the network. In the limit $N \to \infty$, finite-size effects on the spectral quantities disappear and the spectral matrix becomes

$$\mathbf{S} \approx S_0\mathbf{1} + |A|^2 S_{\text{st}}[c\mathbf{U} + (1-c)\mathbf{1}] + c|A|^2 S_{\text{st}}\left[\frac{2\text{Re}\{G\tilde{\alpha}A\} - |G\tilde{\alpha}A|^2}{|1-G\tilde{\alpha}A|^2}\right]\mathbf{U}. \tag{2.87}$$

When the external input is totally uncorrelated, $c = 0$ and the spectral matrix becomes (still with $N \to \infty$)

$$\mathbf{S} \approx (S_0 + |A|^2 S_{\text{st}})\mathbf{1}. \tag{2.88}$$

For $c = 1$, we get after simplifications

$$\mathbf{S} \approx S_0\mathbf{1} + \left|\frac{A}{1-G\tilde{\alpha}A}\right|^2 S_{\text{st}}\mathbf{U}. \tag{2.89}$$

To understand the meaning of these expressions, we consider the population activity of the network, $R_N(t)$. In section 2.7, the limit $N \to \infty$ of this quantity was taken early in the developments, thus removing finite-size fluctuations. For a homogenous neural network, the PSD of R_N, called the *population spectrum*, can be expressed as a function of the single-neuron PSD and of the cross-spectrum [125]:

$$S_{\text{pop}}(\omega) \equiv \lim_{T\to\infty} \frac{\langle|\tilde{R}_N(\omega)|^2\rangle}{T} = S_{\text{cross}}(\omega) + \frac{1}{N}[S(\omega) - S_{\text{cross}}(\omega)], \tag{2.90}$$

where $S_{\text{cross}} = \mathbf{S}_{ij}$ for any $i \neq j$ and $S = \mathbf{S}_{ii}$ for any i. Hence, from the equations above we have

$$S_{\text{pop}} = \left[\frac{S_0}{N} + |A|^2 S_{\text{st}}\left(c + \frac{1-c}{N}\right)\right]\left[\frac{1}{|1-G\tilde{\alpha}A|^2}\right].$$

In the limit $N \to \infty$, we get

$$S_{\text{pop}} = c\left|\frac{A}{1-G\tilde{\alpha}A}\right|^2 S_{\text{st}}. \tag{2.91}$$

When $c = 0$, the large N limit of S_{pop} is zero, unless the denominator $1 - G\tilde{\alpha}A$ itself becomes zero. But this is exactly the condition for the asynchronous state to lose stability in section 2.7. The only difference here is that, with the notation of that section, $\tau_r = \tau_s$ for α synapses. So, when the external input is uncorrelated, there is no spike-to-spike synchrony (i.e., phase-locking between spikes of different neurons) and an oscillatory state appears only when the aforementioned condition is fulfilled. When c is greater than zero, S_{pop} can be nonzero even when the denominator is not equal to zero. This means that the population activity displays coherent fluctuations mediated by the correlated part of the input. Synchronization beyond the firing-rate synchrony of section 1.3.3—and beyond a mean-field treatment—is then possible. Moreover, from Eqs. 2.89 and 2.91 we conclude that the single-neuron PSD is a good indicator of the presence of oscillations at the population level, because a salient peak should protrude above the background whenever the population oscillates. This is illustrated in Fig 2.8C, where it is shown that the model network oscillates when $c = 1$.

Gain control via inhibitory feedback and outlook

To compute the perturbed spectrum we have to evaluate the unperturbed PSD. This, in turn, requires to find the firing rate in the asynchronous state (*cf.* Eq. 2.49). It is obtained just as in section 2.7. We have to use Eq. 2.21, with μ replaced by the effective bias $\mu_{\text{eff}} = \mu + \langle I^{\text{syn}} \rangle$, which must be determined self-consistently like in mean-field theory. The average synaptic current is ($G < 0$)

$$\langle I^{\text{syn}} \rangle = \langle G(\alpha * R_N) \rangle = Gr, \tag{2.92}$$

because the integral of α is normalized to 1 and $\langle R_N \rangle = r$ in the stationary asynchronous state, with r the rate of every neuron. The firing rate is obtained by solving the self-consistent equation [125]

$$r = \nu(\mu + Gr, \sigma, \tau_r). \tag{2.93}$$

Sutherland and coworkers [203] have shown that inhibitory feedback produces divisive gain control by comparing the firing rates extracted from this equation in the cases $G = 0$ and $G < 0$. They used their results to explain the gain control occurring in the ELL as mentioned in section 1.4.4. The experiments that they were referring to actually involved surgical and pharmacological lesioning of the parallel fibers, part of the indirect "feedback" [16, 15]. Calling this pathway a feedback is both true and false. True, because the ELL projects to the Pd, which then projects to the EGp and back to the ELL, constituting a coarse-grained feedback from ELL to ELL. False, because it is the deep cells that project to Pd, and it is the superficial cells that receive the "feedback". Taking this into account, this pathway should be called "feedforward". But DP and SP cells both receive inputs from the P-units, so they are at the same information-processing stage.

In chapter 4, we study gain control in this feedforward setting and find that it supports divisive gain control as well as subtractive and non-monotonic gain control. The direct feedback is more treacherous because it is unclear how feedback and how feedforward it is. In chapter 5, we clarify these fuzzy points. In the next chapter, we will study oscillations in the ELL using the formalism just presented, but considering both ON and OFF cells connected recurrently.

Chapter 3

Linear response theory for two neural populations applied to gamma oscillation generation

Linear response theory (LRT) can be used to compute spectral properties of single and populations of stochastic leaky integrate-and-fire neurons. The effects of inputs, both external and from delayed feedback, can be modeled within that theory when the neural function is sufficiently linearized by noise. It has been used to explain experiments where gamma oscillations are induced by spatially correlated stochastic inputs to a network with delayed inhibitory feedback. Here we expand this theory to include two distinct population types. We first show how to deal with homogeneous networks where both types of neurons are identical. We further tackle the asymmetric case, where noise or bias differ. We also analyze the case where the membrane time constants differ, based on experimental evidence, which requires delicate alterations of the theory. We directly apply the theory to networks of ON and OFF cells in the electrosensory system, which together provide global delayed negative feedback to all cells; however, ON and OFF cells receive external inputs of opposite polarities. Theoretical results are in excellent agreement with numerical simulations of the two population network. In contrast to the case of a single ON cell population with feedback, the more realistic presence of both cell types can significantly reduce the propensity of the delayed feedback network to oscillate for spatially correlated inputs. Our results are further linked to recent predictions from deterministic neural field theory. Among other findings, our work suggests that the observed gamma oscillations could be explained only if the ON and OFF cell feedback pathways are anatomically segregated. Thus our two population LRT can make specific predictions about network topography in specific systems.

3.1 Introduction

Collective oscillations are common in neural assemblies [47, 45]. Mechanisms by which neurons are able to fire rhythmically in more or less synchronized patterns have been explored at the cellular and network levels [219]. A general question is how interactions, both recurrent and with external inputs, determine the properties of coherent oscillations. A well-known type of neural rhythm is gamma oscillations, which occur in the frequency band $30 - 90$ Hz; their exact functions remain elusive [49]. The nonlinear character of nerve cells and the noisy environments in which they evolve often complicate the analysis of population dynamics. However, intrinsic noise may linearize the single-neuron behavior with respect to external inputs and allow useful approximations [199]. Here, such an approach — backed by numerical simulations — will be taken to study gamma oscillations in neural nets involved in the electrosense of weakly electric fish.

Weakly electric fish — and other creatures as well — must be able to ascertain the nature of a given stimulus, be it a prey, a predator, a conspecific or a pattern in their habitat, so that their behavioral response is appropriate. Stimuli differ by their physical properties and dimensions and, as a corollary, by the topology and timing of the signals they send to the electrosensory lateral lobe line (ELL), the first-order processing center of electrosensory stimuli (see [133] for a review). The pyramidal cells of the ELL receive inputs from cutaneous electroreceptors, which sense the amplitude of the electric field surrounding their body. Thus, a prey will excite a small area of their body (local stimulus), whereas a signal from a conspecific will stimulate a large proportion of the electroreceptors (global signal). A global signal may be either spatially correlated (e.g. communication signals from conspecifics), uncorrelated (e.g. large irregular rocks) or anything in between. We assume that a spatially correlated global signal stimulates concurrently all the electroreceptors.

A first step in the sensory discrimination process may be achieved through shifts of the dominant oscillation frequency of the sensory network. By dominant oscillation frequency, we have in mind the frequency at which the power spectrum of the cells' spike train reaches a local maximum. It is established that the ELL pyramidal cells will enter an oscillatory state when submitted to globally extended random signals [75, 54, 77]. These oscillations are in the gamma range with a dominant frequency ≈ 50 Hz and seem to involve a synchronous activity of the cells [75]. Less clear is the effect of a local or a spatially uncorrelated global stimulus, both seemingly eliciting weak low-frequency oscillations ($\approx 15 - 20$ Hz), with the spikes being largely uncorrelated [75, 54, 77].

The ELL cells are subdivided into two classes depending on their response to excitatory

inputs: ON cells (or E cells) increase their firing rate whereas OFF cells (I cells) decrease their firing rate with respect to their baseline activities. Almost every OFF cell is adjacent to an ON cell [132] and both share a common receptive field center. The network studied in this paper may be viewed as a one-dimensional layer of ON and OFF cell pairs receiving inputs from the skin. The ELL's pyramidal cells are mainly connected through a feedback loop, not by direct synaptic connections [29]. Delayed feedback is a potent mechanism by which network activity may oscillate. In the case of ELL pyramidal cells, a delayed inhibitory feedback is responsible for the onset of oscillations [75].

Networks of leaky integrate-and-fire (LIF) neurons of the ON type — connected via delayed inhibitory feedback and receiving external Gaussian noise — were used to model the ELL circuitry [55, 75, 77, 125]. The neural population responsible for the feedback signal together with the excitatory ON cells, form an excitatory-inhibitory network. These networks are fertile breeding grounds for gamma oscillations [49]. Using linear response theory (LRT), Refs. [77], [55] and particularly [125], have explained that gamma oscillations arise in the ELL due to the interplay between a correlated input and the delayed feedback. However, OFF cells were completely ignored in these works. Henceforth, this type of network, where OFF cells are absent, will be called ON-only networks. On the other hand, neural field studies of ON and OFF cell (ON/OFF) networks [120, 121] showed that the latter are less likely to exhibit rhythmic activity compared to ON-only networks under constant, spatially localized inputs.

In this chapter, we derive a LRT for the ON/OFF network and use it to analyze the impact of the OFF cells on the network oscillations. In the symmetric case, for which the properties of the ON and OFF cells are identical except for the way they receive inputs, we find that the OFF cells impede the oscillations found in the ON-only network. This extends some conclusions of [120, 121] to stochastic stimuli. We then try to break that symmetry based on physiological knowledge, to see if we can recover the gamma oscillations. It appears that oscillations can be brought back, but not at an intensity comparable to that of either the ON-only network [125] or the experimental results of [77]. This suggests that the ON and OFF cells belong to segregated pathways, i.e. ON (resp. OFF) cells are connected dominantly through feedback to ON (resp. OFF) cells.

The paper is structured as follows: First, we present the model in detail. Then, we use the LRT to write the perturbed spike trains in terms of the unperturbed activity plus a small quantity. This allows us to obtain approximate analytical expressions for the spectral quantities. For this, we based ourselves heavily on [125]. The analytical results must then be adapted to Gaussian white noise inputs of unlimited bandwidth, which will be the external stimuli used for most of this paper to allow comparison with [125]. We then present the

results of numerical simulations together with the LRT for the various cases mentioned in the preceding paragraph. For the symmetric case, we compare the behavior of the ON-only and the ON/OFF networks. The symmetry between the ON and OFF cells is first broken by changing the bias and the intensity of the internal noise of OFF cells. Also we change the time scale over which the OFF cells evolve with respect to the ON cells, together with an increase of the OFF cells' bias. Beyond the context of electric fish, our results illustrate how LRT can be applied to nonhomogeneous neural nets.

3.2 Model

We consider a network of $2N$ LIF neurons distributed into N ON/OFF pairs, as illustrated in Fig. 3.1. The relevant dynamical variables are the subthreshold membrane potentials and the quantities of interest are the spike trains for each neuron. The LIF scheme prescribes that a spike occurs whenever a given threshold V_T is reached by a membrane potential V. After a refractory time, the membrane potential is reset to a potential V_R. The dynamics of the subthreshold potentials are of the form

$$C_m \dot{V} = -g_L(V - E_L) + \text{noise},$$

with C_m the membrane capacitance, g_L the leak conductance and E_L the reversal potential. In what follows, we will use nondimensionalized versions of the equation above [125]. Namely, we define the bias current μ and the potential v:

$$\mu \equiv \frac{E_L - V_R}{V_T - V_R}, \qquad v \equiv \frac{V - V_R}{V_T - V_R}. \tag{3.1}$$

With this rescaling, the threshold becomes $v_T = 1$, the reset $v_R = 0$ and time is measured in units of the membrane constant, C_m/g_L, which is supposed to be the same for ON and OFF cells for now. The refractory time will be denoted τ_R.

In the absence of any inputs and feedback, the subthreshold spontaneous activities of the ON and OFF neurons are given by:

$$\begin{aligned} \dot{v}_i^{\mathrm{ON}}(t) &= -v_i^{\mathrm{ON}}(t) + \mu + \xi_i^{\mathrm{ON}}(t) \\ \dot{v}_i^{\mathrm{OFF}}(t) &= -v_i^{\mathrm{OFF}}(t) + \mu + \xi_i^{\mathrm{OFF}}(t) + V_0, \end{aligned} \tag{3.2}$$

where $\xi_i^{\mathrm{ON}}(t)$ and $\xi_i^{\mathrm{OFF}}(t)$ are internal noises, and V_0 is the offset of the OFF neurons' activity with respect to the ON neurons [120]. It accounts for the difference in baseline activities of

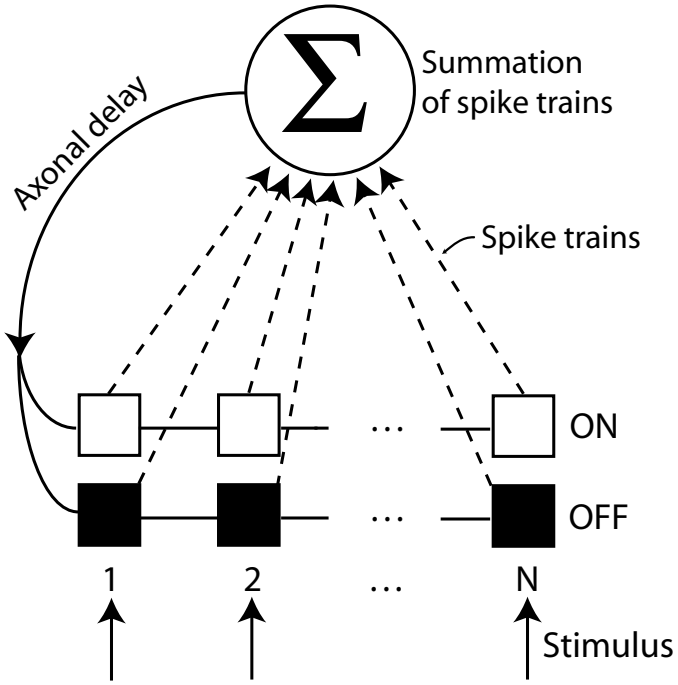

Figure 3.1: Schema of the network. Both ON and OFF cells receive external stimuli. Whenever the firing threshold is reached, a spike is fired. A remote kernel sums the spike trains and feeds them back to the network with a delay in an all-to-all fashion.

ON and OFF cells. According to Krahe et al. [114], this difference is small for baseline rates measured in vivo, without external stimulation and with the feedback intact but not causing oscillations. The internal noises are white, zero-mean and Gaussian with intensity D, i.e.

$$\langle \xi_i^k(t) \rangle = 0 \quad \text{and} \quad \langle \xi_i^k(t) \xi_j^l(t') \rangle = 2D\delta_{ij}\delta_{kl}\delta(t-t'),$$

for $i, j = 1, \ldots, N$ and $k, l = \text{ON}, \text{OFF}$. Note that the intrinsic noise intensities for ON and OFF cells are identical for the moment, but they will be different in Sec. 3.5.2. The brackets $\langle \ldots \rangle$ will always denote ensemble averages.

Sensory pathways to the OFF neurons comprise an interneuron population [29]. It has the approximate effect of inverting any input coming from the electrosensory afferents. Here, we consider a stochastic input of the form

$$I_i(t) = m + \zeta_i(t)$$

with

$$\zeta_i(t) \equiv \sqrt{c}\eta_c(t) + \sqrt{1-c}\eta_i(t),$$

where η_c is the correlated external noise since it is the same for all neurons, η_i is the uncorrelated external noise and $m = \langle I_i(t) \rangle$ is the mean stochastic input; all external noises obey the same statistics. For weakly electric fish the stimuli are zero-mean amplitude modulations of their own electric discharge, so we set $m = 0$. It is assumed that these modulations are linearly coded by neural populations innervating the ELL cells [116]. All noise processes are uncorrelated with each other, i.e. $\langle \eta_c \eta_i \rangle = \langle \eta_i \eta_j \rangle = 0$ for all i, j $(i \neq j)$ at all times. Note that η_i is the same for each cell in an ON/OFF pair. The correlation parameter $c \in [0, 1]$ determines the 'weight' of each of the external noise components: the input is fully correlated (resp. fully uncorrelated) if $c = 1$ (resp. $c = 0$). The prefactors $\sqrt{c}$ and $\sqrt{1-c}$ are chosen so that the total input power does not depend on c.

ON and OFF pyramidal cells of the ELL also receive feedback signals from a distant neural population. For simplicity, the signal processing performed by these neurons is a mere average of all the spike trains received from the network, taking into account the axonal propagation time and the filtering due to synaptic transport. The feedback is inhibitory and diffuse, i.e. each cell of the network receives at time t the following feedback input:

$$f(t) = \frac{G}{2N} \int_{\tau_D}^{\infty} \frac{d\tau}{\tau_S} \left(\frac{\tau - \tau_D}{\tau_S} \right) \exp\left(-\frac{\tau - \tau_D}{\tau_S} \right) \sum_{i=1}^{N} \left[y_i^{\mathrm{ON}}(t-\tau) + y_i^{\mathrm{OFF}}(t-\tau) \right], \qquad (3.3)$$

where $G < 0$ is the feedback strength (rendering $f(t)$ a negative quantity), τ_S is the decay time and τ_D is the delay. The spike trains $y_i^{\mathrm{ON}}(t)$ and $y_i^{\mathrm{OFF}}(t)$ are given by

$$y_i^{\epsilon}(t) = \sum_{t_k} \delta(t - t_k), \qquad (3.4)$$

where the t_k's are the spike times of neuron i $(i = 1, \ldots, N)$ of class ϵ $(\epsilon = \mathrm{ON}, \mathrm{OFF})$. Thus, the feedback is G times the convolution of a delayed alpha function ($\Theta(\cdot) =$ Heaviside function)

$$\alpha(\tau) = \frac{1}{\tau_S} \left(\frac{\tau - \tau_D}{\tau_S} \right) \exp\left(-\frac{\tau - \tau_D}{\tau_S} \right) \Theta(\tau - \tau_D)$$

with the population average of the spike trains. The full model is then:

$$\begin{aligned} \dot{v}_i^{\mathrm{ON}} &= -v_i^{\mathrm{ON}} + \mu + \xi_i^{\mathrm{ON}} + \zeta_i + f \\ \dot{v}_i^{\mathrm{OFF}} &= -v_i^{\mathrm{OFF}} + \mu + V_0 + \xi_i^{\mathrm{OFF}} - \zeta_i + f \\ v(t_k^-) = v_T &\Rightarrow t_{\mathrm{spike}} = t_k \text{ and } v(t_k^+ + \tau_R) = v_R, \end{aligned} \qquad (3.5)$$

where we have included the LIF scheme requirements and omitted the dependence of all functions on time for clarity. Note the minus sign in front of ζ_i for the OFF cell equation. A feedforward version ($G = 0$) of this model with ON cells only has been studied in [71, 191], with particular emphasis on the enhancement of pairwise correlations of neurons following an increase of their firing rates.

In this chapter, we will mainly be concerned with single-neuron power spectra, cross spectra and population spectra. The ON and OFF subnetworks are assumed to be separately homogeneous. Therefore, the power spectra are the same for all neurons of the same type and the cross spectra are identical for each neuron pair of the same type. The power spectrum of any neuron is given by

$$S(\omega) = \lim_{T\to\infty} \frac{\langle |\tilde{y}(\omega)|^2 \rangle}{T}, \tag{3.6}$$

with $\tilde{y}(\omega) = \int_0^T dt\, y(t) e^{i\omega t}$ (we have dropped the ON and OFF superscripts for clarity). The cross spectrum of two neurons is given by

$$S_{\text{cross}} = \lim_{T\to\infty} \frac{\langle \tilde{y}_i^*(\omega) \tilde{y}_j(\omega) \rangle}{T}, \tag{3.7}$$

where both neurons belong to the same class. The average population activity for either the ON or the OFF subnetwork is

$$Y^{\text{ON(OFF)}}(t) = \frac{1}{N} \sum_{i=1}^{N} y_i^{\text{ON(OFF)}}(t),$$

and the population spectra are just the power spectra of these mean activities. Using the homogeneity of the network, the population spectra of either ON or OFF cells can be expressed in terms of S_{cross} and S [125], namely:

$$S_{\text{pop}}(\omega) = S_{\text{cross}}(\omega) + \frac{S(\omega) - S_{\text{cross}}(\omega)}{N}. \tag{3.8}$$

3.3 Linear response theory

Approximate results for the spectral quantities can be obtained using linear response theory (LRT). We write the time evolution of the subthreshold membrane potentials as

$$\begin{aligned} \dot{v}_i^{\text{ON}} &= -v_i^{\text{ON}} + \mu + \langle f \rangle + \xi_i^{\text{ON}} \underline{+ \zeta_i + (f - \langle f \rangle)} \\ \dot{v}_i^{\text{OFF}} &= -v_i^{\text{OFF}} + \mu + V_0 + \langle f \rangle + \xi_i^{\text{OFF}} \underline{- \zeta_i + (f - \langle f \rangle)}. \end{aligned}$$

The underlined terms $\pm\zeta_i + (f - \langle f \rangle)$ are taken as perturbations of the system. The average value of the feedback is $\langle f \rangle = G(r_{\mathrm{ON}} + r_{\mathrm{OFF}})/2$, with $r_{\mathrm{ON(OFF)}} = \langle y_i^{\mathrm{ON(OFF)}} \rangle$ being the time-independent mean firing rate of an ON (OFF) neuron [125]. It means that the unperturbed cells have effective biases given by

$$\begin{aligned} \mu_{\mathrm{ON}} &= \mu + \frac{G}{2}(r_{\mathrm{ON}} + r_{\mathrm{OFF}}) \\ \mu_{\mathrm{OFF}} &= \mu + V_0 + \frac{G}{2}(r_{\mathrm{ON}} + r_{\mathrm{OFF}}). \end{aligned} \tag{3.9}$$

The unperturbed system comprises four relevant parameters, namely the internal noise intensity D, the effective bias currents μ_{ON} and μ_{OFF} and the refractory period τ_R. For $\mu_{\mathrm{ON}} < v_T$ and $\mu_{\mathrm{OFF}} < v_T$, the neurons are in a noise-activated regime [126]. In this case, which will be considered in this paper (except otherwise mentioned), the neurons do not fire at vanishing noise ($r_{\mathrm{ON}} = r_{\mathrm{OFF}} = 0$) so that the resting potentials are $\bar{v}^{\mathrm{ON}} = \mu$ and $\bar{v}^{\mathrm{OFF}} = \mu + V_0$. Therefore, increasing (resp. decreasing) V_0 will favor (resp. hinder) the OFF cells' firing once the noise is set back to a nonzero value. Given the importance of noise in this regime for the neuron firing, it is worth mentioning that the system actually exhibits coherence resonance for some sets of values of the parameters [126]. The cases $\mu_{\mathrm{ON}} > v_T$ and $\mu_{\mathrm{OFF}} > v_T$ are called deterministic firing regimes because the neurons fire in absence of noise.

3.3.1 Unperturbed network

The deviation of the feedback from its mean $(f - \langle f \rangle)$ and the input ζ_i are taken as perturbations of the unperturbed system. When the variances of both these quantities are small, one may assume that the firing rates of the unperturbed system are good approximations to those of the perturbed system [125]. Concretely, this implies that the rates appearing in Eq. (3.9) may be replaced by the rates obtained from the treatment to follow. A 'free' LIF neuron (in open loop and without external input) has a firing rate given by [126]

$$\nu(\mu, D) = \left(\tau_R + \sqrt{\pi} \int_{(\mu - v_T)/\sqrt{2D}}^{(\mu - v_R)/\sqrt{2D}} dx \, \mathrm{e}^{x^2} \mathrm{erfc}(x) \right)^{-1}. \tag{3.10}$$

To get the firing rates of the unperturbed cells, we must replace μ in the above equation by the effective biases of Eq. (3.9). It results that the rates r_{ON} and r_{OFF} must be determined

self-consistently from the coupled equations:

$$\begin{aligned} r_{\mathrm{ON}} &= \nu(\mu + G(r_{\mathrm{ON}} + r_{\mathrm{OFF}})/2, D) \\ r_{\mathrm{OFF}} &= \nu(\mu + V_0 + G(r_{\mathrm{OFF}} + r_{\mathrm{OFF}})/2, D). \end{aligned} \tag{3.11}$$

Analytical expressions for the unperturbed power spectra ($S^{(0)}_{\mathrm{ON}}$ and $S^{(0)}_{\mathrm{OFF}}$) and the susceptibilities ($A(\omega, \mu_{\mathrm{ON}}, D)$ and $A(\omega, \mu_{\mathrm{OFF}}, D)$) are known [125, 123]:

$$S^{(0)}_{\epsilon} = 2\pi r_{\epsilon}^2 \delta(\omega) + r_{\epsilon} \frac{\left|\mathcal{D}_{i\omega}(\frac{\mu_{\epsilon}-v_T}{\sqrt{D}})\right|^2 - \mathrm{e}^{2\Delta_{\epsilon}}\left|\mathcal{D}_{i\omega}(\frac{\mu_{\epsilon}-v_R}{\sqrt{D}})\right|^2}{\left|\mathcal{D}_{i\omega}(\frac{\mu_{\epsilon}-v_T}{\sqrt{D}}) - \mathrm{e}^{\Delta_{\epsilon}}\mathrm{e}^{i\omega\tau_R}\mathcal{D}_{i\omega}(\frac{\mu_{\epsilon}-v_R}{\sqrt{D}})\right|^2}, \tag{3.12}$$

$$A(\omega, \mu_{\epsilon}, D) = \left(\frac{i\omega r_{\epsilon}}{\sqrt{D}(i\omega - 1)}\right)\left[\frac{\mathcal{D}_{i\omega-1}(\frac{\mu_{\epsilon}-v_T}{\sqrt{D}}) - \mathrm{e}^{\Delta_{\epsilon}}\mathcal{D}_{i\omega-1}(\frac{\mu_{\epsilon}-v_R}{\sqrt{D}})}{\mathcal{D}_{i\omega}(\frac{\mu_{\epsilon}-v_T}{\sqrt{D}}) - \mathrm{e}^{\Delta_{\epsilon}}\mathrm{e}^{i\omega\tau_R}\mathcal{D}_{i\omega}(\frac{\mu_{\epsilon}-v_R}{\sqrt{D}})}\right], \tag{3.13}$$

with $\Delta_{\epsilon} = [v_R^2 - v_T^2 + 2\mu_{\epsilon}(v_T - v_R)]/4D$ and $\epsilon \in \{\mathrm{ON}, \mathrm{OFF}\}$. The functions $\mathrm{erfc}(x)$ and $\mathcal{D}_a(x)$ are the complementary error function and the parabolic cylinder functions, respectively.

3.3.2 Response to the perturbation

A LRT is possible if both the amplitude of the external stimulus and the variance of the feedback are not too large. Then, provided that the internal noise is large, the following ansatz can be made for the perturbed spike trains [9, 10, 125, 124, 55]:

$$y_i^{\epsilon}(t) = y^{(0)}_{\epsilon,i}(t) + [A(\mu_{\epsilon}, D) \star (\epsilon\zeta_i + f - \langle f \rangle)](t). \tag{3.14}$$

Here and in the following, ϵ as a subscript or superscript will stand for ON or OFF; as a multiplicative constant, it stands for $+$ in an expression for an ON variable or $-$ for an OFF variable. The $\star$ denotes the convolution operation and the $y^{(0)}_{\epsilon,i}$'s are the unperturbed spike trains. Taking the Fourier transform on each side, we get for $\omega \neq 0$ [1]:

$$\tilde{y}_i^{\epsilon}(\omega) = \tilde{y}^{(0)}_{\epsilon,i}(\omega) + A(\omega, \mu_{\epsilon}, D)\left[\epsilon\tilde{\zeta}_i(\omega) + \frac{F(\omega)}{2N}\left(\sum_{k=1}^{N}\tilde{y}_k^{\mathrm{ON}}(\omega) + \sum_{k=1}^{N}\tilde{y}_k^{\mathrm{OFF}}(\omega)\right)\right] \tag{3.15}$$

with

$$F(\omega) = G\alpha(\omega) = G\frac{\mathrm{e}^{i\omega\tau_D}}{(1 - i\omega\tau_S)^2} \tag{3.16}$$

[1] We consider $\omega \neq 0$ to avoid the cumbersome delta functions that appear in the process.

as the Fourier transform of the delayed alpha function times the feedback strength G.

3.3.3 Calculation of spectral quantities

To compute the relevant spectral quantities, we first form the subsidiary sums:

$$\tilde{Y}_\epsilon^{(0)} = \frac{1}{N}\sum_{i=1}^{N} \tilde{y}_{\epsilon,i}^{(0)}, \quad \tilde{Y}_\epsilon = \frac{1}{N}\sum_{i=1}^{N} \tilde{y}_i^{\epsilon}, \quad \tilde{\zeta} = \frac{1}{N}\sum_{i=1}^{N} \tilde{\zeta}_i. \tag{3.17}$$

Summing Eq. (3.15) over i from 1 to N and dividing by N, we get, writing $A(\omega, \mu_\epsilon, D) \equiv A_\epsilon$:

$$\tilde{Y}_\epsilon = \tilde{Y}_\epsilon^{(0)} + \epsilon A_\epsilon \tilde{\zeta} + \frac{A_\epsilon F}{2}(\tilde{Y}_{\mathrm{ON}} + \tilde{Y}_{\mathrm{OFF}}).$$

These two equations (one for each value of ϵ) can be used to get an expression for the sum $\tilde{Y}_{\mathrm{ON}} + \tilde{Y}_{\mathrm{OFF}}$ intervening in Eq. (3.15). We now have workable expressions for the Fourier transform of the spike trains:

$$\tilde{y}_i^{\epsilon} = \tilde{y}_{\epsilon,i}^{(0)} + A_\epsilon \left[\epsilon\tilde{\zeta}_i + \gamma(\tilde{Y}_{\mathrm{ON}}^{(0)} + \tilde{Y}_{\mathrm{OFF}}^{(0)}) + \gamma(A_{\mathrm{ON}} - A_{\mathrm{OFF}})\tilde{\zeta}\right] \tag{3.18}$$

with

$$\gamma = \frac{F/2}{1 - \frac{(A_{\mathrm{ON}} + A_{\mathrm{OFF}})}{2}F}. \tag{3.19}$$

From Eq. (3.18), it is seen that the perturbed spike train is made of the response of the single neuron to both its external input and to the network's overall activity processed by the feedback kernel.

The following properties are essential to the computation of the spectral quantities:

1. Unperturbed spike trains are uncorrelated with the external stimuli:

$$\lim_{T\to\infty} \frac{\langle \tilde{y}_{\epsilon,i}^{(0)*} \tilde{\eta}_j \rangle}{T} = \lim_{T\to\infty} \frac{\langle \tilde{y}_{\epsilon,i}^{(0)*} \tilde{\eta}_c \rangle}{T} = 0, \ \forall i, j.$$

2. Unperturbed spike trains are uncorrelated among each other:

$$\lim_{T\to\infty} \frac{\langle \tilde{y}_{\epsilon,i}^{(0)*} \tilde{y}_{\epsilon,j}^{(0)} \rangle}{T} = S_\epsilon^{(0)} \delta_{ij}, \quad \lim_{T\to\infty} \frac{\langle \tilde{y}_{\epsilon,i}^{(0)*} \tilde{y}_{-\epsilon,j}^{(0)} \rangle}{T} = 0, \ \forall i, j,$$

where we have written $-\epsilon$ to signify that if $\epsilon = \mathrm{ON}$, $-\epsilon = \mathrm{OFF}$ and conversely.

3. The external stimuli satisfy:

$$\lim_{T\to\infty}\frac{\langle\tilde{\eta}_i^*\tilde{\eta}_j\rangle}{T}=S_{\text{st}}\delta_{ij},\quad \lim_{T\to\infty}\frac{\langle|\tilde{\eta}_c|^2\rangle}{T}=S_{\text{st}},$$
$$\lim_{T\to\infty}\frac{\langle\tilde{\eta}_i^*\tilde{\eta}_c\rangle}{T}=0,\ \forall i,j,$$

with S_{st} the statistics of the noise inputs, which will be described in Sec. 3.4.

After algebraic manipulations, the single-neuron power spectra read (Re $(\cdot)$ = real part of $\cdot$):

$$\begin{aligned}S_\epsilon = S_\epsilon^{(0)}\left[1+\frac{2}{N}\,\text{Re}\,\left(\gamma A_\epsilon\right)+\frac{|\gamma A_\epsilon|^2}{N}\right]+S_{-\epsilon}^{(0)}\frac{|\gamma A_\epsilon|^2}{N}\\ +|A_\epsilon|^2 S_{st}\left\{1+\left[\frac{1}{N}(1-c)+c\right]\left[2\,\text{Re}\,\left[\gamma(A_\epsilon-A_{-\epsilon})\right]+|\gamma(A_\epsilon-A_{-\epsilon})|^2\right]\right\},\end{aligned}\tag{3.20}$$

where we have used the facts that $\epsilon(A_{\text{ON}}-A_{\text{OFF}}) = A_\epsilon - A_{-\epsilon}$ and $|A_{\text{ON}}-A_{\text{OFF}}|^2 = |A_\epsilon - A_{-\epsilon}|^2$. Again $\omega\neq 0$, and $S_\epsilon^{(0)}$ stands for the second term only of Eq. (3.12). In the large N limit, the only terms remaining are the power spectrum of the response of a single neuron to the total external noise ($S_\epsilon^{(0)}+|A_\epsilon|^2 S_{st}$) and the term

$$c|A_\epsilon|^2 S_{st}\left[2\,\text{Re}\,\left[\gamma(A_\epsilon-A_{-\epsilon})\right]+|\gamma(A_\epsilon-A_{-\epsilon})|^2\right],\tag{3.21}$$

which represents the effect of the network response to the correlated noise. The contribution from the identically distributed uncorrelated noises filtered by the feedback kernel vanishes in this limit. Still in this limit, $c=0$ causes the feedback to have the sole effect of shifting the bias current, i.e. only the static part of the feedback matters.

If the offset $V_0 = 0$, $\mu_{\text{ON}} = \mu_{\text{OFF}}$. Hence $A_{\text{ON}} = A_{\text{OFF}}$ and $S_{\text{ON}}^{(0)} = S_{\text{OFF}}^{(0)}$, and the correlation parameter c has no effect on the spectrum, even at finite N. We thus see that within the limits of applicability of the LRT, introducing the OFF cells into the system cancels the effect of the spatial correlation of the external input.

The cross spectrum of two ON or two OFF neurons, noted $S^{\epsilon}_{\text{cross}}$, is

$$\begin{aligned} S^{\epsilon}_{\text{cross}} &= S^{(0)}_{\epsilon}\left[\frac{2}{N}\,\text{Re}\;(\gamma A_{\epsilon}) + \frac{|\gamma A_{\epsilon}|^2}{N}\right] + S^{(0)}_{-\epsilon}\frac{|\gamma A_{\epsilon}|^2}{N} \\ &\quad + |A_{\epsilon}|^2 S_{st}\left\{c + \left[\frac{1}{N}(1-c)+c\right]\left[2\,\text{Re}\;[\gamma(A_{\epsilon}-A_{-\epsilon})] + |\gamma(A_{\epsilon}-A_{-\epsilon})|^2\right]\right\} \\ &= S_{\epsilon} - S^{(0)}_{\epsilon} - (1-c)|A_{\epsilon}|^2 S_{\text{st}}. \end{aligned} \tag{3.22}$$

The cross spectrum is thus given by the difference between the perturbed power spectrum, on the one hand, and the unperturbed spectrum of the neuron plus the transmitted uncorrelated noise on the other hand. It means that the feedback term has the same effect on both the single-neuron power spectrum and the cross spectrum, just like for the ON-only network [125]. For a fully correlated input ($c = 1$), $S^{\epsilon}_{\text{cross}}$ is a measure of the difference between the perturbed and the unperturbed single-neuron spectra.

3.4 Theory for Gaussian white noise stimuli of unlimited bandwidth

Until now, the statistics of the external noises have been left unspecified. For the LRT to be valid, both the variance of these noises and the feedback strength must be small compared to the internal noise. Following [125], we will paradoxically consider Gaussian white noise stimuli of unlimited bandwidth for both the uncorrelated and the correlated noises:

$$\langle \eta_c(t)\eta_c(t')\rangle = \langle \eta_i(t)\eta_i(t')\rangle = 2D_E\delta(t-t'). \tag{3.23}$$

The variance is infinite for white noise and thus does not meet the above requirements. Lindner and co-workers [125] circumvent this difficult by making the following ansatz: whenever $S^{(0)}_{\epsilon}(\omega, D) + 2D_E|A_{\epsilon}(\omega, D)|^2$ appears in the expressions for the spectral measures, one may replace it approximately by

$$S^{(0)}_{\epsilon}(\omega, Q) \equiv S^{(0)}_{\epsilon,Q}, \tag{3.24}$$

with $Q = D + D_E$ the total noise intensity perceived by the neurons. It is also understood that the susceptibilities $A_{\epsilon}(\omega, D)$ must be replaced by $A_{\epsilon}(\omega, Q) \equiv A_{\epsilon,Q}$, and $\gamma(\omega, D)$ becomes

$$\gamma_Q = \frac{F/2}{1 - \frac{A_{\text{ON},Q}+A_{\text{OFF},Q}}{2}F}. \tag{3.25}$$

The firing rates also need to be evaluated at noise intensity Q: $r_\epsilon^{(0)}(\mu_\epsilon, D) \to r_\epsilon^{(0)}(\mu_\epsilon, Q)$. The informal argument behind these substitutions is that a neuron should not make the difference between the statistics of external and internal noises if they are the same. See Lindner et al. [125] for more details. Doing so, we get for the single neuron power spectra:

$$\begin{aligned} S_{\epsilon,Q}(\omega) = S_{\epsilon,Q}^{(0)} \left(1 + \frac{2}{N}\,\mathrm{Re}\,(\gamma_Q A_{\epsilon,Q}) + \frac{1}{N}|\gamma_Q A_{\epsilon,Q}|^2\right) + S_{-\epsilon,Q}^{(0)} \frac{1}{N}|\gamma_Q A_{\epsilon,Q}|^2 \\ + 2D_E|A_{\epsilon,Q}|^2 \left(\frac{1}{N}(1-c)+c\right)\left[2\,\mathrm{Re}\,(\gamma_Q(A_{\epsilon,Q} - A_{-\epsilon,Q})) + |\gamma_Q(A_{\epsilon,Q} - A_{-\epsilon,Q})|^2\right] \\ - 2D_E|A_{\epsilon,Q}|^2\left[\frac{2}{N}\,\mathrm{Re}\,(\gamma_Q A_{\epsilon,Q}) + \frac{1}{N}|\gamma_Q A_{\epsilon,Q}|^2 + \frac{1}{N}|\gamma_Q A_{-\epsilon,Q}|^2\right]. \end{aligned} \tag{3.26}$$

Cross spectra of neurons within a given subpopulation are readily computed by making the appropriate substitutions in Eq. (3.22):

$$S_{\mathrm{cross},Q}^{\epsilon} = S_{\epsilon,Q} - S_{\epsilon,Q}^{(0)} + 2cD_E|A_{\epsilon,Q}|^2, \tag{3.27}$$

and likewise for the population spectra $S_{\mathrm{pop}}^{\epsilon}$.

Numerical integrations of Eq. (3.5) for a particular set of parameters were done using a simple Euler scheme with $\Delta t = 5 \times 10^{-4}$ and between 100 and 200 realizations. The Gaussian white noise used in the simulations is intrinsically bandlimited to the Nyquist frequency, which is 2×10^5 Hz for a membrane time constant of 5 ms.

3.5 Results

3.5.1 Symmetric case

Above, we noted that for $V_0 = 0$, the ON and OFF susceptibilities and unperturbed spectra are equal. In this symmetric case, Eq. (3.20) becomes (replacing γ by its expression and rearranging):

$$S_\epsilon = S_\epsilon^{(0)} + |A_\epsilon|^2 S_{st} + \frac{S_\epsilon^{(0)}}{2N}\frac{2\,\mathrm{Re}\,(A_\epsilon F) - |A_\epsilon F|^2}{|1 - A_\epsilon F|^2}. \tag{3.28}$$

We compare the above equation with $\epsilon = \mathrm{ON}$ to the LRT power spectrum for the ON-only

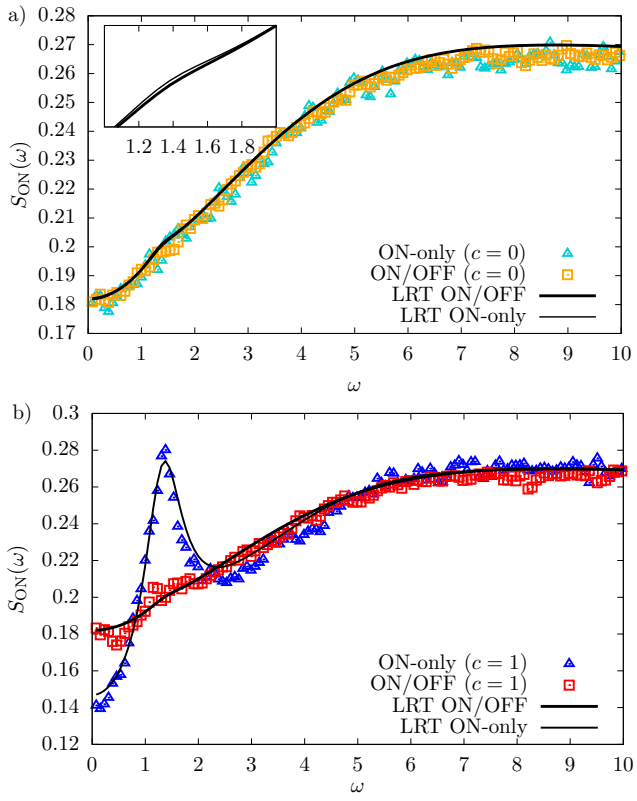

Figure 3.2: Impact of OFF cells on gamma oscillations. Power spectra of ON cells for both the ON-only network ($N = 100$) and the ON/OFF network ($N_{\rm ON} = N_{\rm OFF} = 50$) are displayed for uncorrelated [$c = 0$, plot (a)] and correlated [$c = 1$, plot (b)] inputs. The inset in plot (a) shows the difference between the LRT curves for the ON-only and the ON/OFF networks for $1 < \omega < 2$. Parameters are: $\mu = 0.8$, $D = 0.12$, $D_E = 0.08$, $V_0 = 0$, $\tau_D = 1$, $\tau_S = 0.5$, $\tau_R = 0.1$ and $G = -1.2$; these are the same values as in Fig. 2 of [125] (note the labeling error on that figure). The triangles and squares are from numerical simulations; LRT = Linear Response Theory.

network ([125], equation 22) that we reproduce here for convenience:

$$S = S_0 + |A|^2 S_{st} + c|A|^2 S_{\rm st} \frac{2\,{\rm Re}\,(AF) - |AF|^2}{|1 - AF|^2} + \frac{1}{N}[S_0 + (1-c)|A|^2 S_{\rm st}] \frac{2\,{\rm Re}\,(AF) - |AF|^2}{|1 - AF|^2}.$$

Here, N is the total number of neurons whereas N in Eq. (3.28) is the number of cells of a given type. For $c = 0$ and large N, both expressions are similar if our N is chosen to be equal to half the number of cells in the ON-only network. Figure 3.2a illustrates this fact: the LRT curves are superposed except for a small interval near $\omega = 1.5$, where the LRT for the

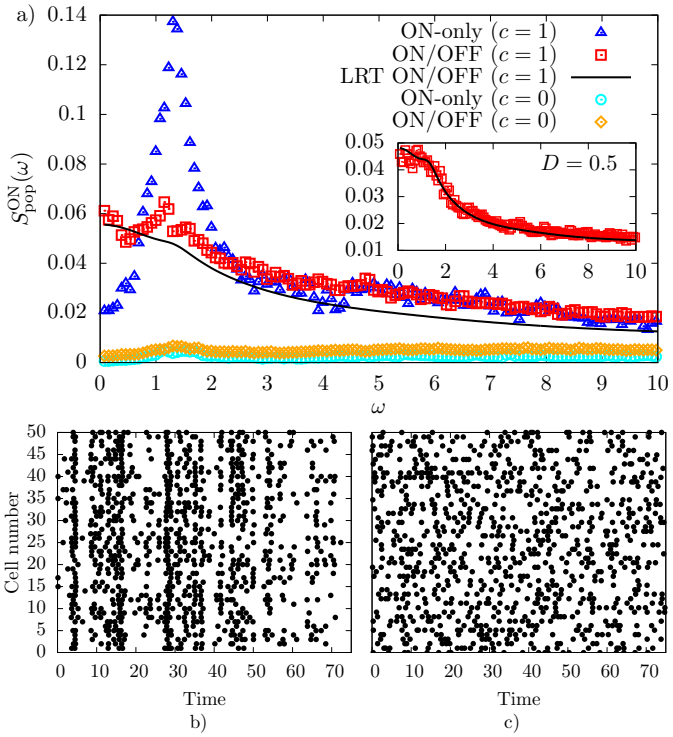

Figure 3.3: Impact of OFF cells on population activity. (a): Comparison of ON population spectra for the ON-only and ON/OFF networks. Parameters are the same as in Fig. 3.2. The only LRT curve displayed is for the ON/OFF network with $c = 1$ (see text for explanation). Inset: LRT for higher internal noise. (b) and (c): Raster plots for ON neurons in the ON/OFF network with $c = 1$ (b) and $c = 0$ (c).

ON-only network is very slightly higher (see inset). On this same plot the corresponding small bump in the numerical curve for the ON/OFF network is difficult to see.

For the case $c = 1$, a drastic contrast is seen between the network types, as depicted in Fig. 3.2b. As in [125], the ON-only network undergoes strong gamma oscillations, embodied in the strong peak at $\omega = 1.5$. However, they are nearly absent for the ON/OFF network. Note that for a membrane time constant of 5 ms, the frequency of the peak corresponds to 48 Hz (for comparison, [77] uses 6 ms and [75], 10 ms). The agreement between the LRT and the numerical results is excellent, although there are minor quantitative discrepancies at low frequency (especially in the case $c = 1$).

At this point, it is also worth noting that adding more neurons yields weaker oscillations in both the ON-only [125] and the ON/OFF network (not shown), although they survive for the ON-only network for $N \to \infty$. So, had we chosen a different number of neurons of each type (e.g. setting $N_{ON} = N_{OFF} = 100$), the conclusion regarding the weakening of

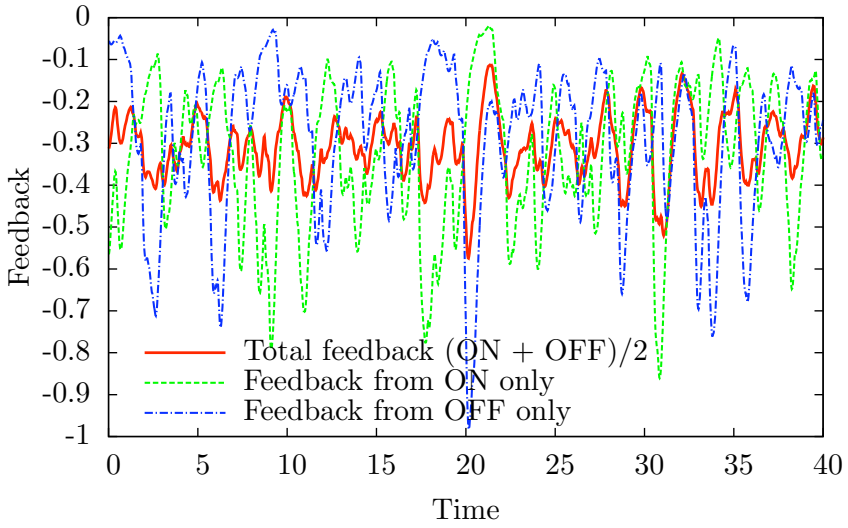

Figure 3.4: Time course of the feedback due to ON (green dashed curve) and OFF (blue dotted-dashed curve) cells and of the total feedback (solid red curve) during one realization. Observe that the peaks and troughs of the blue (dotted-dashed) and green (dashed) traces are roughly in antiphase.

oscillations would have been qualitatively the same.

Perhaps a better way to ascertain the occurrence (or not) of network oscillations is to consider population spectra, as is done in Fig. 3.3a. For a correlated input, one can see a high peak at $\omega = 1.5$ for the ON-only spectrum and the significantly lower peak at about the same frequency for the ON/OFF network. The correlated input triggers high correlations between the spike trains for the ON-only network; for the ON/OFF network, the correlations between ON cells are impeded by the OFF cells, which respond in an opposite way to the external Gaussian noise. In the case $c = 0$ the power is small at each frequency (lowest curves). The slight bump around $\omega = 1.5$ does not represent significant collective oscillations since the power is everywhere small. The vanishing of network oscillations is confirmed by looking at the raster plots in Figs. 3.3b and 3.3c: no stripes of quasi-synchronous spikes are seen in the case $c = 0$.

For a large number of neurons, in both the ON/OFF and ON-only cases, the population spectrum is approximately equal to the cross spectrum (see Eq. (3.8)). Since the cross spectrum is inherently a two-neuron quantity, and given that its evaluation through the LRT implies the use of single-neuron quantities only, the LRT usually fails to reproduce the population spectra (see again [125] for a thorough discussion). An example is provided in Fig. 3.3a (compare the red squares with the black curve). A better fit is obtained for

higher internal noise intensities, as exemplified in the inset. This is due to a decrease of the correlations and to the linearization effect of a high internal noise, so that a LRT will represent more adequately the behavior of the system.

The reason behind the vanishing of network oscillations in the symmetric case is that the ON and OFF cells are totally opposite in their spectral response: when, say, an ON cell has just received the sufficient input to cross the threshold, the probability that the OFF cell in the same pair will fire should be low. This is perhaps best envisioned in Fig. 3.4, where we have plotted the feedback from the ON and the OFF cells for a single trial, together with the total feedback from the kernel. The feedback due to the ON (resp. OFF) cells is a low-pass filtered version of the time-varying ON (resp. OFF) firing rate. The oscillations of these feedbacks are roughly in antiphase and produce a reduced modulation of the network's activity. To put it into the language of [75], the waves of inhibition coming from the feedback kernel are reduced in amplitude due to the antagonistic response of the two classes of cells to external inputs. We further add that, from the dynamical systems perspective, it has been shown that the inclusion of OFF-type responses in continuous neural networks shifts the threshold for oscillatory instabilities under deterministic stimulation [120]. We conclude that we observe here the same phenomenon in a fully noisy and spiking framework.

In the following subsections, we try to get back the strong gamma oscillations seen for the ON-only network under global stimulation by introducing asymmetries between the ON and the OFF cells. First, we will set the offset $V_0 > 0$ with a reduced intrinsic noise for the OFF cells. Then, we will consider unequal membrane time constants, again with $V_0 > 0$. The latter case requires a more delicate analytical treatment. Note that the parameters $G = -1.2$, $\tau_D = 1$, $\tau_S = 0.5$ and $\tau_R = 0.1$ will stay fixed to these values for the remainder of this chapter to allow comparison with results from the symmetric case and the ON-only network. However, we would expect our conclusions to be more generally valid. From [125], we know that gamma oscillations are present over a certain part of parameter space for the ON-only network. Suppose that there is a gamma peak in the ON-only case for a given set of parameters (with $c = 1$). Then, provided that these parameter values are such that the LRT remains a valid approximation, Eq. 3.20 stipulates that gamma-band oscillations should be greatly reduced for the corresponding ON/OFF network in the symmetric case. Moreover, the symmetry breakings discussed below have precise effects on ON and OFF cells. Although we did not explore the full extent of parameter space in search of counterexamples, we suspect these effects — on the grounds of the LRT – to be generic rather than mere artifacts of the given choice of parameter values.

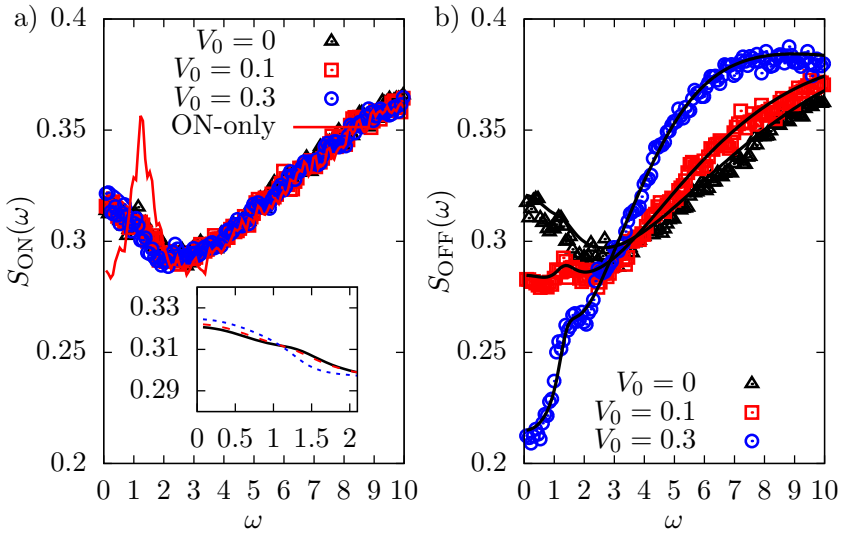

Figure 3.5: ON (a) and OFF (b) power spectra for different values of V_0, with the ON internal noise intensity $D_{\mathrm{ON}} = 0.36$ and $c = 1$. The OFF cells' internal noise intensity, D_{OFF}, is decreased so that the rates remain the same within an absolute error of order 10^{-3}. The values of D_{OFF} for $V_0 = 0, 0.1, 0.3$ are respectively $D_{\mathrm{OFF}} = 0.36, 0.27, 0.125$. All other parameters are as in Fig. 3.2. On plot (a), we also show the power spectrum for the corresponding ON-only network (solid red curve with a distinguishable peak). The inset shows the difference between the LRT curves around $\omega = 1.5$ (solid black: $V_0 = 0$; dashed red: $V_0 = 0.1$; and dotted blue: $V_0 = 0.3$). Other LRT curves are black and fit the numerical results. Note that the effective bias currents stay well below v_T for all curves.

3.5.2 Asymmetry between ON and OFF cells: Offset and OFF cells' intrinsic noise

As stated in Sec. 3.2, physiological experiments have shown that the spontaneous — without external stimulation and with intact feedback — firing rates of ON and OFF cells are approximately the same on average [114]. Also, on average, ON cells have higher spike thresholds (V_T) than OFF cells [218]. For comparable resting and reset potentials for both classes of cells, it would mean, because of Eq. (3.1), that the ON bias current is smaller than that of OFF cells. To account for the larger bias of OFF cells, we try in the present section to increase the offset V_0 of OFF cells. As shown in Appendix 3.7, the ON (resp. OFF) rate is a monotonically decreasing (resp. increasing) function of V_0. For the rates of ON and OFF cells to stay equal, another parameter must be changed concomitantly with V_0. One possibility is to reduce the internal noise intensity of OFF cells, which will be noted D_{OFF}. This implies that the OFF cells are now more deterministic than the ON cells and will fire more reliably.

Figures 3.5a and 3.5b show power spectra for different values of the offset when $c = 1$.

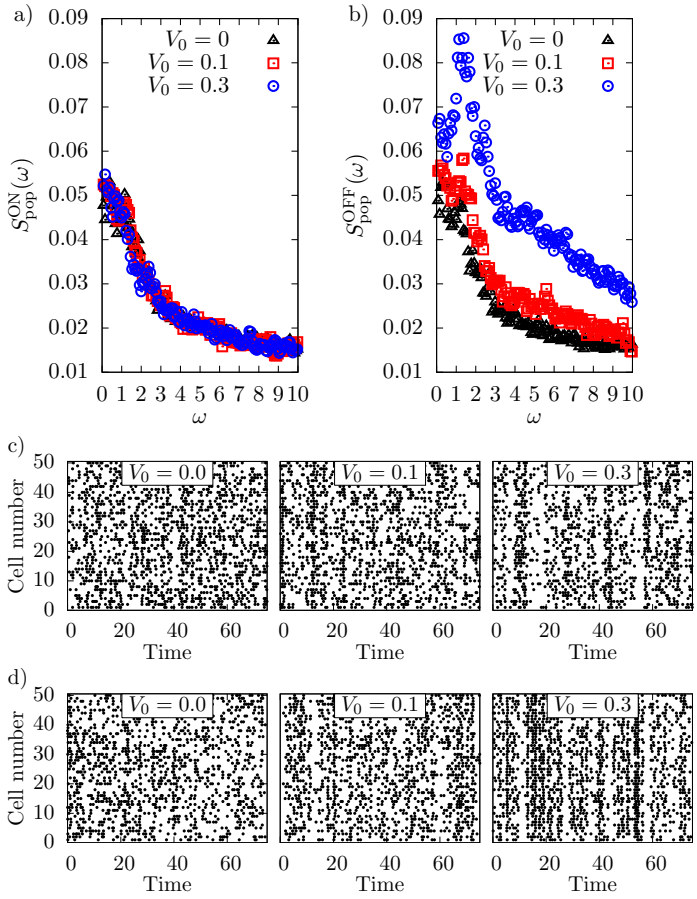

Figure 3.6: Population spectra for ON (a) and OFF cells (b) corresponding to Fig. 3.5 for the ON/OFF network. (c) and (d) are raster plots for ON and OFF cells, respectively. The spectral correlations ($S^{\mathrm{OFF}}_{\mathrm{cross}} \approx S^{\mathrm{OFF}}_{\mathrm{pop}}$) are enhanced for increasing offsets for the OFF cells (b), but not for the ON cells (a). The raster plots of panel (d) show the enhanced synchronization of OFF cells with increasing V_0.

In Fig. 3.5a, we also display the corresponding power spectrum for the ON-only network, showing an obvious gamma peak. The ON cells' power spectra are nearly identical for all values of V_0. The inset of Fig. 3.5a shows that there exists a minute difference between the LRT curves around the vanished gamma peak. On the contrary, for OFF cells, a salient contrast appears between the spectra when the offset increases. Going from $V_0 = 0$ to $V_0 = 0.3$, the OFF cells progressively transfer power from low to high frequency. This is in line with what should happen to a LIF neuron becoming more deterministic. From renewal

theory, we know that the coefficient of variation (CV) is given by [126]:

$$\mathrm{CV} = \sqrt{\frac{\lim_{\omega\to 0} S_{\mathrm{OFF}}(\omega)}{\lim_{\omega\to\infty} S_{\mathrm{OFF}}(\omega)}}. \tag{3.29}$$

From inspection of Fig. 3.5b, we see that the CV decreases for increasing offsets, meaning that greater regularity is achieved by the OFF cells. But this analysis is applicable only if the perturbed spike train does not depart too much from a renewal process. The unperturbed spike train is clearly a renewal process [126], but the perturbed one should not be, because of the feedback. From Eq. (3.22) with $c = 1$, $S^{\mathrm{OFF}}_{\mathrm{cross}} = S_{\mathrm{OFF}} - S^{(0)}_{\mathrm{OFF}}$ and, from Eq. (3.8), $S^{\mathrm{OFF}}_{\mathrm{pop}} \approx S^{\mathrm{OFF}}_{\mathrm{cross}}$. Hence, in the large N limit we have

$$S_{\mathrm{OFF}} \approx S^{(0)}_{\mathrm{OFF}} + S^{\mathrm{OFF}}_{\mathrm{pop}}. \tag{3.30}$$

Therefore, the extent to which the perturbed spike train departs from a renewal process — with the proviso that Eq. (3.22) comes from the LRT — is dictated by the values of $S^{\mathrm{OFF}}_{\mathrm{pop}}$ (see Fig. 3.6b). The maximum value of the ratio $S^{\mathrm{OFF}}_{\mathrm{pop}}/S_{\mathrm{OFF}}$ is about 0.35 for $V_0 = 0.3$. We are thus tempted to analyze the results of Fig. 3.5b in the following way: upon an increase of the offset, the OFF cells become intrinsically more regular due to the larger effective bias, and this allows for the feedback-induced oscillations to gain in power, in comparison to the power of the unperturbed spike train.

The recovery of gamma oscillations for the OFF cells is really a network effect induced by the common action of the feedback and the correlated input. This is clear from Figs. 3.6b and 3.6d. The correlations between gamma-frequency Fourier components increase with V_0, with a maximum around $\omega = 1.5$ (Fig. 3.6b). The cells also tend to synchronize, as revealed by the appearance of well-defined stripes in Fig. 3.6d for $V_0 = 0.3$.

The ON cells' behavior is more enigmatic. Both the single-neuron and the population spectra do not seem to change, as per the numerical results of Figs. 3.5a and 3.6a. In the inset to Fig. 3.5a, however, the LRT curves show a small decrease of power with increasing V_0 around the gamma frequency, to the advantage of the lower frequencies. Also, in Fig. 3.6c, one may notice a hint of attempted synchronization for $V_0 = 0.3$. Comparing Figs 3.6c and 3.6d for $V_0 = 0.3$, we see that whenever the OFF cells are firing, the ON cells are not. When the OFF cells have just received the sufficient push to cross the firing threshold, ON cells receive an opposite external drive which will cause their firing to be less likely. Given that the OFF cells strongly oscillate collectively, the ON cells are bound to their activity. However, the latter do not show gamma oscillations per se because they lack a gamma peak in the power spectrum. This is different than the experimental situation where both the ON and

OFF cells show gamma oscillations [75]. The slight increase in oscillatory strength for the ON/OFF network is compatible with the neural field prediction that an increase in V_0 can enhance the propensity of the deterministic ON/OFF system to exhibit oscillations ([121], Fig. 4c), although in that case the rates are not constrained to be equal.

3.5.3 Asymmetry between ON and OFF cells: Unequal membrane time constants

Studies of mutual information and coherence between spike trains and Gaussian noise have shown that OFF cells are low-pass under local stimulus for all maps constituting the ELL [80, 114, 218]. This means that OFF cells tend to respond more to the low-frequency components of the stimulus. The ON cells are more heteregenous across the maps and can be low-pass, broadband and even high-pass [218]. Thus, on average, the OFF cells are more low-pass than the ON cells. One way to account for this is to reconsider the assumption that we have made hitherto, i.e. that the membrane time constants of ON and OFF neurons are equal. Slower membrane dynamics are likely to promote the low-frequency components of the input. Mathematically, this implies that we consider the following dynamics for the unperturbed OFF membrane potentials:

$$\tau_{\text{ratio}}\dot{v}_i^{\text{OFF}}(t) = -v_i^{\text{OFF}}(t) + \mu_{\text{OFF}} + \xi_i^{\text{OFF}}(t), \tag{3.31}$$

with $\tau_{\text{ratio}} \equiv \tau_{\text{OFF}}/\tau_{\text{ON}}$ the ratio of the membrane time constants. The internal noises are equal as before. Analytical formulas for the rates and spectra of the unperturbed system are given for LIF equations of the form $\dot{v} = -v + \mu + \xi(t)$. The above equation is evidently not of this form. To get rid of the τ_{ratio} prefactor, we define $\hat{t} = t/\tau_{\text{ratio}}$. The consequence is to rescale

1. the noise intensity: $D \to \hat{D} = D/\tau_{\text{ratio}}$;
2. the various time constants of the model:

$$\hat{\tau}_R = \frac{\tau_R}{\tau_{\text{ratio}}} \qquad \hat{\tau}_D = \frac{\tau_D}{\tau_{\text{ratio}}}, \qquad \hat{\tau}_S = \frac{\tau_S}{\tau_{\text{ratio}}};$$

3. and the feedback strength: $\hat{G} = G/\tau_{\text{ratio}}$.

Evidently, the use of these new constants demands corresponding modifications of the spectral quantities. The steps leading to the LRT in presence of unequal membrane time constants are given in Appendix 3.8. We get for the unperturbed power spectrum (see Eq.

(3.42))

$$S^{(0)}_{\mathrm{OFF}} = r_{\mathrm{OFF}} \frac{\left|\mathcal{D}_{i\tau_{\mathrm{ratio}}\omega}\left(\frac{\mu_{\mathrm{OFF}}-v_T}{\sqrt{D/\tau_{\mathrm{ratio}}}}\right)\right|^2 - \mathrm{e}^{2\tau_{\mathrm{ratio}}\delta}\left|\mathcal{D}_{i\tau_{\mathrm{ratio}}\omega}\left(\frac{\mu_{\mathrm{OFF}}-v_R}{\sqrt{D/\tau_{\mathrm{ratio}}}}\right)\right|^2}{\left|\mathcal{D}_{i\tau_{\mathrm{ratio}}\omega}\left(\frac{\mu_{\mathrm{OFF}}-v_T}{\sqrt{D/\tau_{\mathrm{ratio}}}}\right) - \mathrm{e}^{\tau_{\mathrm{ratio}}\delta}\mathrm{e}^{i\omega\tau_R}\mathcal{D}_{i\tau_{\mathrm{ratio}}\omega}\left(\frac{\mu_{\mathrm{OFF}}-v_R}{\sqrt{D/\tau_{\mathrm{ratio}}}}\right)\right|^2}$$

and for the susceptibility (cf. Eq. (3.43))

$$A_{\mathrm{OFF}} = \left(\frac{i\tau_{\mathrm{ratio}}\omega r_{\mathrm{OFF}}}{\sqrt{D/\tau_{\mathrm{ratio}}}(i\tau_{\mathrm{ratio}}\omega - 1)}\right)\left[\frac{\mathcal{D}_{i\tau_{\mathrm{ratio}}\omega-1}\left(\frac{\mu_{\mathrm{OFF}}-v_T}{\sqrt{D/\tau_{\mathrm{ratio}}}}\right) - \mathrm{e}^{\tau_{\mathrm{ratio}}\delta}\mathcal{D}_{i\tau_{\mathrm{ratio}}\omega-1}\left(\frac{\mu_{\mathrm{OFF}}-v_R}{\sqrt{D/\tau_{\mathrm{ratio}}}}\right)}{\mathcal{D}_{i\tau_{\mathrm{ratio}}\omega}\left(\frac{\mu_{\mathrm{OFF}}-v_T}{\sqrt{D/\tau_{\mathrm{ratio}}}}\right) - \mathrm{e}^{\tau_{\mathrm{ratio}}\delta}\mathrm{e}^{i\omega\tau_R}\mathcal{D}_{i\tau_{\mathrm{ratio}}\omega}\left(\frac{\mu_{\mathrm{OFF}}-v_R}{\sqrt{D/\tau_{\mathrm{ratio}}}}\right)}\right]. \tag{3.32}$$

In Sec. 3.8.2 of Appendix 3.8, we prove that r_{OFF} decreases when τ_{ratio} is increased at fixed V_0. Therefore, to fulfill the constraint of equal ON and OFF rates, V_0 must be increased concurrently with τ_{ratio}. For $\tau_{\mathrm{ratio}} = 1$ and $V_0 = 0$, $r_{\mathrm{ON}} = r_{\mathrm{OFF}}$. Now, we have to suppose that as we increase τ_{ratio}, we also increase V_0 so that the rates stay the same (numerically, they will remain the same within statistical error). This means that $\mu_{\mathrm{OFF}} = \mu + V_0 + Gr_{\mathrm{ON}}$ increases with τ_{ratio} whereas μ_{ON} remains the same. Also, $\hat{D}$ diminishes with increasing τ_{ratio}. Hence, the unperturbed activity of the OFF cells tends more and more towards the deterministic firing regime as τ_{ratio} is increased from 1, with the possibility of bifurcating to the deterministic regime at $\mu_{\mathrm{OFF}} = 1$. Finally, given the presence of $\tau_{\mathrm{ratio}}\omega$ instead of ω as a subscript of the parabolic cylinder functions, the unperturbed spectrum will also have the tendency to be compressed toward the $\omega = 0$ axis for large τ_{ratio}. But this compression is not trivial due to the presence of the exponential term $\exp(i\omega\tau_R)$ in the expression above. Something similar is expected for the susceptibility.

We now want to study the effect of τ_{ratio} on the coherence of OFF cells. In the introduction 3.1, we stated that local and spatially uncorrelated global stimuli seem to have similar impacts on ELL neurons. Hence, we here use a global stimulus with $c = 0$. Exceptionally, the external Gaussian noise is bandlimited to 0-120 Hz. The coherence is given by [37]

$$C(\omega) = \frac{|S_{\eta_i y_i}(\omega)|^2}{S_{\mathrm{OFF}}(\omega)S_{\mathrm{st}}(\omega)} \tag{3.33}$$

where $S_{\eta_i y_i}(\omega)$ is the cross spectrum between the spatially uncorrelated noise and the spike train given by Eq. (3.20). Using the LRT to express the various quantities in Eq. (3.33), we get in the limit $N \to \infty$:

$$C(\omega) = \frac{|A_{\mathrm{OFF}}(\omega)|^2 S_{st}(\omega)}{S^{(0)}_{\mathrm{OFF}}(\omega) + |A_{\mathrm{OFF}}(\omega)|^2 S_{st}(\omega)}. \tag{3.34}$$

Figure 3.7 shows the effect of τ_{ratio} on the coherence as a function of ω for two special cases: In

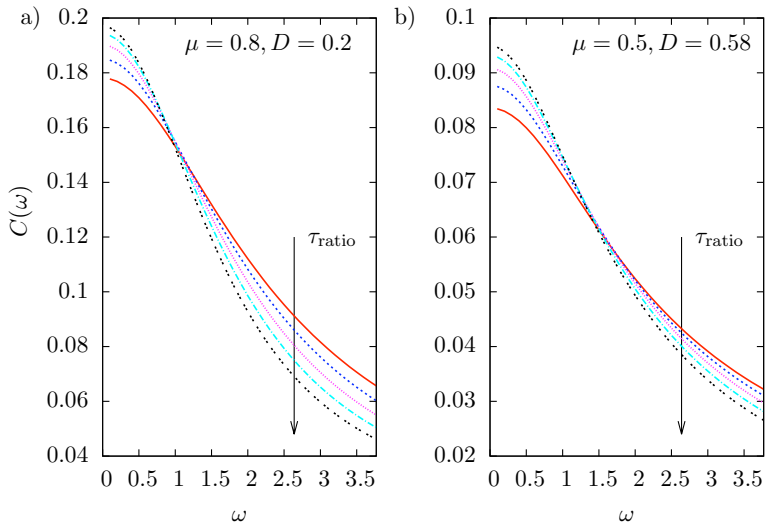

Figure 3.7: Coherence of OFF cells—computed using Eq. (3.34)—for different τ_{ratio} with spatially uncorrelated input ($c = 0$) and in the limit $N \to \infty$. V_0 has been adjusted so that $r_{\text{ON}} = r_{\text{OFF}}$ within an absolute error of order 10^{-3}. Arrow indicates the direction of increasing τ_{ratio}; τ_{ratio} starts at 1 (solid red curves) and stops at 1.8 (double-dotted black curves) with increments of 0.2. Values for μ and D are: a) $\mu = 0.8, D = 0.2$; b) $\mu = 0.5, D = 0.58$. For a), the effective biases range from 0.43 to 0.95, and for b) from 0.08 to 0.83 for increasing τ_{ratio}. Other parameter values are $\tau_D = 1$, $\tau_S = 0.5$, $\tau_R = 0.1$ and $G = -1.2$. Frequency range used for the figure corresponds to 0-120 Hz for $\tau_{\text{ON}} = 5$ ms. Note the different scales used.

Fig. 3.7a, the bias μ is high and D is low and conversely for Fig. 3.7b. In both cases, it is obvious that the OFF cells become more low-pass as τ_{ratio} is increased. For $\tau_{\text{ratio}} = 1$, $A_{\text{ON}} = A_{\text{OFF}}$ and $S^{(0)}_{\text{ON}} = S^{(0)}_{\text{OFF}}$, therefore the red curves of Fig. 3.7 also represent the coherence of the ON cells.

With the ratio different from 1, A_{ON} and A_{OFF} (given by Eq. (3.32)) become different and the feedback may again induce oscillations, provided that the stimulus is spatially correlated [see Eq. (3.20)]. Figure 3.8 shows how the power and population spectra are modified following an increase of τ_{ratio} (the external stimulus is again the Gaussian white noise of unlimited bandwidth). The first thing one may notice is that the ON cells' power spectrum and population spectrum are only slightly modified. The bump at $\omega \approx 1.5$ in the power spectrum (Fig. 3.8a) grows slightly with τ_{ratio}, but reaches nowhere near the one of the ON-only network in Fig. 3.2a. In contrast, there is an appreciable effect on the OFF cells' spectral quantities. For $\tau_{\text{ratio}} = 1$, we have the symmetric case showing a small spectral peak at $\omega = 1.5$ (Fig. 3.8c, red curve); with $\tau_{\text{ratio}} = 2$, $\mu_{\text{OFF}} = 1.001$, and the system has just stepped into the deterministic regime, causing the peak around $\omega = 3$. The firing is then

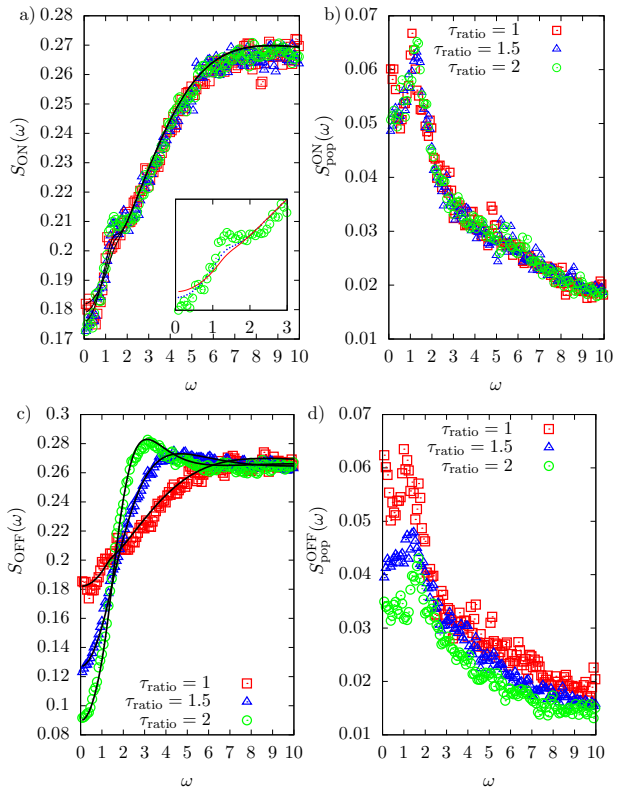

Figure 3.8: Single-neuron spectra and population spectra for ON (top) and OFF cells (bottom) for $\tau_{\text{ratio}} = 1, 1.5, 2$ and $c = 1$. Other parameters are as in Fig. 3.2. V_0 has been adjusted so that $r_{\text{ON}} \approx r_{\text{OFF}}$. The effective OFF biases are $\mu_{\text{OFF}} = 0.48, 0.79, 1.001$ for, respectively, $\tau_{\text{ratio}} = 1, 1.5, 2$; the corresponding values for V_0 are $0, 0.305, 0.52$. On panel (a), the LRT curve with the highest bump corresponds to $\tau_{\text{ratio}} = 2$ and the one with the smallest peak to $\tau_{\text{ratio}} = 1$. Inset to panel (a): zoom on the peaks of the LRT curves (solid red: $\tau_{\text{ratio}} = 1$; dotted blue: $\tau_{\text{ratio}} = 1.5$; dashed green: $\tau_{\text{ratio}} = 2$). We also display the numerical results for $\tau_{\text{ratio}} = 2$ for comparison.

dominated by the unperturbed activity of single OFF neurons, and the feedback-induced bump is no longer visible. Remarkably, for $\tau_{\text{ratio}} = 1.5$, for which $\mu_{\text{OFF}} < 1$, we do not see any recovery of the gamma peak. Obviously, this means that the difference between the ON and OFF susceptibilities is not large enough to permit gamma oscillations.

Figure 3.8d shows that spectral correlations between OFF neurons are impeded when τ_{ratio} is increased. On the other one hand, the spectral correlations between ON neurons stay about the same for all values of the ratio (Fig. 3.8b). This is even more obvious from the raster plots of Fig. 3.9. Increasing τ_{ratio} uncorrelates the firing of OFF cells and we progressively lose the striped pattern of $\tau_{\text{ratio}} = 1$ (compare Fig. 3.9d and 3.9f). The OFF cells still oscillate, but

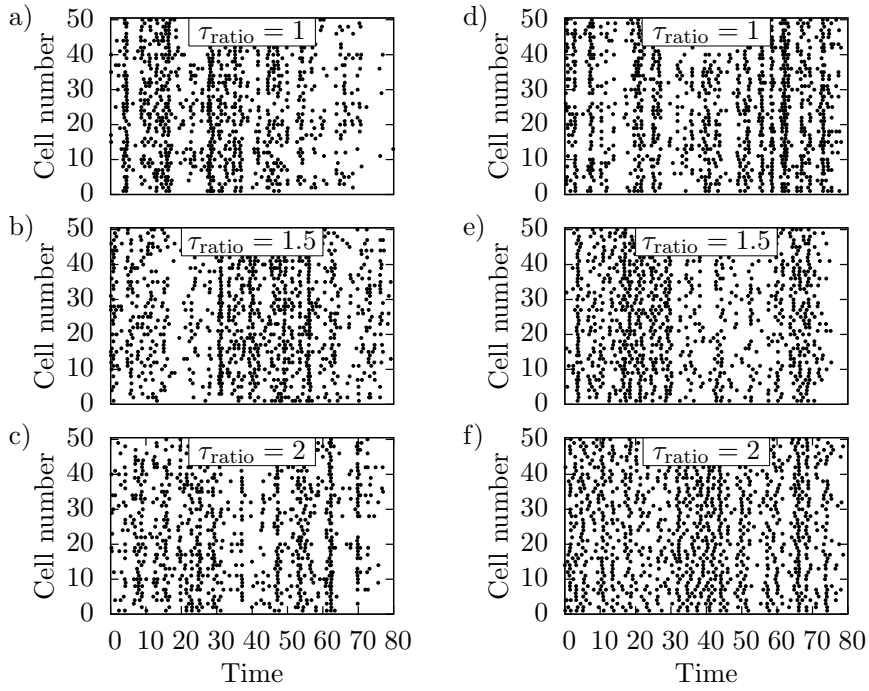

Figure 3.9: Raster plots of the activity of ON (a, b, c) and OFF (d, e, f) cells corresponding to Fig. 3.8. Notice on the right column (plots d to f) how the OFF cells' spikings desynchronize with increasing τ_{ratio}.

in a more independent manner; the gamma oscillations are somewhat buried underneath the dominant unperturbed activity. From Sec. 3.5.2, we know that OFF neurons with higher bias and lower internal noise tend to synchronize more in presence of spatially correlated noise. In the present case, however, their more deterministic behavior is triggered in part by the slower dynamics for $\tau_{\text{ratio}} > 1$ (the other part being due to $V_0 > 0$). This not only affects the intensity of both the internal and external noises — D and D_E become D/τ_{ratio} and D_E/τ_{ratio}, respectively—but also the feedback strength perceived by OFF neurons. With a lower strength G/τ_{ratio}, the feedback is less able to modulate the collective activity of OFF cells. Our results thus indicate that the loss of feedback here dominates over the increased synchrony.

3.6 Summary and Conclusion

In this chapter, we have derived a linear response theory (LRT) for networks containing two classes of neurons, and for which the interactions are predominantly through an all-to-all feedback loop. We obtained expressions for the power spectrum of these cells when submitted to external noise. The primary motivation for this work was the processing done by the electrosensory lateral line lobe (ELL) of weakly electric fish, in particular its ability to generate gamma oscillations under spatially correlated stimuli. The ELL contains ON and OFF cells which respond in opposite ways to inputs. We have studied the impact of the OFF cells on the network's oscillatory behavior under stochastic inputs. The model assumed that both the ON and OFF cells send their spike trains to a remote kernel which sums all the incoming spike trains and feeds them back with a delay to all the cells of the network. When only ON cells are taken into account (ON-only network), this diffuse inhibitory feedback is sufficient to trigger network oscillations when the global input is spatially correlated [125]. We have shown analytically and numerically that including the OFF cells (ON/OFF network) hinders the onset of oscillations for a spatially correlated stimulus.

Neural field theory already showed that gamma oscillations are less prevalent in an ON/OFF network compared to ON-only networks. In that theory, a certain parameter can be computed, R ([120], Eq. (5)), which acts as a global gain parameter that depends on the spatial extent of a spatially homogeneous input. When R exceeds a critical value, oscillations occur. However, this is applicable in the context of deterministic inputs. In the LRT presented here, it is possible to deal with stochastic inputs, and both the feedback strength G and the correlation parameter c explicitly determine the occurrence of an oscillatory state (cf. the expression for the power spectrum Eq. (3.20)).

Having concluded that bringing the OFF cells into the picture has a negative effect on network oscillations, we then tried to break the symmetry of the ON and OFF cells to see if we can recover the gamma oscillations obtained with an ON-only network [77, 125]. In doing so, we had to ensure that the rates of ON and OFF cells are approximately equal to satisfy physiological constraints. The symmetry breakings were also done on the grounds of physiological considerations. First, we increased the bias of OFF cells (through an increase of the offset V_0). This was motivated by their lower spike threshold compared to ON cells. To keep the rates equal, we also decreased the internal noise intensity of OFF cells. The OFF cells were then more deterministic and their spectral cross-correlations (as seen on the population spectra) increased with V_0. Such a regularity enhanced the feedback-induced gamma oscillations of the OFF cells (see Figs. 3.5b and 3.6b). However, the recovered gamma oscillations were far from reaching the same level as that of either the corresponding ON-only

network [125] or the experimental results of [77].

Second, we increased the OFF membrane time constant with respect to the ON's. The goal was to mimic the low-pass properties of OFF cells. If the OFF cells are more low-pass compared to the ON cells, then it may make room for the latter to oscillate. Again, the constraint of equal firing rates had to be satisfied, and it was achieved by increasing the bias of OFF cells. The cross spectra of ON cells were not modified (Fig. 3.8b) but, contrary to the first asymmetric case, the spectral cross-correlations of OFF cells were impeded (Fig. 3.8d). This was interpreted as a reduction of the synchronization of the OFF cells (and confirmed by the corresponding raster plots, Figs. 3.9d-f). The power spectra of ON cells showed a slight increase of the height of the feedback-induced peak (Fig. 3.8a); for the OFF cells, this peak disappears when τ_{ratio} departs from 1 (Fig. 3.8c). Another peak appears for the OFF cells at higher frequency, which is a direct consequence of the more regular spiking due to the higher effective bias.

Interestingly, both symmetry breakings yielded more regularity in the OFF neurons' firing, but with drastically different results on the collective behavior of the network. On the one hand, the higher bias and the lower intrinsic noise allowed for less random OFF cells' spikings, which consequently helped to synchronize the OFF cells (and to a lesser extent, the ON cells). On the other hand, a slower dynamics for OFF cells, together with a higher bias, strives to reduce the synchronization. Since these symmetry breakings have opposite effects on the feedback-induced gamma oscillations, it seems unlikely that combining them in any way may allow for oscillations as strong as in the ON-only network [125] or the experimental results [77].

Note that, as in [125], we completely ignored the activity of the neurons responsible for the diffuse inhibitory feedback. According to [139], however, this should not invalidate the analysis found in both [125] and in the present paper. Indeed, they show that the dynamics of LIF neurons constituting the feedback kernel have only a marginal effect on the response of the network.

Another point is that both the ON and OFF cells show gamma oscillations experimentally [75]. Hence, one may wonder why we tried to invoke these oscillations by actually privileging one class of cells with respect to the other. The first reason is that the frequency range for gamma oscillations is large (usually $30 - 90$ Hz). So, it is not excluded that ON and OFF cells show gamma oscillations of different frequencies. Also, as stated in the previous paragraph, even though a combination of the two attempts may not allow to recover the full strength of gamma oscillations, it may at least bring back some symmetry between the power spectra of ON and OFF cells. Finally, we also think that the mathematical apparatus derived here may

be used in studies of other feedback-driven systems like the thalamocortical circuit (see for instance [196]) or of correlations in recurrent networks [164].

In the end, the simplest hypothesis that preserves the feedback-induced gamma oscillations is that the ON and OFF cells are segregated, i.e. ON cells would be connected (through feedback) only to ON cells and the same for OFF cells. This would be the easiest way to recover the gamma oscillations, but we lack the necessary physiological evidence for such connectivity. Alternatively, future work may require more detailed anatomical and physiological analysis of the ELL feedback circuitry to elucidate the origin of gamma oscillations.

Acknowledgements

The authors acknowledge support from the Fonds de Recherche en Sciences Naturelles et Technologies du Québec, the Natural Sciences and Engineering Research Council of Canada, and the Canadian Institutes of Health Research. They also wish to thank Jason Middleton and Maurice J. Chacron for discussions.

3.7 Appendix A: Proof that the ON (resp. OFF) rate is a monotonically decreasing (resp. increasing) function of V_0

Using Eq. (3.9), we write

$$\begin{aligned}\frac{dr_{\mathrm{ON}}}{dV_0} &= \frac{dr_{\mathrm{ON}}}{d\mu_{\mathrm{ON}}}\frac{d\mu_{\mathrm{ON}}}{dV_0} = \frac{dr_{\mathrm{ON}}}{d\mu_{\mathrm{ON}}}\frac{G}{2}\left(\frac{dr_{\mathrm{ON}}}{dV_0}+\frac{dr_{\mathrm{OFF}}}{dV_0}\right)\\ \frac{dr_{\mathrm{OFF}}}{dV_0} &= \frac{dr_{\mathrm{OFF}}}{d\mu_{\mathrm{OFF}}}\left[1+\frac{G}{2}\left(\frac{dr_{\mathrm{ON}}}{dV_0}+\frac{dr_{\mathrm{OFF}}}{dV_0}\right)\right].\end{aligned} \tag{3.35}$$

Adding the two equations above, isolating $\frac{dr_{\mathrm{ON}}}{dV_0}+\frac{dr_{\mathrm{OFF}}}{dV_0}$ and replacing in Eq. (3.35) yields

$$\begin{aligned}\frac{dr_{\mathrm{ON}}}{dV_0} &= \frac{dr_{\mathrm{ON}}}{d\mu_{\mathrm{ON}}}\frac{G}{2}\left[\frac{\frac{dr_{\mathrm{OFF}}}{d\mu_{\mathrm{OFF}}}}{1-\frac{G}{2}\left(\frac{dr_{\mathrm{ON}}}{d\mu_{\mathrm{ON}}}+\frac{dr_{\mathrm{OFF}}}{d\mu_{\mathrm{OFF}}}\right)}\right]\\ \frac{dr_{\mathrm{OFF}}}{dV_0} &= \frac{dr_{\mathrm{OFF}}}{d\mu_{\mathrm{OFF}}}\left[1+\frac{G}{2}\frac{\frac{dr_{\mathrm{OFF}}}{d\mu_{\mathrm{OFF}}}}{1-\frac{G}{2}\left(\frac{dr_{\mathrm{ON}}}{d\mu_{\mathrm{ON}}}+\frac{dr_{\mathrm{OFF}}}{d\mu_{\mathrm{OFF}}}\right)}\right].\end{aligned}$$

As proven in subsection 3.8.2 of Appendix 3.8, $\frac{dr_{\mathrm{ON}}}{d\mu_{\mathrm{ON}}}$ and $\frac{dr_{\mathrm{OFF}}}{d\mu_{\mathrm{OFF}}}$ are always positive. Since $G < 0$, it is easy to see that $\frac{dr_{\mathrm{ON}}}{dV_0} < 0$. For $\frac{dr_{\mathrm{OFF}}}{dV_0}$, the term in brackets will be positive for

$$2 + |G| \left(\frac{dr_{\mathrm{ON}}}{d\mu_{\mathrm{ON}}} + \frac{dr_{\mathrm{OFF}}}{d\mu_{\mathrm{OFF}}} \right) > |G| \frac{dr_{\mathrm{OFF}}}{d\mu_{\mathrm{OFF}}},$$

which is true for all V_0, hence $\frac{dr_{\mathrm{OFF}}}{dV_0} > 0$.

3.8 Appendix B: Effect of unequal membrane time constants on the LRT

In this appendix, we detail how the dynamics of OFF cells is affected by an increase of their membrane time constant with respect to that of the ON cells. We also delineate its consequences on the analytical results for the rate, the unperturbed spectrum, the susceptibility and the LRT.

3.8.1 Transformation of the OFF cells' dynamics

The unperturbed dynamics of the subthreshold OFF voltage is

$$\tau_{\mathrm{ratio}} \ddot{v}_i^{\mathrm{OFF}}(t) = -v_i^{\mathrm{OFF}}(t) + \mu_{\mathrm{OFF}} + \xi_i^{\mathrm{OFF}}(t).$$

Defining $\hat{t} = t/\tau_{\mathrm{ratio}}$ leads to

$$\frac{d\hat{v}_i^{\mathrm{OFF}}}{d\hat{t}}(\hat{t}) = -\hat{v}_i^{\mathrm{OFF}}(\hat{t}) + \mu_{\mathrm{OFF}} + \hat{\xi}_i^{\mathrm{OFF}}(\hat{t}), \tag{3.36}$$

where we have defined

$$\hat{\xi}_i^{\mathrm{OFF}}(\hat{t}) \equiv \xi_i^{\mathrm{OFF}}(\tau_{\mathrm{ratio}}\hat{t}) \text{ and } \hat{v}_i^{\mathrm{OFF}}(\hat{t}) \equiv v_i^{\mathrm{OFF}}(\tau_{\mathrm{ratio}}\hat{t}).$$

The processes $\hat{\xi}_i^{\mathrm{OFF}}(\hat{t})$ obey

$$\langle \hat{\xi}_i^{\mathrm{OFF}}(\hat{t}) \hat{\xi}_i^{\mathrm{OFF}}(\hat{t}') \rangle = 2\hat{D}\delta(\hat{t} - \hat{t}'),$$

with the rescaled noise intensity $\hat{D} = D/\tau_{\mathrm{ratio}}$. Note the important fact that $\mu_{\mathrm{OFF}} = \mu + V_0 + G(r_{\mathrm{ON}} + r_{\mathrm{OFF}})/2$ is invariant with respect to the transformation $t \to t/\tau_{\mathrm{ratio}}$. This is readily

seen by considering the transformation of the feedback function $f(t)$. Putting $t = \tau_{\text{ratio}}\hat{t}$ in Eq. (3.3) gives

$$\hat{f}(\hat{t}) \equiv f(\tau_{\text{ratio}}\hat{t}) = \frac{G}{2N}\int_{\tau_D}^{\infty}\frac{d\tau}{\tau_S}\left(\frac{\tau-\tau_D}{\tau_S}\right)\exp\left(-\frac{\tau-\tau_D}{\tau_S}\right)\sum_{i=1}^{N}\left[y_i^{\text{ON}}(\tau_{\text{ratio}}\hat{t}-\tau)+y_i^{\text{OFF}}(\tau_{\text{ratio}}\hat{t}-\tau)\right].$$

From the definition of the spike trains, Eq. (3.4), we have

$$\begin{aligned} y(\tau_{\text{ratio}}\hat{t}) &= \sum_{t_i}\delta(\tau_{\text{ratio}}\hat{t}-t_i) \\ &= \frac{1}{\tau_{\text{ratio}}}\sum_{\hat{t}_i}\delta\left(\hat{t}-\hat{t}_i\right) \equiv \frac{1}{\tau_{\text{ratio}}}\hat{y}(\hat{t}). \end{aligned} \tag{3.37}$$

Here, the $\hat{t}_i$'s are the spike times in units of the OFF membrane time constant and $\hat{y}(\hat{t}) \equiv \sum_{\hat{t}_i}\delta\left(\hat{t}-\hat{t}_i\right)$ is the corresponding spike train. We replace this in the expression for the feedback and change variable to $\tau \to \tau/\tau_{\text{ratio}}$ to give

$$\hat{f}(\hat{t}) = \frac{G}{2N}\int_{\tau_D/\tau_{\text{ratio}}}^{\infty}\frac{d\tau}{\tau_S}\left(\frac{\tau\tau_{\text{ratio}}-\tau_D}{\tau_S}\right)\exp\left(-\frac{\tau\tau_{\text{ratio}}-\tau_D}{\tau_S}\right)\sum_{i=1}^{N}\left[\hat{y}_i^{\text{ON}}(\hat{t}-\tau)+\hat{y}_i^{\text{OFF}}(\hat{t}-\tau)\right].$$

We define $\tau_D = \hat{\tau}_D\tau_{\text{ratio}}$, $\tau_S = \hat{\tau}_S\tau_{\text{ratio}}$ and $G = \hat{G}\tau_{\text{ratio}}$. Doing so, the feedback keeps the same form as before:

$$\hat{f}(\hat{t}) = \frac{\hat{G}}{2N}\int_{\hat{\tau}_D}^{\infty}\frac{d\tau}{\hat{\tau}_S}\left(\frac{\tau-\hat{\tau}_D}{\hat{\tau}_S}\right)\exp\left(-\frac{\tau-\hat{\tau}_D}{\hat{\tau}_S}\right)\sum_{i=1}^{N}\left[\hat{y}_i^{\text{ON}}(\hat{t}-\tau)+\hat{y}_i^{\text{OFF}}(\hat{t}-\tau)\right].$$

The static part of the feedback is given by

$$\langle\hat{f}\rangle = \frac{\hat{G}}{2}(\hat{r}_{\text{ON}}+\hat{r}_{\text{OFF}}) = \frac{G}{2}(r_{\text{ON}}+r_{\text{OFF}}),$$

where we have transformed the rates measured in units of the ON membrane time constant (u.ON.m.t.c.) to the ones measured in unit of the OFF time constant (u.OFF.m.t.c.), i.e. $\hat{r} = \tau_{\text{ratio}}r$. Therefore, it appears that the static part of the feedback is invariant with respect to the transformation $t \to t/\tau_{\text{ratio}}$, and this is the case for the effective bias μ_{OFF} as well.

3.8.2 Transformation of the analytical results

When $\tau_{\text{ratio}} > 1$, the formula for the rate of a LIF neuron, Eq. (3.10), must be replaced by

$$g(\mu, \tau_{\text{ratio}}) = \left(\tau_R + \tau_{\text{ratio}} \sqrt{\pi} \int_{(\mu-1)/\sqrt{2D/\tau_{\text{ratio}}}}^{\mu/\sqrt{2D/\tau_{\text{ratio}}}} \exp(z^2) \text{erfc}(z) dz \right)^{-1}. \tag{3.38}$$

This equation is obtained by replacing $D \to \hat{D}$ and $\tau_R \to \hat{\tau}_R = \tau_R/\tau_{\text{ratio}}$ in Eq. (3.10), and by using the conversion between rates measured in different units. Note that for Gaussian white noise stimuli of unlimited bandwidth, we would use $Q = D + D_E$ as the noise intensity. We now want to prove that r_{OFF} decreases with τ_{ratio}, i.e. that $\frac{dr_{\text{OFF}}}{d\tau_{\text{ratio}}} < 0$. In what follows we keep V_0 fixed. The rates are obtained by solving the following coupled equations in a self-consistent way (compare with Eq. (3.11)):

$$r_{\text{ON}} = \nu[\mu_{\text{ON}}(\tau_{\text{ratio}})] \quad \text{and} \quad r_{\text{OFF}} = g[\mu_{\text{OFF}}(\tau_{\text{ratio}}), \tau_{\text{ratio}}]$$

with ν given by Eq. (3.10) and g given by Eq. (3.38) above. Note that both rates are expressed in u.ON.m.t.c. and that only $g(\mu_{\text{OFF}}, \tau_{\text{ratio}})$ explicitly depends on τ_{ratio}. The derivatives of the rates with respect to τ_{ratio} are

$$\begin{aligned}
\frac{dr_{\text{ON}}}{d\tau_{\text{ratio}}} &= \frac{dr_{\text{ON}}}{d\mu_{\text{ON}}} \frac{d\mu_{\text{ON}}}{d\tau_{\text{ratio}}} = \frac{G}{2} \frac{dr_{\text{ON}}}{d\mu_{\text{ON}}} \left(\frac{dr_{\text{ON}}}{d\tau_{\text{ratio}}} + \frac{dr_{\text{OFF}}}{d\tau_{\text{ratio}}} \right) \\
\frac{dr_{\text{OFF}}}{d\tau_{\text{ratio}}} &= \frac{G}{2} \frac{dr_{\text{OFF}}}{d\mu_{\text{OFF}}} \left(\frac{dr_{\text{ON}}}{d\tau_{\text{ratio}}} + \frac{dr_{\text{OFF}}}{d\tau_{\text{ratio}}} \right) + \frac{\partial r_{\text{OFF}}}{\partial \tau_{\text{ratio}}}.
\end{aligned}$$

Adding them together and isolating $\left(\frac{dr_{\text{ON}}}{d\tau_{\text{ratio}}} + \frac{dr_{\text{OFF}}}{d\tau_{\text{ratio}}} \right)$ yields:

$$\left(\frac{dr_{\text{ON}}}{d\tau_{\text{ratio}}} + \frac{dr_{\text{OFF}}}{d\tau_{\text{ratio}}} \right) = \frac{\frac{\partial r_{\text{OFF}}}{\partial \tau_{\text{ratio}}}}{1 - \frac{G}{2} \left(\frac{dr_{\text{ON}}}{d\mu_{\text{ON}}} + \frac{dr_{\text{OFF}}}{d\mu_{\text{OFF}}} \right)}.$$

Hence,

$$\frac{dr_{\text{OFF}}}{d\tau_{\text{ratio}}} = \frac{\partial r_{\text{OFF}}}{\partial \tau_{\text{ratio}}} \left[\frac{\frac{dr_{\text{OFF}}}{d\mu_{\text{OFF}}}}{1 - \frac{G}{2} \left(\frac{dr_{\text{ON}}}{d\mu_{\text{ON}}} + \frac{dr_{\text{OFF}}}{d\mu_{\text{OFF}}} \right)} + 1 \right]. \tag{3.39}$$

Let us define

$$\frac{\mu_{\text{OFF}}}{\sqrt{2D/\tau_{\text{ratio}}}} \equiv b \qquad \frac{\mu_{\text{OFF}} - 1}{\sqrt{2D/\tau_{\text{ratio}}}} \equiv a$$

with $b = a + \sqrt{\frac{\tau_{\text{ratio}}}{2D}} > a$. On the one hand, we have

$$\frac{\partial r_{\text{OFF}}}{\partial \tau_{\text{ratio}}} = -r_{\text{OFF}}^2 \sqrt{\pi} \left\{ \int_a^b \exp(z^2)\text{erfc}(z)dz + \frac{b}{2}\exp(b^2)\text{erfc}(b) - \frac{a}{2}\exp(a)^2\text{erfc}(a) \right\},$$

which is smaller than zero because the function $h(x) = x\exp(x^2)\text{erfc}(x)$ is strictly increasing and the first term between the curly brackets is positive. On the other hand,

$$\frac{dr_{\text{ON}}}{d\mu_{\text{ON}}} = \sqrt{\frac{\pi}{2D}} r_{\text{ON}}^2 \Big[\exp\left(\frac{a^2}{\tau_{\text{ratio}}}\right) \text{erfc}\left(\frac{a}{\sqrt{\tau_{\text{ratio}}}}\right) - \exp\left(\frac{b^2}{\tau_{\text{ratio}}}\right) \text{erfc}\left(\frac{b}{\sqrt{\tau_{\text{ratio}}}}\right) \Big]$$

and

$$\frac{dr_{\text{OFF}}}{d\mu_{\text{OFF}}} = \sqrt{\frac{\tau_{\text{ratio}}^3 \pi}{2D}} r_{\text{OFF}}^2 \left[\exp(a^2)\text{erfc}(a) - \exp(b^2)\text{erfc}(b)\right]$$

are always positive since $\exp(x^2)\text{erfc}(x)$ is monotonically decreasing. Therefore, the term in square brackets in Eq. (3.39) is always positive and r_{OFF} decreases with τ_{ratio}.

The unperturbed power spectrum for OFF neurons in u.OFF.m.t.c., $\hat{S}^{(0)}(\hat{\omega}, \mu_{\text{OFF}}, \hat{D}) \equiv \hat{S}_{\text{OFF}}^{(0)}(\hat{\omega})$, is obtained from Eq. (3.12) — without the term proportional to $\delta(\omega)$ — by replacing the parameters and variables by their hatted counterparts. Note in this respect that $\hat{\delta}$ and $\hat{\omega}$ are linked to the corresponding quantities in u.ON.m.t.c. by

$$\hat{\delta} = \tau_{\text{ratio}}\delta \quad \text{and} \quad \hat{\omega} = \tau_{\text{ratio}}\omega. \tag{3.40}$$

It is useful for graphing and computation purposes to have the OFF power spectrum in u.ON.m.t.c.. One can show that this is obtained from the spectrum in u.OFF.m.t.c. by replacing $\hat{\omega} = \tau_{\text{ratio}}\omega$ and dividing by τ_{ratio}:

$$S_{\text{OFF}}(\omega) = \frac{1}{\tau_{\text{ratio}}} \hat{S}_{\text{OFF}}(\tau_{\text{ratio}}\omega).$$

The intuitive rationale behind the prefactor $1/\tau_{\text{ratio}}$ is that the power spectrum is measured in Hz when the second is the unit of time. Therefore, changing the way we count time changes the scaling factor of the power spectrum too. Using the same intuitive argumentation, we have for the OFF susceptibility

$$A_{\text{OFF}}(\omega) = \frac{1}{\tau_{\text{ratio}}} \hat{A}_{\text{OFF}}(\tau_{\text{ratio}}\omega). \tag{3.41}$$

This follows at once from the facts that D_E has dimension [time] when the voltage is dimen-

sionless and time is measured in seconds, and that $2D_E|A|^2$, the first order response to noise, must have dimension $[\text{time}]^{-1}$.

Gathering all these results leads to the unperturbed power spectrum of the OFF cells in u.ON.m.t.c.:

$$S^{(0)}_{\text{OFF}} = r_{\text{OFF}} \frac{\left|\mathcal{D}_{i\tau_{\text{ratio}}\omega}\left(\frac{\mu_{\text{OFF}}-v_T}{\sqrt{\hat{D}}}\right)\right|^2 - \mathrm{e}^{2\tau_{\text{ratio}}\delta}\left|\mathcal{D}_{i\tau_{\text{ratio}}\omega}\left(\frac{\mu_{\text{OFF}}-v_R}{\sqrt{\hat{D}}}\right)\right|^2}{\left|\mathcal{D}_{i\tau_{\text{ratio}}\omega}\left(\frac{\mu_{\text{OFF}}-v_T}{\sqrt{\hat{D}}}\right) - \mathrm{e}^{\tau_{\text{ratio}}\delta}\mathrm{e}^{i\omega\tau_R}\mathcal{D}_{i\tau_{\text{ratio}}\omega}\left(\frac{\mu_{\text{OFF}}-v_R}{\sqrt{\hat{D}}}\right)\right|^2}, \tag{3.42}$$

and to the OFF susceptibility (in u.ON.m.t.c.)

$$A_{\text{OFF}} = \tau_{\text{ratio}}\left(\frac{i\omega r_{\text{OFF}}}{\sqrt{\hat{D}}(i\tau_{\text{ratio}}\omega - 1)}\right)\frac{\mathcal{D}_{i\tau_{\text{ratio}}\omega-1}\left(\frac{\mu_{\text{OFF}}-v_T}{\sqrt{\hat{D}}}\right) - \mathrm{e}^{\tau_{\text{ratio}}\delta}\mathcal{D}_{i\tau_{\text{ratio}}\omega-1}\left(\frac{\mu_{\text{OFF}}-v_R}{\sqrt{\hat{D}}}\right)}{\mathcal{D}_{i\tau_{\text{ratio}}\omega}\left(\frac{\mu_{\text{OFF}}-v_T}{\sqrt{\hat{D}}}\right) - \mathrm{e}^{\tau_{\text{ratio}}\delta}\mathrm{e}^{i\omega\tau_R}\mathcal{D}_{i\tau_{\text{ratio}}\omega}\left(\frac{\mu_{\text{OFF}}-v_R}{\sqrt{\hat{D}}}\right)}. \tag{3.43}$$

3.8.3 Transformation of the LRT

The perturbed time evolution for the OFF cells is now given by

$$\frac{d\hat{v}_i^{\text{OFF}}}{d\hat{t}}(\hat{t}) = -\hat{v}_i^{\text{OFF}}(\hat{t}) + \mu_{\text{OFF}} + \hat{\xi}_i^{\text{OFF}}(\hat{t}) + \hat{f}(\hat{t}) - \langle\hat{f}(\hat{t})\rangle - \hat{\zeta}_i(\hat{t}).$$

This equation produces spike trains $\hat{y}_i^{\text{OFF}}(\hat{t})$ in u.OFF.m.t.c.. The LRT equation for the OFF cells corresponding to Eq. (3.15) is ($\epsilon = -1$ and the tilde denotes the Fourier transform):

$$\tilde{\hat{y}}_i^{\text{OFF}}(\hat{\omega}) = \tilde{\hat{y}}^{(0)}_{\text{OFF},i}(\hat{\omega}) + \hat{A}(\hat{\omega}, \mu_{\text{OFF}}, \hat{D})\left[-\tilde{\hat{\zeta}}_i(\hat{\omega}) + \frac{\hat{F}(\hat{\omega})}{2N}\left(\sum_{k=1}^{N}\tilde{\hat{y}}_k^{\text{ON}}(\hat{\omega}) + \sum_{k=1}^{N}\tilde{\hat{y}}_k^{\text{OFF}}(\hat{\omega})\right)\right],$$

with

$$\hat{F}(\hat{\omega}) = \hat{G}\frac{\mathrm{e}^{i\hat{\omega}\hat{\tau}_D}}{(1-i\hat{\omega}\hat{\tau}_S)^2} = \frac{G}{\tau_{\text{ratio}}}\frac{\mathrm{e}^{i\omega\tau_D}}{(1-i\omega\tau_S)^2} = \frac{1}{\tau_{\text{ratio}}}F(\omega)$$

and

$$\begin{aligned}\tilde{\hat{y}}_i^{\text{OFF}}(\hat{\omega}) &\equiv \int_0^{T/\tau_{\text{ratio}}} \hat{y}_i^{\text{OFF}}(\hat{t})\mathrm{e}^{i\hat{\omega}\hat{t}}d\hat{t} \\ &= \frac{1}{\tau_{\text{ratio}}}\int_0^T \hat{y}_i^{\text{OFF}}(t/\tau_{\text{ratio}})\mathrm{e}^{i\omega t}dt = \tilde{y}_i^{\text{OFF}}(\omega),\end{aligned}$$

because of Eq. (3.37). For $\tilde{\hat{\zeta}}_i(\hat{\omega}) = \sqrt{c}\tilde{\hat{\eta}}_c(\hat{\omega}) + \sqrt{1-c}\tilde{\hat{\eta}}_i(\hat{\omega})$, we have

$$\begin{aligned}\tilde{\hat{\zeta}}_i(\hat{\omega}) &= \int_0^{T/\tau_{\text{ratio}}} \hat{\zeta}_i(\hat{t})\mathrm{e}^{i\hat{\omega}\hat{t}}d\hat{t} \\ &= \frac{1}{\tau_{\text{ratio}}}\int_0^T \hat{\zeta}_i(t/\tau_{\text{ratio}})\mathrm{e}^{i\omega t}dt = \frac{1}{\tau_{\text{ratio}}}\int_0^T \zeta_i(t)\mathrm{e}^{i\omega t}dt,\end{aligned}$$

so that

$$\tilde{\hat{\zeta}}_i(\hat{\omega}) = \tilde{\zeta}_i(\omega)/\tau_{\text{ratio}}. \tag{3.44}$$

In the end, Eqs. (3.18) remain unchanged except for the replacement of the A_{OFF} of Eq. (3.13) by the one given by Eq. (3.43). The accuracy of this transformed LRT can be seen on Figs. 3.8a and 3.8c.

Chapter 4

Subtractive, divisive and non-monotonic gain control in feedforward nets linearized by noise and delays

The control of input-to-output mappings, or gain control, is one of the main strategies used by neural networks for the processing and gating of information. Using a spiking neural network model, we studied the gain control induced by a form of feedforward circuitry—designated here as "open-loop feedback"—which has been experimentally observed in a cerebellum-like structure in weakly electric fish. We found, both analytically and numerically, that this network displays three different regimes of gain control: subtractive, divisive, and non-monotonic. Subtractive gain control was obtained when noise is low in the network. Also, it was possible to change from divisive to non-monotonic gain control by simply modulating the strength of the open-loop feedback, which may be achieved via long-term synaptic plasticity. The particular case of divisive gain control has been previously observed *in vivo* in weakly electric fish. These gain control regimes were robust to the presence of temporal delays in the open-loop feedback pathway, which were found to linearize the input-to-output mappings (or f-I curves) via a novel variability-increasing mechanism. Our findings highlight the feedforward-induced gain control analyzed here as a highly versatile mechanism of information gating in the brain.

4.1 Introduction

The mapping between the input arriving to a neuron and its evoked firing rate has constituted one of the major interests in the study of neural systems over the last decades [162, 189, 184]. In many situations, neurons are able to perform a scaling operation on their response to input. For instance, contrast invariance of signal cancellation [143] and object representation [190], receptive field properties [5], and orientation selectivity [86] require a contrast-dependent scaling of responses in sensory areas. Gaze direction also scales the spiking response rate of neurons in the primary visual [208] and posterior parietal [8] cortices. Scaling can also be context-specific, as found in the auditory pathway of crickets [104]. Finally, it is known that cortical circuits are able to modulate their response gain depending on the input frequency [1, 211, 182] by means of short-term synaptic plasticity.

The scaling and control of input-output behavior of neural systems is often characterized by the so-called f-I curve, which displays the output firing rate vs input current to the neuron. In particular, the slope of such a dependency, or *gain*, constitutes a useful indicator of the behavior of the neuron. If the gain of the f-I curve is high, a small change in the input current will be mapped by the cell into a large change in the output firing rate, increasing the sensitivity of the neuron to weak stimuli. On the other hand, a low gain of the f-I curve translates large changes in the input current to small changes in the output firing rate, and this allows the neuron to encode a broad range of stimulus intensities into a physiologically plausible range of firing rates. This framework is strictly valid only when the stimulus evolves slowly compared to the integration time scale of the neuron, although the applicability of this formalism to fast varying stimuli has also been considered [130].

Certain mechanisms allow a neuron to modify its f-I curve, a phenomenon which is known as gain control. Figure 4.1 shows three possible examples of gain control on the behavior of a neuron. A simple gain control effect that one could think of is the subtractive gain control. In this case, the f-I curve experiences a subtractive (or additive) shift towards larger (or smaller) values of the input current without varying its slope. A large number of mechanisms, such as the introduction of some level of shunting inhibition [107], are able to produce this form of gain control.

Mechanisms providing the other two types of gain control shown in Fig. 4.1 have been more elusive. Divisive gain control is often assumed in rate models of neurons and neural populations [51, 57]. However, biophysical mechanisms for such a gain modulation have been hard to identify [107, 58, 78]. On the other hand, non-monotonic gain control has been reported to be induced by short-term depression [73, 122]. However, general biophysical

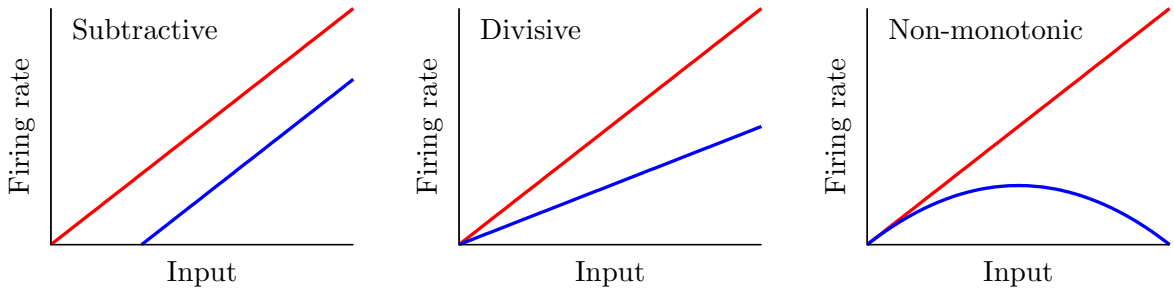

Figure 4.1: Three different gain control effects commonly observed in real neural systems. The red curve corresponds to the f-I curve before the gain modulation, and the blue curve shows it after the modulation. Subtractive gain control implies a shift of the original f-I curve, while divisive control leads to a change in slope. Non-monotonic gain control usually reflects more sophisticated input-output properties of neurons and synapses.

mechanisms which cover all these gain control strategies have been difficult to identify and characterize up to date.

In this work, we present a computational model of a neural circuit which is able to display these three types of gain control (subtractive, divisive and non-monotonic). We consider a generic neural circuit in which neurons receive a given stimulus both directly, i.e. from sensory receptors, and indirectly, via inhibitory interneurons driven by the same stimulus (a pathway referred here as *open-loop feedback*). Our results, both numerical and analytical, show that these neurons can exhibit the three different types of gain control described above (and shown in Fig. 4.1). The particular type of gain control exhibited by the system depends on (*i*) the noise level present on neurons, and (*ii*) the strength of the negative input provided by the inhibitory neurons, which may be easily modulated in real circuits via long-term synaptic plasticity and therefore provides a highly versatile gain control mechanism. We also carefully analyze the conditions under which the f-I curve of the neural system becomes non-monotonic.

In addition to these results, we also study the above gain control mechanism in the case in which the open-loop feedback introduces a temporal delay in the signal transmission (reflecting the presence of finite propagation times in the real brain). The existence of such a delay increases the variability of the effective current arriving to the neurons, which in turn leads to a linearization of the f-I curves. The gain control induced by the open-loop feedback is also present in this more realistic scenario. Finally, we use our model to explain the divisive gain control observed *in vivo* in the superficial pyramidal (SP) neurons of the electrosensory lateral-line lobe (ELL), a cerebellum-like structure of the weakly electric fish

Apteronotus leptorhynchus [[16]; see also Fig. 4.2A].

4.2 Materials & Methods

A simplified scheme of the cerebellum-like circuit that we consider in this study is shown in Fig. 4.2B. This circuit resembles the electrosensory lateral-line lobe (ELL), a primary sensory area in the brain of the weakly electric fish *Apteronotus leptorhynchus* [135], although our results can be easily generalized to any circuit in the brain presenting a degree of feedforward inhibition comparable to the one present in the circuit of Fig. 4.2B.

Briefly, sensory inputs coming from electroreceptors arrive at two different pools of neurons: the superficial pyramidal (SP) and the deep pyramidal (DP) neurons [133]. Both populations are mainly feedforward, and neurons within a population do not connect between themselves (i.e. no recurrent connectivity). From the population of DP neurons, the stimulus is transmitted to the nucleus praeminentialis (nP) and then to a population of granule cells called the eminentia granularis (EGp). A subset of neurons in the nP directly connects to SP cells and provides an inhibitory feedforward signal. Granule cells of the EGp transmit the signal they receive from a subgroup of nP cells to the SP cells via a massive set of parallel fibers [29]. Apart from a direct excitatory synaptic connection with SP cells, the parallel fibers also synapse onto inhibitory interneurons which then project onto the soma and/or apical dendrites of SP cells and strongly inhibit them. As a consequence of this strong inhibition (or, more precisely, disynaptic inhibition), the overall contribution of parallel fibers to SP neurons is mainly inhibitory [16]. Note that another subset of nP neurons is excitatory and connects directly to the SP cells as well as disynaptically through inhibitory interneurons. A compelling option is to amalgamate the latter pathway with the other one going directly from the nP to the SP cells. We shall make no further reference to this third pathway. The pathway followed by the sensory stimulus going from DP neurons to SP neurons via the nP and the EGp is called open-loop feedback. (Both populations receive the stimuli, and thus lie at the same processing stage. However, DP cells send signals to SP cells but do not receive them back.) Finally, neurons in the SP population project onto higher areas of the brain.

We consider a population of N deep pyramidal (DP) neurons receiving sensory input from electroreceptors. The membrane potential $V_i^D(t)$ of the DP neuron i follows a simple leaky integrate-and-fire (LIF) dynamics,

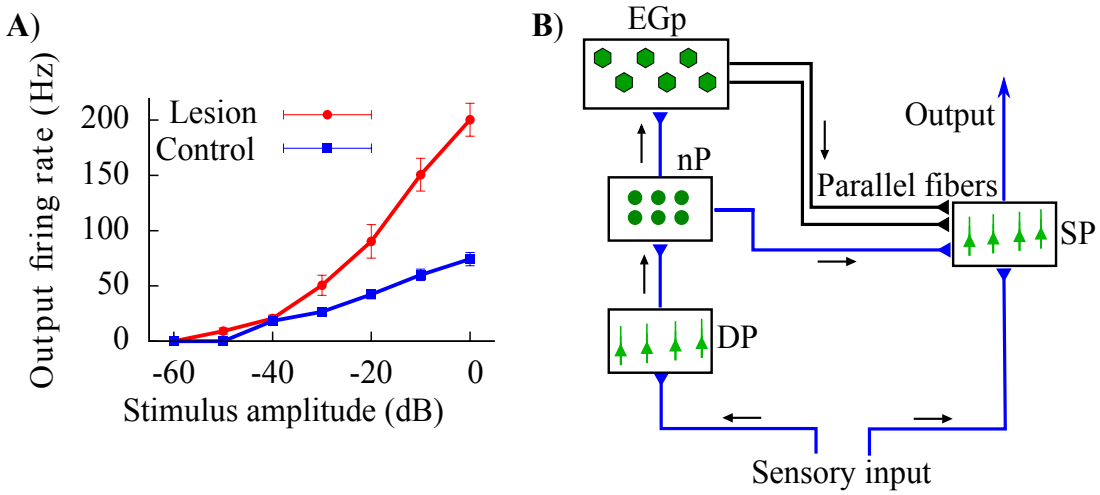

Figure 4.2: (**A**) *In vivo* experimental recording of the firing rate (relative to baseline activity) of SP neurons in the ELL of the weakly electric fish, as a function of the stimulus intensity. Data from control fish are compared with data from fish for which synapses providing the open-loop feedback via the cerebellar parallel fibers have been removed (lesion). This implies that the effect caused by the parallel fibers is mainly a divisive gain control. Data taken from [16]. (**B**) Scheme of the neural circuit considered in our study. Black arrows show the direction of the input transmission. Note that both SP and DP cells receive the sensory input.

$$\tau_m \frac{dV_i^D(t)}{dt} = -V_i^D(t) + \mu + \sqrt{\tau_m}\,\sigma\,\xi_i^D(t), \tag{4.1}$$

where τ_m is the membrane time constant, μ is the sensory input (that we consider constant for simplicity), and $\xi_i^D(t)$ is a Gaussian white noise of zero mean and delta-type autocorrelation $\left\langle \xi_i^D(t)\,\xi_j^D(t+\tau)\right\rangle = \delta_{ij}\,\delta(\tau)$, with δ_{ij} and $\delta(t)$ being the Kronecker and Dirac delta, respectively. This noise term was included in the model to reflect the intrinsic stochasticity of the deep neurons [20]. The factor σ reflects the noise intensity of this stochastic term.

Following the typical dynamics of the LIF model, when the membrane potential reaches a certain threshold V_{th}, an action potential is generated by the neuron and the membrane potential is reset to V_r, and remains there for a refractory period τ_r.

In the real neural circuit, the population of DP neurons projects onto the nP, then some nP cells project directly onto SP cells while others connect with the EGp, whose granule cells make synaptic contact with the SP neurons both directly and through disynaptic inhibition. In our model, we simplify these intricate feedforward connections by assuming that the activity of the DP population ultimately drives the dynamics of the SP cells, so that the input

that a SP neuron receives from the open-loop feedback is given by

$$f(t) = \tau_m \, G \, \frac{1}{N} \sum_{i=1}^{N} \sum_{t_{i,k}} s(t - t_{i,k}). \tag{4.2}$$

Here, the first sum runs over all DP neurons, and the second sum runs over the spike times of each presynaptic DP neuron (i.e. $t_{i,k}$ is the k-th spike time from the i-th presynaptic neuron). The factor G/N may be identified as the effective strength of synapses onto the SP cells. Since the net polarity of the open-loop feedback seems mainly inhibitory [16, 75], we assume $G \leq 0$.

The synaptic filter function, $s(\tau)$, describes the effect of a given presynaptic spike on the postsynaptic potential, and it is given by a delayed alpha function

$$s(\tau) = \left(\frac{\tau - \tau_d}{\tau_s^2} \right) \, e^{-\frac{\tau - \tau_d}{\tau_s}} \, \Theta(\tau - \tau_d), \tag{4.3}$$

where τ_s is the synaptic transmission time scale, τ_d is a temporal delay, and $\Theta(x)$ is the Heaviside step function (i.e. $\Theta(x) = 1$ if $x > 0$, and $\Theta(x) = 0$ otherwise). Note that $s(\tau)$ is normalized so that integration over a large enough time window yields unity.

Equation 4.2 can also be written by using a convolution of the population average for the DP cells and the synaptic filter. The population average is

$$Y^D(t) = \frac{1}{N} \sum_{i=1}^{N} y_i^D(t), \tag{4.4}$$

where $y_i^D(t) = \sum_{t_{i,k}} \delta(t - t_{i,k})$ is the spike train of DP neuron i. Hence,

$$f(t) = \tau_m \, G \, (s * Y^D)(t), \tag{4.5}$$

where $(a * b)(t)$ is the convolution of functions $a(t)$ and $b(t)$.

Since all SP neurons will receive, on average, the same input from sensory receptors and from the open-loop feedback, we will consider the response of a typical SP cell, and the output statistics will be valid for all the other SP neurons. The membrane potential $V^S(t)$ of a typical SP neuron is described, as in the case of DP neurons, by a simple leaky integrate-and-fire (LIF) dynamics,

$$\tau_m \frac{dV^S(t)}{dt} = -V^S(t) + \mu + f(t) + \sqrt{\tau_m}\,\sigma\,\xi^S(t), \tag{4.6}$$

where, as for the DP neurons, μ is the sensory input and the last term considers the intrinsic stochasticity of the neuron (with $\xi^S(t)$ being a Gaussian white noise of zero mean and delta-type autocorrelation). The term $f(t)$ is the input coming from the DP cells and is already known. Again, when the membrane potential reaches the threshold V_{th}, a spike is generated and $V^S(t)$ is reset to V_r during a period τ_r.

Unless specified otherwise, we choose the following values for the time constants: $\tau_m = 10$ ms, $\tau_r = 1$ ms, $\tau_s = 5$ ms and $\tau_d = 10$ ms. For parameters related to the membrane potential we consider, without loss of generality, dimensionless units. We set $V_{th} = 1$ and $V_r = 0$, and therefore μ, σ and G will be in "resting-to-threshold" units. For simulations, we used a DP population of $N = 500$ neurons unless another size is specified.

4.3 Results

4.3.1 Subtractive gain control

We start our analysis by considering the limit case in which neurons in our system are deterministic (i.e. $\sigma = 0$). In this case, the mean firing rate of the DP neurons may be easily obtained by solving Eq. 4.1 for a given neuron. Since the input is the same for all DP neurons, their firing rate will also be the same. One obtains [212]

$$r_D = \left[\tau_r + \tau_m \log\left(\frac{V_r - \mu}{V_{th} - \mu}\right)\right]^{-1} \Theta(\mu - V_{th}). \tag{4.7}$$

For convenience, we define the *effective bias* to the SP cell as the sum of the sensory input plus the open-loop feedback contribution, i.e. $\mu_{\text{eff}} = \mu + f(t)$. In the diffusion limit [40], and ignoring fluctuation terms, the effective bias can be written as

$$\mu_{\text{eff}} = \mu + \tau_m\, G\, r_D. \tag{4.8}$$

In the case of a deterministic system, it takes the form

$$\mu_{\text{eff}} = \mu + \frac{\tau_m\, G\, \Theta(\mu - V_{th})}{\tau_r + \tau_m \log\left(\frac{V_r - \mu}{V_{th} - \mu}\right)}. \tag{4.9}$$

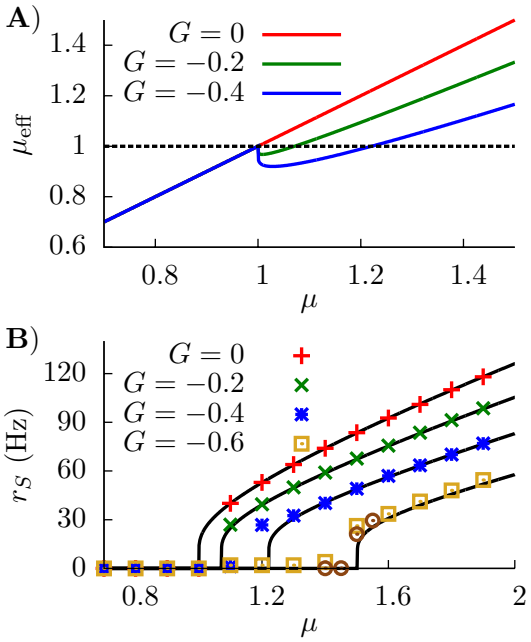

Figure 4.3: (**A**) Effective bias entering the SP neuron as a function of the sensory input μ, for the deterministic case ($\sigma = 0$). Red line ($G = 0$) indicates the effective bias for the case in which no feedback is considered. (**B**) Response of the SP neuron to sensory input (i.e. the f-I curve of the SP cell circuit), for the deterministic case. Numerical simulations of a network of $N = 100$ DP neurons (symbols) show a very good agreement with our theoretical estimations (lines). Circles denote the $G = -0.6$ case for a network of $N = 500$ DP neurons, indicating that the small discrepancies between simulations and theory in the firing onset are due to finite-size effects of the simulation.

Finally, the expression for the SP firing rate is given by

$$r_S = \left[\tau_r + \tau_m \log\left(\frac{V_r - \mu_{\text{eff}}}{V_{th} - \mu_{\text{eff}}}\right)\right]^{-1} \Theta(\mu_{\text{eff}} - V_{th}). \tag{4.10}$$

The dependence of the effective bias on the sensory input, given by Eq. 4.9, is depicted in Fig. 4.3A. By careful inspection of this figure, one can see that the effective bias increases linearly with μ until $\mu = V_{th}$. Beyond that point, DP cells start firing and the effective bias starts decreasing, therefore frustrating the firing of the SP cells. For large enough values of μ, the denominator in the second term of the r.h.s. of Eq. 4.9 tends to τ_r, and the effective bias becomes $\mu_{\text{eff}} \simeq \mu + \tau_m\, G/\tau_r$. This constitutes a simple shift of the input μ, and one should therefore expect to observe a subtractive effect on the f-I curve of the SP neurons. This is indeed what we observe in Fig. 4.3B, where the results of numerical simulations of

the model for $\sigma = 0$ show a very good agreement with our theoretical expression given by Eq. 4.10. We can observe that the displacement of the firing onset (the value of μ for which the neuron starts to fire) depends on G as theoretically predicted, with more negative values of G leading to larger shifts of the onset of firing. Therefore, the modulation of the feedback strength G leads to a subtractive gain control in the f-I curve of the deterministic ($\sigma = 0$) system.

4.3.2 Divisive gain control

After analyzing the deterministic case, we can now study the more general case in which neurons in the circuit present some level of stochasticity (i.e. $\sigma > 0$) in their dynamics. Such a stochasticity may be due, for instance, to the noisy dynamics of ion channels [222], or to synaptic bombardment from other surrounding neurons [105], for instance (see also [128] for a review).

For the stochastic case $\sigma > 0$, the mean firing rate of the DP neurons may be obtained by solving Eq. 4.1, and it is given by [213]

$$r_D = \left[\tau_r + \tau_m \int_{z_r^D}^{z_{th}^D} \sqrt{\pi}\, e^{z^2}\, (1 + \mathrm{erf}(z)) dz \right]^{-1}, \tag{4.11}$$

where $z_{th}^D \equiv \frac{V_{th} - \mu}{\sigma}$, $z_r^D \equiv \frac{V_r - \mu}{\sigma}$, and $\mathrm{erf}(z)$ is the error function. The effective bias is again $\mu_{\mathrm{eff}} = \mu + \tau_m\, G\, r_D$, although the explicit form is more complex now. Finally, we can write the mean firing rate of SP neurons as

$$r_S = \left[\tau_r + \tau_m \int_{z_r^S}^{z_{th}^S} \sqrt{\pi}\, e^{z^2}\, (1 + \mathrm{erf}(z)) dz \right]^{-1}, \tag{4.12}$$

where $z_{th}^S \equiv \frac{V_{th} - \mu_{\mathrm{eff}}}{\sigma}$ and $z_r^S \equiv \frac{V_r - \mu_{\mathrm{eff}}}{\sigma}$.

The main effect of the addition of noise to the DP neuron model is the appearance of a smooth *linearization* around the onset of the f-I curve for these neurons [78]. This is due to noise-induced firing, which is especially important when the neuron is slightly below the firing threshold.

As a consequence of this linearization, the effective bias μ_{eff} may be approximately described, for values of μ close to the DP firing onset, as

$$\mu_{\text{eff}} = \mu + \tau_m \, G \, r_D \simeq \mu + \tau_m \, G \, (C_1 \mu + C_2), \tag{4.13}$$

where C_1 and C_2 are constants. The effective bias, therefore, would be approximately linear, with a slope that depends on G. Such a linear dependence is shown in Fig. 4.4A, where we can see that μ_{eff} is indeed well approximated by a linear relationship with μ, and that the slope depends on G. In particular, as G goes to more negative values the slope of the $\mu_{\text{eff}} - \mu$ relationship decreases. We will restrict ourselves here to the case of G having relatively small absolute values, so that μ_{eff} will be an increasing function of μ in all cases. Large negative values of G, which could compromise this tendency, will be addressed in the next section.

Interestingly, since the f-I curve of an isolated SP neuron is also linearized by the presence of stochasticity, the above approximation can be applied again: the multiplicative effect that G has on the slope of μ_{eff} will also affect the slope of the SP f-I curve in the same way. This is shown in Fig. 4.4B, where one can observe that increasing the strength of the inhibitory feedback has a divisive effect on the f-I curve of the SP neuron, as predicted. The figure also shows the good agreement between numerical simulations of the model and the theoretical expression, given by Eq. 4.12. Figure 4.4C shows more clearly the relation between G and the SP gain, and we observe that larger negative values of G produce a smaller slope on the SP f-I curve. Consequently, the modulation of the feedback strength (in the range of small $|G|$) provides a divisive gain control to the system when stochasticity is considered.

In the mechanistic description presented above, the divisive gain modulation strongly depends on the fact that DP neurons drive the response of SP neurons without any restrictions (such as feedback from other cells to the DP neurons). This approach differs from previous attempts to explain the divisive gain control found *in vivo* [16], such as the one presented in [203]. In this previous approach, Sutherland *et al.* considered a unique population of ELL pyramidal cells, without distinguishing between DP and SP neurons. The system was then assumed to display a closed-loop feedback circuit, in which the population of pyramidal neurons projected their output to a feedback kernel, which in turn inhibited the activity of the population. Because the negative feedback was affecting all neurons in this case, a high firing rate would be prevented by the closed-loop inhibition. In the approach we present here, on the other hand, DP neurons (which drive the gain control of SP cells) do not receive inhibitory feedback, and therefore they are able to raise their firing rate higher and produce a stronger and more effective modulation of the SP firing rate than in the closed-loop scenario. This can be seen in Fig. 4.4D, where our open-loop model is compared with a closed-loop version of the same model for the same values of the parameters (the case $G = 0$ is also shown for comparison). In the closed-loop model, the firing rate of the unique pyramidal cell population ν is given by

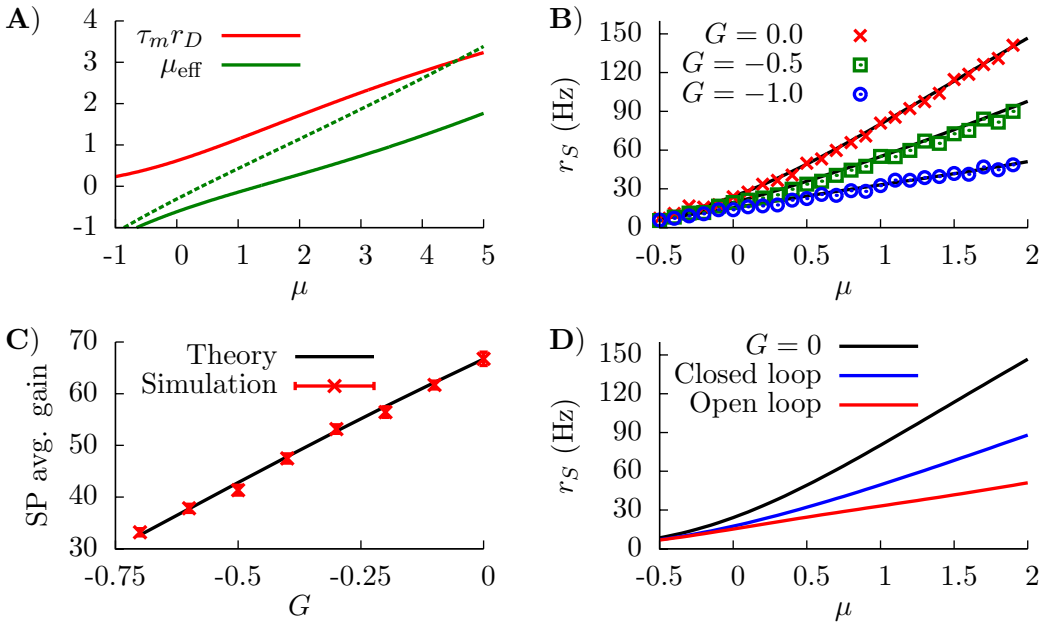

Figure 4.4: (**A**) DP firing rate (red line) and effective bias (solid green line) as a function of the sensory input μ. One can see that both curves can be considered as approximately linear with μ (this is valid as long as $|G|$ is not very large). Parameters are $\sigma = \sqrt{3}$, $G = -1$. The green dotted curve is the same as the solid green line, but for $G = -0.5$. (**B**) SP firing rate as a function of the sensory input, for different values of the feedback strength and $\sigma = 1$. The divisive gain control occurs as G takes larger negative values. (**C**) Average gain of the SP f-I curve as a function of G, for $\sigma = 1$. The average gain is obtained by fitting the f-I curve with a linear function over a range of μ values for which the curve is approximately linear. The displayed results differ by the origin of the f-I curves used. For the solid black line, the rates are computed from the analytical formula [Eq. 4.12]. For the red symbols, the f-I curves are extracted from simulations of network. The error bars come from the fitting procedure, and serve to illustrate how linear the f-I curves are for the chosen μ ranges. (**D**) Comparison between the divisive gain control in our model (open-loop feedback) and the one studied in [203] (closed-loop feedback), for $\sigma = 1$. As the panel shows, the open-loop feedback mechanism studied here provides a stronger divisive effect than the one studied by Sutherland et al. for the same feedback strength and conditions. Both the open-loop and closed-loop cases correspond to the same feedback strength, $G = -1$. For all panels, solid or dotted lines come from theoretical (analytical) expressions, whereas symbols come from simulations.

$$\nu = \left[\tau_r + \tau_m \int_{z_r}^{z_{th}} \sqrt{\pi}\, e^{z^2}\, (1 + \mathrm{erf}(z))dz\right]^{-1}, \tag{4.14}$$

where $z_{th} = \frac{V_{th} - \mu - \tau_m G\nu}{\sigma}$ and $z_r = \frac{V_r - \mu - \tau_m G\nu}{\sigma}$. It is worth noting that both z_{th} and z_r depend on the population firing rate ν, and therefore Eq. 4.14 has to be solved recursively. As Fig. 4.4D shows, the open-loop model proposed here allows for a stronger modulation than the closed-loop model for the same parameter values and, in particular, for the same value of G in both cases.

4.3.3 Non-monotonic gain control

In the previous section, we assumed small absolute values of G to simplify our treatment. This allowed us to understand, from a qualitative point of view, the origin of the divisive gain control in our open-loop feedback system, and our findings were supported by both our numerical simulations and our theoretical description. However, these results could be different for larger values of G, as a strong inhibitory feedback driven by sensory input could overcome the excitatory effects of this same sensory input on SP cells. If that were the case, μ_{eff} would no longer increase with μ (or, at least, not for all values of μ), and the effect of G on the f-I curve might change. In order to explore such a possibility, we need to consider a more careful analysis of the stochastic case (i.e. $\sigma > 0$) for large values of G.

We start by identifying potential extrema in the f-I curve of the SP neuron, which would be an expected effect of inhibitory feedback overcoming excitation in SP cells for a given value of μ. Since the SP firing rate is a monotonically increasing function of μ_{eff} (see Eq. 4.12), finding the extrema of the effective bias as a function of μ would be equivalent to finding the extrema of the f-I curve for SP cells. The condition for extrema of the effective bias $\mu_{\mathrm{eff}} = \mu + \tau_m\, G\, r_D$ can be obtained from $\frac{d\mu_{\mathrm{eff}}}{d\mu} = 0$, and it gives

$$\frac{dr_D}{d\mu} = -\frac{1}{\tau_m\, G}, \tag{4.15}$$

or, in a more explicit form,

$$\sigma + \sqrt{\pi}\, \tau_m^2\, G\, r_D^2 \left[e^{z_{th}^2}(1 + \mathrm{erf}(z_{th})) - e^{z_r^2}(1 + \mathrm{erf}(z_r))\right] = 0. \tag{4.16}$$

We will consider the condition of extrema in its simplified version (i.e. Eq. 4.15) for clarity. The relationship r_D vs. μ, given by Eq. 4.11, takes a sigmoidal-like shape: r_D tends to zero for large negative μ, and to $1/\tau_r$ for large positive μ. The maximum slope reached by r_D (which

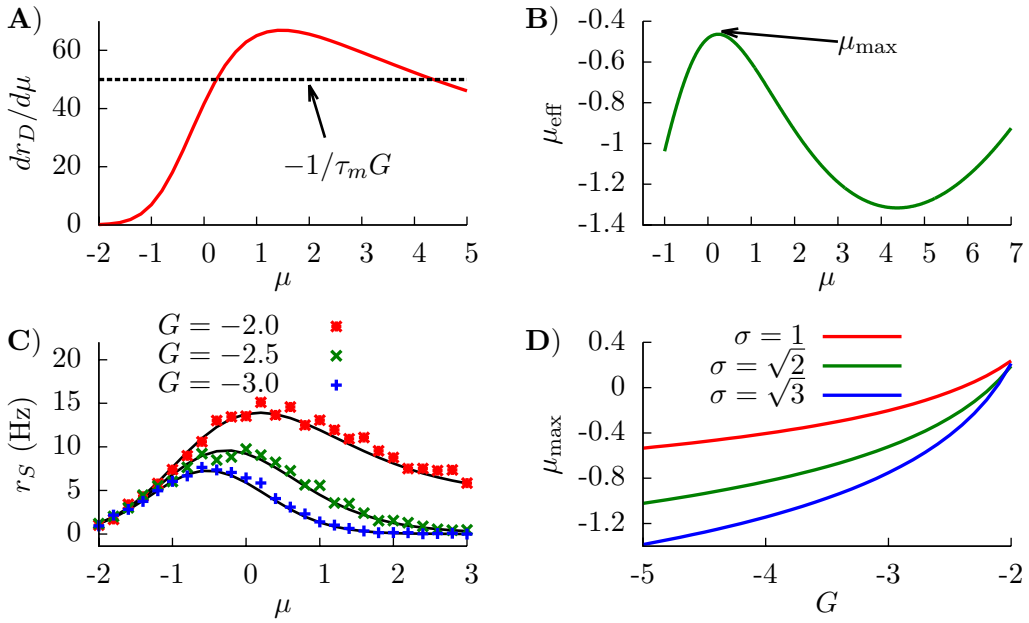

Figure 4.5: (**A**) Graphical representation of the condition for the existence of extrema (Eq. 4.15). We can see that solutions for this equation exist only when γ (the peak value of the curve) is larger that $\frac{1}{\tau_m|G|}$ (dashed line). Parameters are $\sigma = 1$ and $G = -2$. (**B**) Effective bias as a function of μ, showing the existence of a maximum and a minimum. Parameters are $\sigma = 1$ and $G = -2$. (**C**) SP firing rate as a function of the input μ, for different values of G, displaying non-monotonic gain control. We choose $\sigma = \sqrt{2}$. (**D**) Position of the SP firing rate peak (shown in panel C) as a function of G, for different values of σ. In all panels, solid lines come from theoretical (analytical) expressions, whereas symbols come from simulations.

we will denote as γ) would be located somewhere between these two limits. Therefore, $\frac{dr_D}{d\mu}$ would follow approximately a bell shape: it is zero at both ends of the μ axis, and it has a maximum (of value γ) for a moderate value of μ.

With this information, and together with Eq. 4.15, it is easy to see that extrema will exist only when

$$\gamma > \frac{1}{\tau_m\,|G|}, \tag{4.17}$$

which typically occurs for large enough absolute values of the feedback strength. In this case, we will have two extrema (since $\frac{dr_D}{d\mu}$ is bell-shaped and will cross the constant level $\frac{1}{\tau_m\,|G|}$ twice). Figure 4.5A shows $\frac{dr_D}{d\mu}$ as a function of μ, in a situation in which the two extrema exist (i.e. Eq. 4.17 is satisfied).

The concrete shape of the effective bias with μ is shown in Fig. 4.5B. The two extrema correspond to a maximum (at small μ) and a minimum (at larger μ). This yields a bell-shape dependence of the effective bias on μ (at low values), followed by an increase with μ (for high values). The cause for this behavior is the following: at very small values of μ, the firing rate of the DP neurons is not high enough, and the main contribution to the effective bias is the sensory stimulus μ, with the feedback playing a minor role. As μ makes the DP firing rate increase, the feedback term becomes dominant (since G is negative and large) and the effective bias starts to decrease, completing the bell-shape profile observed in Fig. 4.5B. After that, and for very large values of μ, the DP firing rate approaches its saturation value $1/\tau_r$. At this point the sensory stimulus μ becomes the dominant term and μ_{eff} begins to rise again. It is worth noting that this second rising of μ_{eff} occurs at input levels where the firing rate of DP neurons saturates, and such sensory input levels are beyond the range in which biologically relevant information can be linearly processed in the system, according to electroreceptor input-output characteristics [98]. The attainment of the maximum rate, however, occurs for biologically sound biases, as exemplified in Fig. 4.5C.

Since the firing rate of SP cells monotonically increases with the effective bias, the maximum and minimum found in μ_{eff} are also observed for the SP firing rate, as shown in Fig. 4.5C. This implies that, for a stochastic system, enhancing the negative feedback strength may drive the system from a divisive to a non-monotonic gain control regime. The figure also shows that the location and height of the peak of the SP firing rate is also modulated by G. This is to be expected, since the stronger the feedback, the sooner the DP firing rate will start decreasing the effective bias. Large negative values of G will shift the location of the SP firing rate peak towards lower values of μ, as can be observed in Fig. 4.5C and, with more detail, in Fig. 4.5D. The level of stochasticity has also an impact on the position of the peak: larger values of σ will increase the overall firing rate of DP neurons, and the resulting increment in the feedback will shift the peak towards even lower values of μ, as Fig. 4.5D also shows.

Our neural circuit, therefore, is able to display three different regimes, corresponding to subtractive, divisive, and non-monotonic gain control. As the theory and numerical simulations show, both the level of stochasticity σ and the feedback strength G play an important role in the behavior of the system. Figure 4.6 shows a phase diagram of the system as a function of these two parameters. The subtractive gain control is observed only for deterministic systems ($\sigma = 0$), since the introduction of some level of noise would smooth the DP firing onset and we would go into other regimes. For $\sigma > 0$ the system can be in the divisive or non-monotonic regime, depending on the value of G. As discussed in the previous section, small absolute values of G would linearize the DP rate, effective bias

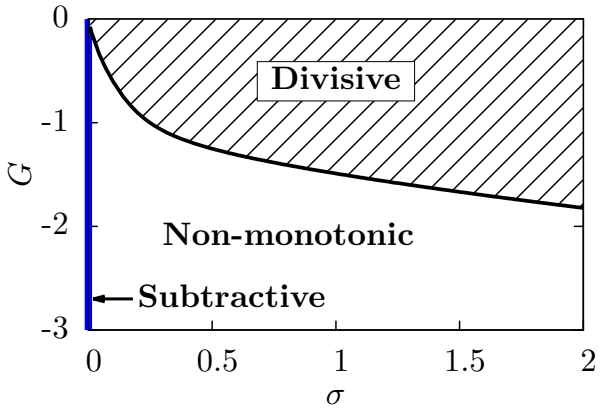

Figure 4.6: Phase diagram of the neural system studied. The subtractive gain control regime corresponds to the $\sigma = 0$ axis, and the critical line separating the divisive and the non-monotonic gain control regimes is given by Eq. 4.16.

and SP rate with respect to μ, leading to the appearance of divisive gain control. For large negative G, the feedback can eventually overcome the input μ to the SP neuron and the circuit will display non-monotonic gain control. The critical line separating the divisive and non-monotonic regimes is given by Eq. 4.16.

We simplified the above analysis of the system by assuming that the SP and DP cells share the same properties, namely the same bias (μ) and the same level of noise (σ). We argue that this analysis can be extended to more general contexts. In real systems, neural attributes are not uniform. For instance, in the cortex, inhibitory interneurons fire irregularly in response to constant inputs *in vitro* [202], whereas pyramidal neurons are relatively regular [131]. In the ELL, even though both the SP and the DP cells receive a common external input, they do not possess the same variability [20], which may translate into different biases and noise intensities. Above, we amalgamated the intrinsic bias and the applied current (I_{app}, say) under the umbrella variable μ. To accommodate different neural properties, we could use two separate intrinsic biases (μ_S and μ_D) and noise intensities (σ_S and σ_D), and study the SP firing rate as a function of I_{app}. This could produce slightly different results than above. As an example, for the noiseless case ($\sigma_S = \sigma_D = 0$), the SP cells could start to fire before the DP neurons—provided that μ_S is large enough—stay active for a limited range of I_{app} values, become silent when the DP cells attain a sufficient rate, and then fire again for higher I_{app}. This would constitute a mixed regime showing both non-monotonic and subtractive gain controls. Note that this mixed regime is also observable for uniform neural properties when the noise is low. In the latter case, the isolated blob of nonzero firing rate is due to noise.

4.3.4 Time-dependent stimuli and synaptic delays

Real neural systems deal with time-dependent stimuli. In sensory networks, neurons are responsible for encoding the temporal features of information coming from the environment. When dealing with time-dependent signals, a number of factors have to be considered in addition to the ones included up to now in our study. For instance, the existence of temporal delays in the transmission of information, due to the finite transmission speed of action potentials along the axons, becomes important. In the particular case of the ELL, signals from the DP cells have to travel through different neural populations before arriving to SP cells. Traveling such a distance causes a temporal delay in the transmission that, in some cases, may reach tens of milliseconds [135, 32]. Although this temporal delay did not have an effect in our previous results, with μ being constant in time (data not shown), it might have a significant impact in more realistic situations, when the sensory input presents temporal fluctuations and some level of autocorrelation. Different synaptic filter functions $s(\tau)$ might have an effect in these realistic situations as well.

To test our model for these conditions we use, as sensory input, a time-varying quantity $\hat{\mu}(t)$ of the form

$$\hat{\mu}(t) = \mu + \zeta(t), \tag{4.18}$$

where μ is the constant bias, and $\zeta(t)$ is a gaussian low-pass filtered noise of zero mean and standard deviation σ_c. This noise term is generated using a Butterworth fourth-order digital filter with cut-off frequency of 100 Hz. In the following, to explore the behavior of our model for other synaptic filter functions, we consider a delta-type synaptic function $s(\tau) = \delta(\tau - \tau_d)$.

When the temporal delay τ_d is set to zero, the main effect of considering this slowly fluctuating input is a slight increase in the SP firing rate along the f-I curve, as Fig. 4.7A shows. The black line in the figure shows the f-I curve for $\sigma_c = 0$, and the red line shows the same f-I curve for $\sigma_c = 1$. This firing rate increase is simply due to the presence of the extra fluctuating term $\zeta(t)$.

The figure also shows, interestingly, that considering temporal delays τ_d larger than zero leads to a significant increment in the SP firing rate, especially around the firing onset. To understand this phenomenon, it is useful to consider a slow (with respect to τ_m) positive fluctuation in $\hat{\mu}(t)$, arriving at DP and SP neurons simultaneously. If the temporal delay τ_d of our circuit from DP cells to SP cells is zero, the transient increment in the sensory input received by the SP neuron will be compensated almost simultaneously (apart from the

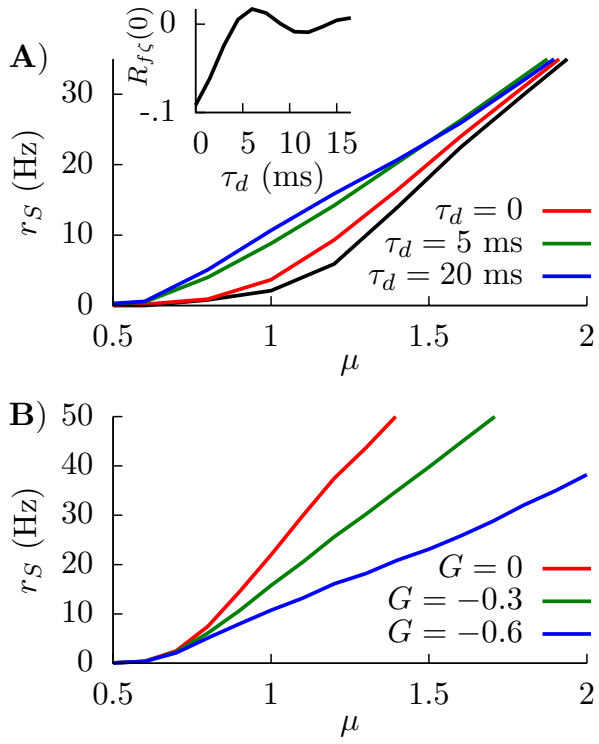

Figure 4.7: (**A**) Simulation results of the effect of the temporal delay on the f-I curve of the system at $G = -0.6$. The black line corresponds to the case in which the sensory input is just a constant bias and only the white noise is present. For other curves, the fluctuating term $\zeta(t)$ is also present, and the effect of varying τ_d is shown. Inset: Theoretical dependence of the covariance of $f(t)$ and $\zeta(t)$ on the temporal delay. (**B**) Simulation results of the SP firing rate vs constant bias for different values of G, a white noise term, and a slow fluctuating term $\zeta(t)$, for a temporal delay $\tau_d = 20$ ms. In both panels, we have $\sigma = 0.1$, $\sigma_c = 0.3$ and $\tau_m = 15$ ms.

response time) by the transient increment of the DP firing rate, which will have an inhibitory effect on the SP neuron since $G < 0$. The effective bias received by the SP neuron, therefore, will be mainly unaffected by such a fluctuation.

In the presence of some level of temporal delay, however, this "compensation" from the DP cells may arrive late, maybe once the transient increment in SP firing due to the fluctuation is over. Even more, this late compensation from DP cells could cause a transient decrement in the effective bias that the SP neuron receives. As a consequence, the effective bias arriving at the SP cells will be highly fluctuating, and these input fluctuations will cause an increment in the SP firing rate. This novel variability-increasing effect would be especially important around the SP firing onset, and our simulation results confirm this point (Fig.

4.7A). The effect of the temporal delay on the level of fluctuations in the effective bias can also be mathematically derived (see Appendix, section 4.5). Indeed, our calculations show (inset in Fig. 4.7A) that the covariance of $f(t)$ and $\zeta(t)$, and therefore the fluctuations of the effective bias, increases with the temporal delay τ_d (with some minor oscillatory component), supporting these findings on the delay-induced effects on SP firing rate.

In Fig. 4.7B, one can see an example of the divisive gain control for this more realistic situation with delay. In this case, the sensory input has a fluctuating part $\zeta(t)$ and a temporal delay of 20 ms. The behavior of the system is qualitatively similar to the one observed previously (i.e. large negative values of G induce a divisive effect on the gain of the f-I curve). The non-monotonic behavior also holds for this more realistic model (data not shown).

In our treatment of the effect of delay on the f-I curve, we assumed that the feedback strength G is a constant. That is, G had the same magnitude for all frequencies contained in the signal. This may not be the case for real systems. For instance, in the weakly electric fish, the EGp parallel fibers are known to be inactive when sinusoidal signals of higher frequency than about 20 Hz are applied to the network [32]. On the other hand, the direct inhibition coming from the nP seems to be active for an extended frequency range [77]. So, G may generally depend on frequency in a non trivial way. This fact can easily be included into the result for the covariance of $f(t)$ and $\zeta(t)$ appearing in the Appendix (section 4.5), provided that one knows — at least approximately — how G changes with frequency. If, for example, G is zero above a given frequency, then, obviously, signals only containing frequencies above that cutoff should have no effect on the f-I curve.

4.3.5 Comparison with experimental data

We now apply our model to *in vivo* data showing gain control in SP cells of weakly electric fish [16]. In Fig. 4.2A, reproduced from [16], the firing rates of SP cells in control fish are compared with those in fish for which the EGp has been lesioned. The result of this procedure is to eliminate the signal coming from the parallel fibers. Bastian's results show that the EGp feedback modulates the f-I curve of SP neurons, producing a divisive gain control. Since this form of gain control has been identified as one of the behaviors displayed by our model, it might be interesting to test whether the model is able to quantitatively explain the experimental data. Here, we shall restrict ourselves to analytical f-I curves (Eq. 4.12) since they have been shown to be in good agreement with those extracted from numerical simulations of Eq. 4.1 and Eq. 4.6. The experimental results are given as relative firing rates versus the stimulus intensity in decibels. To apply our model, we first need to convert these relative firing rates to absolute ones and relate the experimental stimuli to our μ.

The stimuli are amplitude modulations (AMs, in mV/cm) of the fish's electric field. In mathematical expressions we shall denote the intensity of these AMs by the symbol E_{AM}. The zero-decibel reference is $E_{\mathrm{AM}}^{\mathrm{ref}} = 2$ mV/cm [16], so that the number of decibels is given by $20\log_{10}(E_{\mathrm{AM}}/E_{\mathrm{AM}}^{\mathrm{ref}})$. Relative firing rates for the lesion and control cases are given by the rates for given AMs from which is subtracted the rate in the spontaneous regime (i.e. when $E_{\mathrm{AM}} = 0$). More precisely, absolute firing rates are given by

$$r_S^{C(L)}(E_{\mathrm{AM}}) = r_S^{C(L)}(0) + \Delta r_S^{C(L)}(E_{\mathrm{AM}}), \tag{4.19}$$

where $\Delta r_S^{C(L)}(E_{\mathrm{AM}})$ represents relative firing rates in the control (lesion) case. From [15], it appears that $r_S^C(0) - r_S^L(0) \sim 10$ Hz, which means that the net effect of the EGp parallel fibers is actually excitatory in the spontaneous regime. On the other hand, it is clear from Fig. 4.2A that the EGp feedback is inhibitory for larger E_{AM}. Therefore, this feedback changes sign somewhere between 0 mV/cm and 2 mV/cm. From inspection of Fig. 4.8A, in which absolute firing rates are plotted, this transition occurs for a small AM value (where the lesion and control curves intersect).

We now relate the AMs to our μ. The SP and DP cells receive input from electrosensory afferents, which encode the AMs of the electric field. We consider that μ encompasses an intrinsic bias μ_0 and a simple summation of these afferent inputs. The average input to a ELL pyramidal cell is then

$$\mu = \mu_0 + \tau_m G_{\mathrm{aff}} r_{\mathrm{aff}}(E_{\mathrm{AM}}) \equiv \mu_0 + \gamma(E_{\mathrm{AM}}), \tag{4.20}$$

where G_{aff} is the afferent synaptic strength times the number of afferents (compare with Eq. 4.8). The afferents' firing rate, $r_{\mathrm{aff}}(E_{\mathrm{AM}})$, is a sigmoid taking the form of a Boltzmann function [26]:

$$r_{\mathrm{aff}}(E_{\mathrm{AM}}) = \frac{r_{\max}}{1 + e^{-k(E_{\mathrm{AM}} - h)}}, \tag{4.21}$$

where $r_{\max}$ is the maximum firing rate. The slope of the linear part of this curve is given by $r_{\max}k/4$, whereas h is the AM value at which $r_{\mathrm{aff}} = r_{\max}/2$. Replacing Eq. 4.21 into Eq. 4.20 provides the required relationship between μ and E_{AM}. We further define the constant $A \equiv \tau_m G_{\mathrm{aff}} r_{\max}$.

To fit the data corresponding to a lesioned fish, i.e., which lacks the set of parallel fibers connecting the EGp with the SP cells, we set $G = 0$. This is not rigorously true since, from Fig. 4.2B, the direct inhibitory signal from the nP could still be active, and hence G should not be strictly zero. However, given the overall large effect that the parallel fibers have on the SP cells' f-I curve (Fig. 4.2A), it seems plausible to neglect the nP component of the total

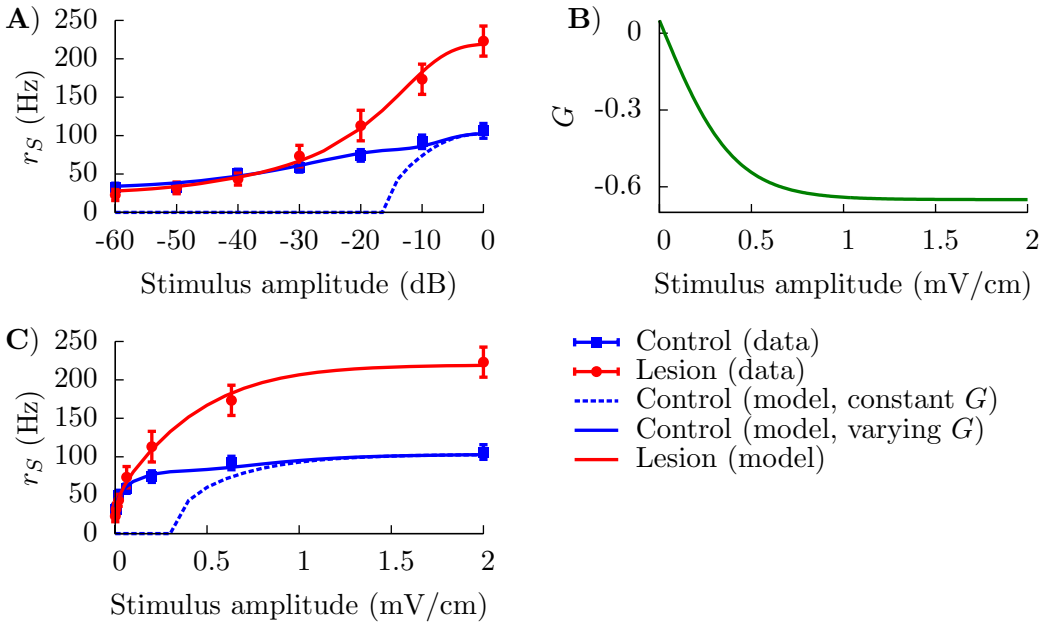

Figure 4.8: Divisive gain control measured in SP cells *in vivo* and the theoretical predictions of our model. (**A**) Absolute firing rates (in contrast to Fig. 4.2A where rate changes relative to baseline were plotted) as a function of the stimulus amplitude in decibels. Experimental absolute rates (symbols) are obtained from the relative rates appearing in Fig. 4.2A by adding the baseline rates for the lesion and control cases (see text for details, especially Eq. 4.19). To simulate the lesion (i.e. absence of parallel fibers) in the model, we set $G = 0$ (red line). The theoretical control curves (solid and dashed blue lines) differ by the assumptions on G. For the dashed blue line, $G = -0.65$ for all amplitude modulations; for the solid one, G is given by the curve illustrated in panel (**B**) ($G(E_{\rm AM}) = G_0 + B/[1 + \exp(-\kappa E_{\rm AM})]$, with $G_0 = 0.75$, $B = -1.4$ and $\kappa = 5$). Panel (**C**) contains the same information as panel (**A**), but with a linear scale for the x axis.

feedback in first approximation. Heuristically, we could write $G = G_{EGp} + G_{nP}$, where G_{EGp} (respectively G_{nP}) is the strength of the feedback coming from the EGp (respectively nP). Then the effective bias would be expressed as $\mu_{eff} = \mu_{0,S} + \gamma(E_{\rm AM}) + \tau_m(G_{EGp} + G_{nP})r_D$, with $\mu_{0,S}$ the AM-independent bias of the SP cell, which is assumed to be different than that of the DP cells. Our approximation consists in setting $\mu_{0,S} + \tau_m G_{nP} r_D \approx \mu_{0,D} \equiv \mu_0$ and identifying G_{EGp} with G.

All parameters other than G are assumed identical for the lesion and control cases. These parameters were determined by inspection of the lesion data and by comparison with other studies of the ELL pyramidal neurons [32, 143, 75, 77]. The time constants are $\tau_m = 6$ ms and $\tau_r = 0.8$ ms. The noise intensity σ is set to 0.01. Such a small value is prescribed by the overall aspect of the lesion firing rate when plotted against a linear scale for the AMs (Fig.

4.8C). To model the input μ, we set $\mu_0 = -0.16$, $A = 2.31$ and $k = 3$ cm/mV. A range of plausible values for k has been established from an inspection of figure 4.6D in [26]. The parameters μ_0 and A are then set in such a way that the theoretical f-I curve for the lesion case adequately fits with the corresponding experimental results. Also, h is set to 0 mV/cm to simplify. From Fig. 4.8A (or C), one sees that the red line (lesion) is a good fit to the experimental curve.

For the control case, making G a constant for the whole range of AMs is incompatible with the experimental results for low AMs. The dotted blue curves in panels A and C of Fig. 4.8 illustrate this point. For instance, fixing $G = -0.65$ makes the theoretical f-I curve agree with the experimental results for large AMs, but it fails to model what happens for smaller AMs. The agreement for large AMs is mainly due to the saturation of the afferents' activity, i.e. μ does not vary appreciably under strong stimulations. Consequently, the divisiveness of the gain control at large amplitudes results from a different mechanism than the one presented up to now in the paper. For small AMs, given the very low noise intensity, a moderately high negative feedback strength naturally yields a subtractive gain control. To get a divisive gain control for small AMs, we have to take into account the variation of G with the stimulus amplitude E_{AM}.

As discussed above, the feedback strength G changes from negative to positive values as E_{AM} decreases from 2 mV to 0 mV. The theoretical firing rate for the control case fits with the experimental one for $E_{\mathrm{AM}} = 2$ mV if we set $G = -0.65$. It is reasonable to assume that G saturates for large positive and negative AM values. A simple function comprising these characteristics is a sigmoid akin to that given in Eq. 4.21, but with a negative '$r_{\max}$'. It is displayed in Fig. 4.8B for positive AMs. The solid blue curves in panels A and C of Fig. 4.8 show a good agreement between theory and experiment when this function is used. We stress, however, that this function has been determined in a purely *ad hoc* way to illustrate that a varying G is necessary. Our simple model is able to explain the divisive gain control obtained *in vivo* for the SP cells as long as the feedback strength is allowed to change with the stimulus amplitude. As stated in section 4.2, the SP cells receive direct excitation and disynaptic inhibition from the parallel fibers. These contributions have different respective weights depending on the stimulus intensity, yielding a variable net feedback strength [122]. Also, SP cells possess receptive fields of the ON center-OFF surround type [133]. Center and surround have different activation thresholds [17], so the proportion of inhibition and excitation depends on the intensity of AMs. Such a differential activation might have an impact on phenomena depending on a neuron's activity, namely the synaptic plasticity.

4.4 Discussion

The modulation of input-output characteristics is a major strategy used by neural systems for the processing and gating of information. Neurons able to change their f-I (firing rate versus input) curve stand as a prominent example of this. While very general mechanisms are known to produce subtractive gain control, mechanisms providing other forms of modulation, such as divisive or non-monotonic gain control (see Fig. 4.1), have been harder to identify [107, 57, 58, 78, 129, 203]. In this work, we presented a biologically plausible mechanism which yields subtractive, divisive and non-monotonic gain control. We analyzed, both with theoretical approaches and numerical simulations, the conditions for which each behavior may appear. The mechanism is based on the interaction between noise and some level of input-driven feedforward inhibition, and constitutes an example of a generic neural strategy which can produce different types of computationally-desirable modulations of f-I curves.

The concrete neural circuit considered in the study is very general: a feedforward neural assembly inhibits target neurons, and both populations receive a common external input. To emphasize this latter fact, the feedforward inhibition is here called 'open-loop feedback inhibition', since both populations strictly lie at the same processing stage in the network. Such a circuit has already been used to study the role of feedforward inhibition in decorrelating neural activity in the somatosensory cortex of rats [145]. In the present work, it is used to model the ELL of the weakly electric fish. The feedforward population is composed of so-called deep pyramidal (DP) neurons, whose targets are neurons that inhibit superficial pyramidal (SP) cells. It is the modulation of the SP cells' firing activity that is studied here. However, the simplicity of the network and neuron models suggests that our results could apply to other systems with a similar structure.

As stated before, the f-I behavior of the neural network depends on the level of neural noise and the strength of the inhibitory open-loop feedback. In most situations, when neurons present some level of stochasticity, the feedback strength will determine whether the gain control is divisive or non-monotonic. This makes the system highly versatile, since the open-loop feedback strength, which can be seen as efficacies of the synaptic connections involved, could be modulated by long-term synaptic plasticity. This might be the case in the ELL of the weakly electric fish, as different types of long-term plasticity have been identified in the parallel fiber synapses with SP neurons in this and other species of electric fish [102, 99]. If that were the case, then as the inhibitory feedback contribution becomes stronger, the divisive gain control observed experimentally in SP cells would turn into a non-monotonic gain control. Indeed, non-monotonic f-I curves for SP neurons have been recently observed *in vitro* under direct parallel fiber stimulation (Fig. 4.3C in [122]), supporting this

hypothesis. Further experimental work, however, is needed to adequately test our prediction of non-monotonic gain control emerging for strong inhibitory feedback.

The gain control mechanism presented here holds for more realistic situations, such as when considering a time-dependent sensory stimulus. In the presence of complex input signals, the presence of temporal delays in the systems becomes an important factor, and we have considered it in our study as well. The introduction of a temporal delay in our inhibitory feedback term revealed a novel effect in the dynamics of the system: as the delay increases, an excitatory input fluctuation arriving to a SP cell from the senses does not coincide in time with its inhibitory counterpart arriving from the feedback pathway. As a consequence, the level of fluctuations arriving at SP cells increases, which in turn increases the SP firing rate, especially near the firing onset. Our theoretical approach is able to explain this effect as well (see Appendix, section 4.5). In spite of these new effects in the dynamics of SP cells, the feedback-driven gain control remains the same as in the original model, indicating that the mechanism can work as well for time-dependent stimuli.

Finally, we found that the basic model used to obtain the aforementioned gain control regimes is able to describe the divisive gain control observed *in vivo* in the SP cells of the ELL [16] for moderate to large stimuli. For low amplitude stimuli, the control has a significant subtractive component. Once the model is augmented with an open-loop feedback strength that depends on stimulus intensity, it exhibits divisive behavior at all amplitudes (i.e. even for low amplitude stimuli). This illustrates the versatility of our model, and shows how it can be modified to depict more subtle gain control phenomena in real systems.

Acknowledgments

This work was funded by the Natural Sciences and Engineering Research Council of Canada (NSERC) and in particular its Discovery Accelerator program, the Canadian Institutes for Health Research (CIHR) and the Fonds de Recherche en Sciences Naturelles et Technologies du Québec (FQRNT).

4.5 Appendix

In this Appendix, we consider the situation where the SP cell receives the external input $\mu + \zeta(t)$ (*cf.* Eq. 4.18), together with the delayed synaptic inputs from the DP cells, $f(t)$. Since $\zeta(t)$ is zero-mean, the SP cell receives on average a signal $\mu + \tau_m G r_D$. The goal here is to determine the intensity of the fluctuations of the total input with respect to that average. These fluctuations play the role of an effective noise which alters the f-I curve of the SP cell. In particular, we determine the effect of the delay on the magnitude of the fluctuations.

We shall compute the auto-covariance function of $f(t) + \zeta(t)$ (we discard the constant μ) defined by

$$\begin{aligned} C_{f+\zeta}(\tau) &\equiv \langle [f(t) + \zeta(t)][f(t+\tau) + \zeta(t+\tau)] \rangle - \tau_m^2 G^2 r_D^2 \\ &= C_f(\tau) + R_{f\zeta}(\tau) + R_{f\zeta}(-\tau) + R_\zeta(\tau), \end{aligned} \tag{4.22}$$

where $C_f(t) \equiv \langle f(t) f(t+\tau) \rangle - \tau_m^2 G^2 r_D^2$ is the auto-covariance function of $f(t)$, $R_{f\zeta}(\tau) \equiv \langle f(t)\zeta(t+\tau) \rangle$ is the cross-correlation function of $f(t)$ and $\zeta(t)$, and $R_\zeta(\tau) \equiv \langle \zeta(t)\zeta(t+\tau) \rangle$ is the auto-correlation function of $\zeta(t)$. The brackets $\langle \cdot \rangle$ signify averages over realizations of the internal (ξ_i^D) and external (ζ) noises. We compute these quantities separately below.

4.5.1 Auto-covariance of $f(t)$

We consider both delta and alpha synapses. We first get the power spectrum of $f(t)$ and then Fourier transform it to get the auto-correlation function, from which we subtract $\tau_m^2 G^2 r_D^2$ to get the auto-covariance function. Consider first the following restriction of all y_i^D to the interval $[0, T]$, i.e.

$$y_{i,T}^D \equiv \begin{cases} y_i^D(t) & \text{for } t \in [0, T] \\ 0 & \text{elsewhere} \end{cases}.$$

We define the Fourier operator, $\mathscr{F}$, as

$$\tilde{x}(\omega) = \mathscr{F}\{x\}(\omega) = \int_{-\infty}^{\infty} x(t) e^{i\omega t} dt$$

and its inverse, $\mathscr{F}^{-1}$, as

$$x(t) = \mathscr{F}^{-1}\{\tilde{x}\}(t) = \int_{-\infty}^{\infty} \tilde{x}(\omega) e^{-i\omega t} \frac{d\omega}{2\pi}.$$

With $f_T(t) \equiv \tau_m \frac{G}{N} \sum_i (s * y_{i,T}^D)(t)$, we have

$$\tilde{f}_T(\omega) = \mathscr{F}\{f_T\}(\omega) = \tau_m \frac{G}{N} \tilde{s}(\omega) \sum_i \tilde{y}_{i,T}^D(\omega). \tag{4.23}$$

The power spectrum of $f(t)$ is defined as

$$S_f(\omega) = \lim_{T\to\infty} \frac{\langle |\tilde{f}_T(\omega)|^2 \rangle}{T}. \tag{4.24}$$

Using Eq. 4.23, we get

$$S_f(\omega) = \tau_m^2 \frac{G^2}{N^2} |\tilde{s}(\omega)|^2 \sum_{i,j} \lim_{T\to\infty} \frac{\langle \tilde{y}_{i,T}^D(\omega) \tilde{y}_{j,T}^{D*}(\omega) \rangle}{T}.$$

Analogously to Eq. 4.24, we define the cross-spectrum of two DP cells as

$$S_D^{\text{cross}}(\omega) = \lim_{T\to\infty} \frac{\langle \tilde{y}_{i,T}^D(\omega) \tilde{y}_{j,T}^{D*}(\omega) \rangle}{T} \qquad (i \neq j). \tag{4.25}$$

Given the homogeneity of the deep cell layer, all DP cells share the same power spectrum, $S_D(\omega)$, and all cross-spectra are the same for any neuron pairs. Hence,

$$S_f(\omega) = \tau_m^2 \frac{G^2}{N} |\tilde{s}(\omega)|^2 [S_D(\omega) + (N-1) S_D^{\text{cross}}(\omega)]. \tag{4.26}$$

One may interpret $|\tilde{s}(\omega)|^2$ as the energy spectrum of the synaptic filter. For $N = 1$, we get the usual output power spectrum of a linear system with impulse response function $\tau_m G s(t)$.

If we assume that DP cells' spike trains are wide-sense stationary random processes, it should be the case for $f(t)$ as well. By the Wiener-Khintchin theorem, the auto-correlation function of $f(t)$ is then

$$\begin{aligned} R_f(\tau) &= \mathscr{F}^{-1}\{S_f\}(\tau) \\ &= \tau_m^2 \frac{G^2}{N} \Big\{ \mathscr{F}^{-1}\{|\tilde{s}|^2\} * [R_D + (N-1) R_D^{\text{cross}}] \Big\}(\tau). \end{aligned}$$

Since

$$\mathscr{F}^{-1}\{|\tilde{s}(\omega)|^2\}(\tau) = (\underline{s} * s)(\tau),$$

where $\underline{s}(t) \equiv s(-t)$, we get

$$R_f(\tau) = \tau_m^2 \frac{G^2}{N}\Big\{(\underline{s} * s) * \left[R_D + (N-1)R_D^{\text{cross}}\right]\Big\}(\tau).$$

On the one hand, for α synapses,

$$\int_{-\infty}^{\infty}(\underline{s} * s)(t)dt = \left(\int_{-\infty}^{\infty} s(t)dt\right)^2 = 1.$$

On the other hand, for δ synapses, $|\tilde{s}(\omega)|^2 = 1$, so that

$$\mathscr{F}^{-1}\{|\tilde{s}(\omega)|^2\}(\tau) = \delta(\tau) \Rightarrow \int_{-\infty}^{\infty}(\underline{s} * s)(t)dt = 1. \tag{4.27}$$

Therefore, we can write

$$C_f(\tau) = \tau_m^2 \frac{G^2}{N}\Big\{(\underline{s} * s) * \left[R_D + (N-1)R_D^{\text{cross}} - Nr_D^2\right]\Big\}(\tau),$$

or

$$C_f(\tau) = \tau_m^2 G^2 \left\{(\underline{s} * s) * \left[\frac{R_D - R_D^{\text{cross}}}{N} + C_D^{\text{cross}}\right]\right\}(\tau). \tag{4.28}$$

For α synapses,

$$(\underline{s} * s)(t) = \frac{1}{4\tau_s}\left[\left(1 + \frac{t}{\tau_s}\right)e^{-\frac{t}{\tau_s}}\Theta(t) + \left(1 - \frac{t}{\tau_s}\right)e^{\frac{t}{\tau_s}}\Theta(-t)\right]. \tag{4.29}$$

For δ synapses, $(\underline{s} * s)(t) = \delta(t)$ from Eq. 4.27, hence

$$C_f(\tau) = \tau_m^2 G^2 \left[\frac{R_D(\tau) - R_D^{\text{cross}}(\tau)}{N} + C_D^{\text{cross}}(\tau)\right]. \tag{4.30}$$

Since $R_D(\tau)$ contains a Dirac delta function, $\delta(\tau)$, there is a white noise component due to the synaptic bombardment in the case of δ synapses. This component is low-pass filtered in the case of α synapses.

4.5.2 Autocorrelation of ζ

Given that ζ is a zero-mean, low-pass filtered Gaussian white noise of standard deviation σ_c, its auto-correlation function is

$$R_\zeta(\tau) = \sigma_c^2 \frac{\sin(\omega_c \tau)}{\omega_c \tau}, \tag{4.31}$$

where $\omega_c = 2\pi f_c$, with f_c being the cutoff frequency.

4.5.3 Cross-correlation function $(R_{f\zeta}(\tau) + R_{f\zeta}(-\tau))$

We now compute $R_{f\zeta}(\tau) + R_{f\zeta}(-\tau)$ in Eq. 4.22. We will conclude that this term only depends on the delay τ_d. We will then obtain an expression for the covariance $R_{f\zeta}(0)$, which represents the intensity of the effective noise caused by the interaction of f and ζ. This noise is responsible for the linearization of the f-I curve observed in Fig. 4.7A.

Since $S_{f\zeta}(-\omega) = S^*_{f\zeta}(\omega)$, we have

$$R_{f\zeta}(\tau) + R_{f\zeta}(-\tau) = 2\mathscr{F}^{-1}\{\mathrm{Re}[S_{f\zeta}(\omega)]\}(\tau), \tag{4.32}$$

where

$$S_{f\zeta}(\omega) = \lim_{T\to\infty} \frac{\langle \tilde{f}_T(\omega)\tilde{\zeta}^*_T(\omega)\rangle}{T}. \tag{4.33}$$

Using Eq. 4.23, we get

$$S_{f\zeta}(\omega) = \tau_m \frac{G}{N}\tilde{s}(\omega) \sum_i \lim_{T\to\infty} \frac{\langle \tilde{y}^D_{i,T}(\omega)\tilde{\zeta}^*_T(\omega)\rangle}{T} = \tau_m G\tilde{s}(\omega) S_{D\zeta}(\omega), \tag{4.34}$$

where again we have used the homogeneity of the DP cells. $S_{D\zeta}(\omega)$ is the cross-spectrum of any DP cell with the external input. Provided that the variance of ζ is not too large, we can write

$$\langle \tilde{y}^D_{i,T}(\omega)\rangle_{|\zeta_T} = \tilde{A}^D_{\mu,\sigma}(\omega)\tilde{\zeta}_T(\omega) + 2\pi r_D \delta(\omega), \tag{4.35}$$

where the notation $\langle x\rangle_{|\zeta_T}$ means the average of x for a given sample function restricted to $[0, T]$ of the process ζ ('frozen noise'), taken over realizations of the internal noise. In the

above equation, r_D is the rate evaluated with Eq. 4.11 without the time-dependent external input $\zeta(t)$. The susceptibility (or transfer function) $\tilde{A}^D_{\mu,\sigma}(\omega)$ depends on μ, σ and τ_r [125]:

$$\tilde{A}^D_{\mu,\sigma}(\omega) = \left(\frac{i\omega r_D\sqrt{2}}{\sigma(i\omega-1)}\right)\left[\frac{\mathcal{D}_{i\omega-1}(\frac{\mu-V_{th}}{\sigma/\sqrt{2}}) - e^{\Delta}\mathcal{D}_{i\omega-1}(\frac{\mu-V_r}{\sigma/\sqrt{2}})}{\mathcal{D}_{i\omega}(\frac{\mu-V_{th}}{\sigma/\sqrt{2}}) - e^{\Delta+i\omega\tau_r}\mathcal{D}_{i\omega}(\frac{\mu-V_r}{\sigma/\sqrt{2}})}\right],$$

with $\Delta = [V_r^2 - V_{th}^2 + 2\mu(V_{th}-V_r)]/2\sigma^2$, and $\mathcal{D}_\alpha(z)$ are parabolic cylinder functions. Multiplying both sides of Eq. 4.35 by $\tilde{\zeta}^*_T(\omega)$, averaging over realizations of ζ_T, dividing by T and taking the limit $T\to\infty$ gives

$$S_{D\zeta}(\omega) = \tilde{A}^D_{\mu,\sigma}(\omega)S_\zeta(\omega), \tag{4.36}$$

where $S_\zeta(\omega)$ is the power spectrum of ζ. Substituting this into Eq. 4.34 and using Eq. 4.32 yields

$$R_{f\zeta}(\tau) + R_{f\zeta}(-\tau) = 2\tau_m G\mathscr{F}^{-1}\{\mathrm{Re}[\tilde{s}(\omega)\tilde{A}^D_{\mu,\sigma}(\omega)S_\zeta(\omega)]\}(\tau). \tag{4.37}$$

Since $S_\zeta(\omega) = (\sigma_c^2/2f_c)[\Theta(\omega+\omega_c) - \Theta(\omega-\omega_c)]$, we obtain

$$R_{f\zeta}(\tau) + R_{f\zeta}(-\tau) = \frac{\tau_m G\sigma_c^2}{f_c}\int_{-\omega_c}^{\omega_c}\mathrm{Re}[\tilde{s}(\omega)\tilde{A}^D_{\mu,\sigma}(\omega)]e^{-i\omega\tau}\frac{d\omega}{2\pi}. \tag{4.38}$$

Of all the contributions to $C_{f+\zeta}(\tau)$, only $R_{f\zeta}(\tau) + R_{f\zeta}(-\tau)$ depends on the delay τ_d (through the Fourier transform of $s(t)$, which makes $e^{i\omega\tau_d}$ appear). In the inset of Fig. 4.7A we display the covariance of f and ζ, $R_{f\zeta}(0)$, for delta synapses. Writing $\tilde{A}^D_{\mu,\sigma}(\omega) = |\tilde{A}^D_{\mu,\sigma}(\omega)|e^{i\phi_A(\omega)}$ and $\tilde{s}(\omega) = e^{i\phi_s}$, with $\phi_s = \omega\tau_d$, we get, from Eq. 4.38,

$$R_{f\zeta}(0) = \frac{\tau_m G\sigma_c^2}{2\omega_c}\int_{-\omega_c}^{\omega_c}|\tilde{A}^D_{\mu,\sigma}(\omega)|\cos[\phi_A(\omega)+\phi_s(\omega)]d\omega. \tag{4.39}$$

The expression above can be calculated with MAPLE (Maplesoft, Waterloo, Canada). A LIF neuron model can be seen as a low-pass filter. Precisely, $|\tilde{A}^D_{\mu,\sigma}(\omega)|$ is an even function which decreases to zero for large ω values. For a small ω_c, $\phi_A \approx \beta\omega$ on $[-\omega_c,\omega_c]$, with β a positive proportionality constant, so that

$$R_{f\zeta}(0) \approx \frac{\tau_m G\sigma_c^2}{2\omega_c}\int_{-\omega_c}^{\omega_c}|\tilde{A}^D_{\mu,\sigma}(\omega)|\cos[(\beta+\tau_d)\omega]d\omega. \tag{4.40}$$

For $\tau_d < -\beta + \pi/2\omega_c$, the first zero-crossing of $\cos[(\beta+\tau_d)\omega]$ lies outside the interval $[-\omega_c,\omega_c]$. When τ_d increases from 0 to $-\beta+\pi/2\omega_c$, $R_{f\zeta}(0)$ increases monotonically with τ_d (because $G<0$). This increase continues until the second zero-crossing of $\cos[(\beta+\tau_d)\omega]$

reaches $\pm\omega_c$. Hence, a maximum occurs at $\tau_d = -\beta + 3\pi/(2\omega_c)$. Beyond that value, $R_{f\zeta}(0)$ oscillates as more and more of the function $\cos[(\beta + \tau_d)\omega]$ gets compressed into $[-\omega_c, \omega_c]$. For very large delays, the integral tends to zero. In summary, this calculation shows that the noise component coming from the delayed interaction of the feedback function $f(t)$ with the time-varying input $\zeta(t)$ increases monotonically for small delays. For larger delays, it starts to show damped oscillations around zero (see inset of Fig. 4.7A).

Chapter 5

Oscillatory-like behavior in feedforward neuronal nets

We demonstrate how rhythmic activity can arise in neural networks from feedforward rather than recurrent circuitry and, in so doing, we provide a mechanism capable of explaining the temporal decorrelation of gamma-band oscillations. We compare the spiking activity of a delayed recurrent network of inhibitory neurons with that of a feedforward network with the same neural properties and axonal delays. Paradoxically, these very different connectivities can yield very similar spike-train statistics in response to correlated input. This happens when neurons are noisy and axonal delays are short. A Taylor expansion of the feedback network's susceptibility—or frequency-dependent gain function—can then be stopped at first order to a good approximation, thus matching the feedforward net's susceptibility. The feedback network is known to display oscillations; these oscillations imply that the spiking activity of the population is felt by all neurons within the network, leading to direct spike correlations in a given neuron. On the other hand, in the output layer of the feedforward net, the interaction between the external drive and the delayed feedforward projection of this drive by the input layer causes indirect spike correlations: spikes fired by a given output layer neuron are correlated only through the activity of the input layer neurons. High noise and short delays partially bridge the gap between these two types of correlation, yielding similar spike-train statistics for both networks. This similarity is even stronger when the delay is distributed, as confirmed by linear response theory.

5.1 Introduction

Rhythms in neural, genetic and other networks can arise from intrinsically oscillating components, feedback connectivity, external input, or a combination of these factors [219, 6]. In the neural context, input characteristics, connectivity and intrinsic dynamics determine the sequence of times at which neurons spike. Diverse statistics may be computed from these times to understand single-neuron and population activity along with their correlations. Such analyses can assess the network's information-processing capabilities [37, 166], the origin of oscillatory activity and the correlation structure of the input [77, 127]. The interplay of intrinsic, input and connectivity properties plays a particularly important role in the context of synchronous and oscillatory activity [221, 217, 220, 100, 90].

There is a general expectation that connectivity strongly determines auto- and cross-correlations and leads to specific neural dynamics [200, 156, 163, 227], even when neurons, synapses and external inputs have identical properties [31, 4]. For example, oscillations are associated with negative (i.e. inhibitory) feedback (Fig. 5.1A). Here we show that, paradoxically, a neuron that is not intrinsically oscillatory can exhibit oscillatory-like dynamics when embedded in a feedforward network (Fig. 5.1B). In the ubiquitous feedforward connectivity, the target 'postsynaptic' neuron receives input directly from an external source as well as via another—usually inhibitory—population, the latter input being slightly delayed. The field of gamma rhythm research is already accustomed to the idea that different topologies can give rise to similar activity, e.g. via the hotly debated PING or ING feedback mechanisms [219, 35]. A special consequence of our analysis is the existence of a third contender involving feedforward rather than feedback circuitry. Delayed correlations are key to this effect, and provide another avenue to explain the genesis and temporal decoherence of noisy neural rhythms [68, 21, 22].

Temporal decorrelation of gamma rhythms is not systematic and strongly depends on the attributes of the stimulation evoking them and the brain region involved [117, 68]. Theoretical explanations of this decoherence have implicated transition to chaos due to increased excitation between neural populations whose respective oscillations are caused by local inhibition [21, 22]. Also, noise—due, for instance, to finite-size effects—can blur the phase transition between oscillatory and non-oscillatory states, leading to damped oscillations in recurrent nets [42]. Below, we propose a drastically different scheme to explain the temporal decorrelation of neural rhythms, by showing that actual oscillations are not necessary to produce neural activity with oscillatory-like features. We emphasize from the onset that the feedforward net *does not* oscillate in the strict sense.

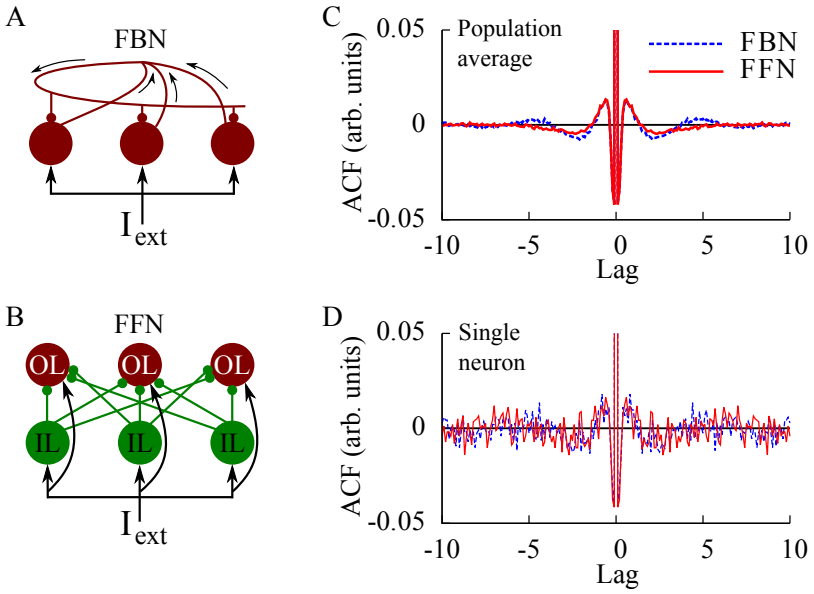

Figure 5.1: (A-B) Schematics of the fully recurrent/feedback net (FBN, A) and of the feed-forward net (FFN; B). For the FFN, the feed-forwarding IL neurons (green discs, bottom of panel B) project only to the OL cells (maroon discs, top). (C-D) Spike-train autocorrelation functions (ACF) averaged over the population (C) and for a single neuron (D). Solid red curves relate to the OL neurons of the FFN, whereas dashed blue curves pertain to neurons of the FBN. Parameters are: $\mu_O = \mu_{FB} = 0.432$, $\mu_I = 0.032$, $\sigma = 0.8$, $\sigma_{\text{ext}} = 0.4$, $\tau_r = 0.1$, $\tau_d = 1$, $\tau_s = 0.5$, $G = -2$. The number of FBN, OL and IL neurons is $N = 100$.

5.2 Methods and model

The spiking activity of a neuron, embodied in its spike train $y(t) = \sum_f \delta(t - t_f)$ with firing times $\{t_f\}$, can be said to oscillate when the occurrence of a spike at time t indicates a greater-than-chance probability to find a spike at time $t + nT$, where T is the period of oscillation and $n \in \mathbb{N}$. This statement follows from the definition of the autocorrelation function (ACF) of spike trains, namely $Q(\tau) = \langle y(t+\tau)y(t)\rangle - r^2$, where brackets denote time average and r is the stationary firing rate of the neuron. For a sufficiently small time interval, r^2 is the probability density that a spike occurs at $t + \tau$ and at t by chance (i.e., solely due to the firing rate). A purely periodic spike train is one for which $Q(\tau) \propto \sum_n \delta(\tau - nT)$. Typically, however, various sources of noise (be it channel noise or synaptic noise) eliminate the occurrence of such idealized spike sequences, and make Q a smooth function of τ which goes to zero in the limit $\tau \to \infty$. In other words, a stochastically oscillating activity is characterized by the presence of a damped oscillation in its ACF (e.g. the dashed blue curve in Fig. 5.1C).

With this criterion in hand, we compared the dynamics of two networks of leaky integrate-and-fire (LIF) inhibitory neurons: a fully recurrent network (Fig. 5.1A; FBN) and a feedforward network (Fig. 5.1B; FFN). The subthreshold dynamics of each neuron obeys $\dot{v} = \mu - v + \sigma\xi + I_{\text{syn}} + I_{\text{ext}}$. A spike is fired whenever a threshold $v_T = 1$ is reached from below; the membrane potential v is then reset to 0 for a refractory time τ_r and integration proceeds thereafter. All neurons in both networks receive the same external input I_{ext}, namely, the same Gaussian white noise (GWN) of intensity σ_{ext}, i.e. the input is perfectly correlated across cells. We also introduce 'internal' individual GWN sources, $\sigma\xi$, which are uncorrelated across neurons and thus act as desynchronizing forces.

A spike fired by a given neuron elicits a synaptic current in all of its targets after an axonal propagation delay τ_d. The total synaptic current impinging on neuron i is

$$I_{\text{syn}} = \sum_k W_{ik}(s * y_k)(t),$$

where $s(t) = \frac{t-\tau_d}{\tau_s^2} \exp\left(-\frac{t-\tau_d}{\tau_s}\right)$ for $t > \tau_d$ and 0 otherwise, y_k is the spike train of presynaptic neuron k, and $*$ denotes convolution. W_{ik} is an element of the connectivity matrix $\mathbf{W}$; it is equal to $G/N < 0$ if neuron k projects to neuron i and 0 otherwise. For the FBN, a neuron's synaptic current comprises contributions from all neurons in the network; for the FFN, $I_{\text{syn}} \equiv 0$ for all input layer (IL) neurons.[1] The output layer (OL) of the FFN receives synaptic inputs from all IL neurons.

The constant bias μ is the same for the recurrent neurons of Fig. 5.1A and for the OL neurons of Fig. 5.1B, but different for the IL neurons. We require all neurons to have the same average firing rate r, in order to allow comparison between networks. The average effective bias of a neuron is $\mu + \langle I_{\text{syn}} \rangle$ for zero-mean internal noise and external input. Since the integral over time of s is unity, $\langle I_{\text{syn}} \rangle = Gr$. For FBN neurons and OL neurons, $\langle I_{\text{syn}} \rangle = Gr$; for IL neurons, $\langle I_{\text{syn}} \rangle = 0$. Therefore, setting $\mu_I = \mu + Gr$ yields, in first approximation[2], identical expected rates r because all other parameters are identical for all neurons. We set μ so that all neurons operate in the subthreshold, noise-activated regime.

[1] Please note that, in this chapter, DP cells are called IL neurons, and SP cells are called OL neurons.

[2] Auto- and cross-interactions between synaptic currents and the external current act as extraneous sources of noise. Since neurons belonging to different populations do not receive the same synaptic current, small discrepancies across the neural populations involved may occur in the firing rates obtained from simulations. However, this does not change our conclusions.

5.3 Results

Recurrent neurons connected through delayed inhibitory synapses and collectively driven by an external noise current have well-known dynamics [125]. A transient excitation from the external current tends to synchronize the neurons, and may cause a good proportion of them to cross threshold and fire. After a delay, they receive their own inhibition, which reduces their probability of spiking for a time interval of approximately $2\tau_s$—the time span of the synaptic function. After this inhibitory current wears off, neurons are back to their original state with higher firing probability, and oscillations ensue (Fig. 5.2A). The strength of this oscillation depends on the proportion of correlated input received by the neurons. Each neuron receives a total effective noise $\sigma_{\text{eff}} = \sqrt{\sigma^2 + \sigma^2_{\text{ext}}}$. The ratio $c \equiv \sigma^2_{\text{ext}}/\sigma^2_{\text{eff}}$ quantifies the proportion of the total noise that is correlated, with c varying between 0 (totally uncorrelated) and 1 (fully correlated). *In vivo*, we typically have $0 \leq c \leq 0.3$ [71]; here, $c = 0.2$. The result is an ostensibly damped collective oscillation of the neural population activity (Fig. 5.1C).

A crucial element of the FFN's dynamics is that both the feed-forwarding layer and its target are driven by the external input (Fig. 5.1B). If at time t the external input provides excitation to all IL cells with such amplitude that most of them fire, then, concomitantly, OL neurons receive the same excitation, and a number of them are likely to fire as well. However, at $t + \tau_d$, OL cells are inhibited by IL cells for an interval of order $2\tau_s$, during which the OL cells' spiking probability is greatly reduced (see for instance in Fig. 5.2B the burst of activity for the IL neurons around $t = 40$ and the resulting large inhibitory current). Hence, the delayed feedforward inhibition mediates anticorrelations in the spike trains of OL neurons. These correlations could be interpreted as *second-order* (or indirect) correlations since the OL spikes are actually correlated through the joint firing of OL and IL cells, which itself is caused by the fluctuations of the external input. Put differently, the cross-covariance between the synaptic and external currents (Fig. 5.2D) shapes the OL neurons' spike trains into burst-like firing epochs. This translates in the ACF into a salient bump at short time lag (due to intraburst spiking), followed by a marked undershoot. These two features of the OL neurons' spike trains match with those of the FBN. However, for the latter, the ACF continues to shallowly oscillate around the time axis past this point.

The Fourier transform of the ACF of a spike train is its power spectral density (PSD). It represents the relative strength of the frequency components of a spike train. As expected from the ACFs, the PSDs of both networks present similar features, namely a distinctive peak. We formalize these similarities by using the linear response theory (LRT) of neural activity [125]. The LRT relates the PSD of neurons whose spontaneous activity—embodied in an "unperturbed" PSD—is perturbed by an external input. The change in activity due to this

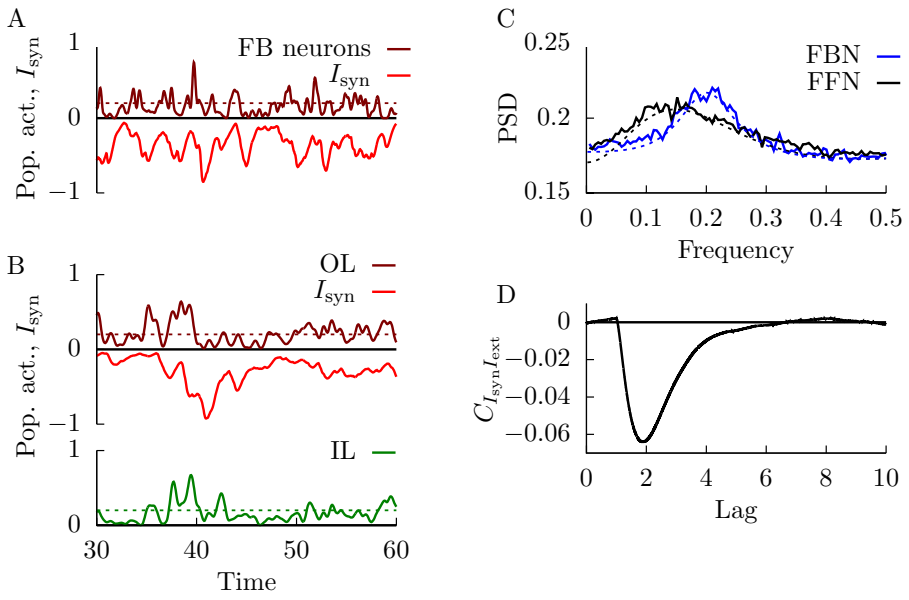

Figure 5.2: (A) Population activity (Pop. act.) for the FBN (top solid maroon curve) for a fraction of the total simulation time, together with the corresponding inhibitory synaptic current I_{syn} (bottom red curve). (B) Same as panel A but for the FFN. We display the population activity of the OL layer (top solid maroon curve) and of the IL layer (bottom solid green curve), together with I_{syn} (middle red curve). Value of the population activity at time t is given by the fraction of neurons with a spike in bin $(t - \Delta t/2, t + \Delta t/2)$ divided by $\Delta t = 0.05$. This population activity was then filtered by a Hann window of unit area and unit length to give the curves displayed above. Dashed curves in panels A and B represent the expected analytical rate of 0.2 spike/unit time. (C) Power spectral densities (PSD) of single-neuron spike trains. For the FFN (leftmost curve), it is the PSD of OL neurons. Linear response theory estimates of the PSDs are given by dashed lines. (D) Cross-covariance between the synaptic current and the external input for the FFN, i.e. $\langle [I_{syn}(t+\tau) - \langle I_{syn}\rangle][I_{ext}(t) - \langle I_{ext}\rangle]\rangle$, as a function of τ. Parameters are as in Fig. 5.1.

external input is quantified by a susceptibility. This susceptibility describes the amplitude of the modulation of the discharge rate of a neuron as a function of the frequency of an applied sinusoidal input current. We used LIF neurons because analytical expressions for their susceptibility and unperturbed PSD are readily available [125]. The PSD of an OL cell in the FFN and of any FBN neuron can be expressed as follows when the number of neurons is large (see Appendix 5.6):

$$S(f) = S^{(0)}(f) + |A_{\text{eff}}(f)|^2 S_{\text{ext}}(f), \tag{5.1}$$

where $S^{(0)}$ is the unperturbed PSD, i.e. the PSD of its spontaneous activity, A_{eff} is the *effective* susceptibility and S_{ext} is the power spectrum of the external input. The effective susceptibility

depends on the connectivity of the circuit. For the FBN, A_{eff} is given by (see Appendix 5.6, Eq. 5.11)

$$A_{\text{eff}}(f) = \frac{A(f)}{1 - Gs(f)A(f)}, \tag{5.2}$$

where A is the susceptibility of a single LIF neuron with intrinsic noise, i.e., neglecting all synaptic input fluctuations, and $s(f)$ is the Fourier transform of the synaptic response function. For the FFN, the effective susceptibility is given by (see Appendix 5.6, Eq. 5.14)

$$A_{\text{eff}}(f) = A(f)[1 + Gs(f)A(f)]. \tag{5.3}$$

Given that all neurons in both the FFN and FBN have the same average input (from all sources), the effective susceptibility for the FFN can be viewed as a first-order Taylor approximation of the effective susceptibility for the FBN [209]. In fact, for M successive layers of inhibitory neurons all receiving a correlated external input, the effective susceptibility of the Mth layer is readily seen to be (*cf.* Appendix 5.6, Eq. 5.15)

$$A_{\text{eff}}(f) = A(f)\left(\sum_{n=0}^{M-1}[Gs(f)A(f)]^n\right). \tag{5.4}$$

In the limit $M \to \infty$, we recover the FBN result in Eq. 5.2, provided that $|GsA| < 1$. That is, we can construct layer by layer a feedforward network whose output layer always possesses more oscillatory-like properties than the penultimate layer; the 'last' layer of an infinite net thus built oscillates. Analogous expansions have already been reported in Refs. [163, 209]. Here, we claim that for high noise and small delays a low-order expansion of Eq. 5.2 is sufficient—i.e., putting this expansion back into Eq. 5.1 yields a PSD in good agreement with the exact FBN's PSD. An obvious corollary is that the FFN is itself a good approximation of the FBN, as seen in Fig. 5.2C.

Coherent oscillatory-like activity prevails at the level of the neural population in the FFN (Fig. 5.3). As depicted in Fig. 5.3B, the PSD of the OL population activity—i.e., the summation of all OL neurons' spike trains, divided by the number of neurons—shows a distinctive peak around $f = 0.15$. Although this peak is of lower amplitude and to the left of the FBN peak, it still represents correlated activity within the output layer of the FFN. Also, synchronous spiking occurs in this layer, as substantiated by the zero-lag peak in the crosscorrelation function (CCF; Fig. 5.3A); the amount of synchronization seems to be identical to that of the FBN. This synchronization of the OL neurons implies that they can intensely modulate the spiking of yet another neural assembly, postsynaptic to the OL neurons (not depicted in Fig. 5.1B). Assuming that this postsynaptic population contains pyramidal cells, one could measure the local field potential (LFP). In the simplest model, the

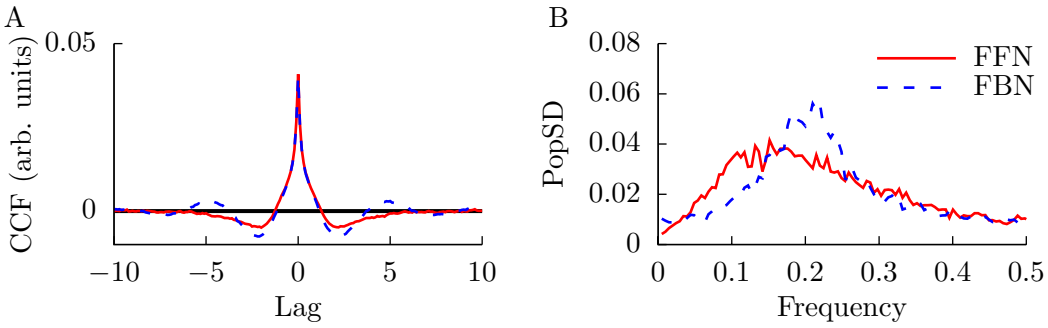

Figure 5.3: Coherent population activity. (A) Cross-correlation functions (CCF) of OL (solid red line) and recurrent neurons (dashed blue line). If y_i and y_j are the (stationary) spike trains of two given neurons i and j, their CCF is $\langle y_i(t)y_j(t+\tau)\rangle - r_i r_j$. Panel A displays the CCF averaged over all pairs of either OL neurons or FBN neurons. (B) PSD of the population activity for the OL and FBN populations. Population activity is defined by $Y(t) = (1/N)\sum_{n=1}^{N} y_n(t)$, where the sum runs over all neurons belonging to the population. Parameters are as in Fig. 5.1.

LFP is proportional to a convolution of the presynaptic (OL) population activity with a filter representing the processing performed by synapses onto pyramidal cells [142]. For synapses acting like low-pass filters, Fig. 5.3B indicates that the PSD of the LFP signal would show a peak which could be interpreted as a neural rhythm; the frequency of this rhythm depends on the chosen value for the membrane time constant (see end of Discussion below).

Although a Gaussian white noise with perfect spatial correlations was used as input in the model, we expect other correlated signals to yield similar results. The similarity between the FBN and the FFN relies on their corresponding effective susceptibilities, which depend on the single-neuron susceptibility, the network connectivity and the synaptic function, but not on the temporal dynamics of the signal. However, since the FBN must oscillate, the input must minimally contain power in a frequency range where the FBN susceptibility can elicit oscillations, and this range depends on $A(f)$ and $s(f)$. On the other hand, spatial correlations are necessary for the IL neurons to coherently shape the spiking activity of OL neurons, but *perfect* correlations are not necessary: we could include into I_{ext} a fraction of the uncorrelated internal noise, view this newly composed signal as the external input and still obtain the same phenomenon as above.

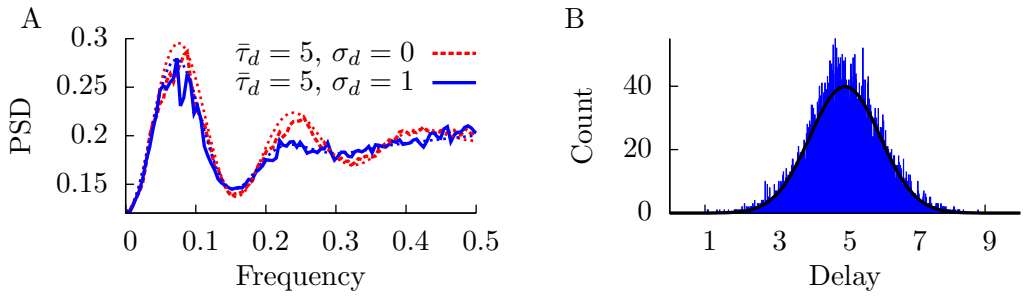

Figure 5.4: Heterogeneous delays smooth the PSD. (A) Power spectral density for uniform (thick dashed red curve) and distributed delays (solid blue curve) together with the corresponding LRT estimates (dotted curves). (B) Gaussian distribution of delays. The histogram contains samples of the parent distribution (solid black curve). Parameters: $\mu_S = 0.85$, $\mu_D = 0.4$, $\sigma = 0.5$; other parameters are as in Fig. 5.1.

5.3.1 Heterogeneous delays

In general, PSDs of oscillating recurrent nets display a dominant peak at the fundamental oscillation frequency and possibly smaller harmonic peaks, depending on how noisy the neurons are. For the FFN, the bursting activity of the output neurons yields a damped oscillation in their PSDs. (In Fig. 5.2C, τ_d is too short for the oscillation to be noticeable.) This damping is mild, so the lowest-frequency peak does not usually stand out with respect to the other satellite bumps (which are not harmonics; Fig. 5.4A, solid red curve). We now ask whether it is possible to make a given FFN's PSD look more like a FBN's PSD by making the dominant peak more salient with respect to the satellite peaks. We found that it is so when axonal delays follow a distribution (in our case a Gaussian, see Fig. 5.4B).

Heterogeneity in the delays was implemented as follows. For any IL-to-OL connection, the delay is given by $\tau_d + x$, where x is sampled from a zero-mean normal distribution with standard deviation σ_d (Fig. 5.4B). Each of the N^2 connections has a different x. Distributed delays imply that even though the IL cells are strongly correlated *via* the external input, the ensuing inhibitory current impinging on the OL cells will lose the precise timing it had for identical delays. Quantitatively, these considerations translate into a reshaping of the OL neurons's PSD (Fig. 5.4A). The distributed delays have a direct impact on the effective susceptibility of the OL neurons in the FFN: instead of Eq. 5.3, we now have (see Appendix 5.6, Eq. 5.19)

$$A_{\text{eff}}(f) = A(f)[1 + Gs(f)A(f)\psi(f)], \tag{5.5}$$

where $\psi(f)$ is the characteristic function of the delay distribution. Thus, heterogeneous delays modulate the effective susceptibility, and through Eq. 5.1, the PSD itself. Figure 5.4A

illustrates that distributed delays considerably flatten the satellite peaks in the PSD; the effect on the dominant peak is less pronounced. For a Gaussian distribution, $\psi(f)$ contains a Gaussian modulation, $\exp(-(2\pi\sigma_d f)^2/2)$, where σ_d is the variance of the delays. Hence, any frequency component greater than $\sim (2\pi\sigma_d)^{-1}$ will be attenuated through Eqs. 5.5 and 5.1. We predict that any delay distribution for which this cutoff frequency lies between the first and second peak in the FFN's PSD will bring it closer to that for the FBN.

5.4 Discussion

In this paper, we showed that feedforward (FFN) and recurrent/feedback (FBN) nets can display similar second-order spike statistics when stimulated by a spatially correlated input. For short time scales—of the order of the axonal delay—the feedforward net presents oscillatory-like properties comparable to those of the feedback net. Using linear response theory, we showed that this is due to the feedforward net being a first-order approximation to the feedback net. This result is more valid for shorter delays; for larger delays the coherence of oscillations in the FBN typically increases [125], so that the mismatch at larger time scale between this network configuration and the FFN is more pronounced. Also, neurons must be noisy for the FBN autocorrelation function (ACF) to decay rapidly.

The existence of this relationship between the recurrent and feedforward nets—under certain conditions—means that one may presumably complement the list of attributes of inhibitory feedforward nets with some of the functions usually ascribed to oscillating recurrent neurons. Feedforward inhibition is already known to sharpen the temporal precision of the target neurons [168, 181], to extend their dynamic range through diverse gain control mechanisms [167], and to keep correlations low between the postsynaptic neurons, helping population coding [145]. On the other hand, recurrent nets of inhibitory neurons enable gamma oscillation generation (a delay-based ING mechanism). Therefore, feedforward nets could also exhibit some of the putative functionalities of these rhythms [49].

Our results are not bound to any rhythm in particular. For a given set of parameters in our model, the oscillation frequency (in SI units) of the recurrent net is determined by the value given to the membrane time constant τ_m; e.g., setting $\tau_m = 5$ ms yields 40-Hz rhythms; slower membrane time constant produces slower rhythms. But since gamma rhythms are known to be transient and low-amplitude [68, 117, 21], our conclusions may be especially applicable to them and may explain their temporal decorrelation.

Acknowledgments

This work was supported by NSERC and CIHR. We thank Richard Naud for useful discussions.

5.5 Appendix A: Firing rates

The analytical formula for the firing rate of a noisy LIF neuron whose subthreshold dynamics reads

$$\frac{dv(t)}{dt} = \mu - v(t) + \sigma\xi(t), \tag{5.6}$$

is given, in our units, by [213]

$$r = \left(\tau_r + \sqrt{\pi}\int_{(\mu-1)/\sigma}^{\mu/\sigma} e^{z^2}\mathrm{erfc}(z)dz\right)^{-1}, \tag{5.7}$$

where erfc is the complementary error function. In our case the subthreshold dynamics also comprises synaptic and external currents. With Gaussian white noise as the external input, a simplified approximation for the rate reads [40, 125]:

$$r = \left(\tau_r + \sqrt{\pi}\int_{(\mu_{\mathrm{eff}}-1)/\sigma_{\mathrm{eff}}}^{\mu_{\mathrm{eff}}/\sigma_{\mathrm{eff}}} e^{z^2}\mathrm{erfc}(z)dz\right)^{-1}, \tag{5.8}$$

where $\mu_{\mathrm{eff}} = \mu + \langle I_{\mathrm{syn}}\rangle$ is an effective bias and $\sigma_{\mathrm{eff}} = \sqrt{\sigma^2 + \sigma_{\mathrm{ext}}^2}$ is the intensity of an effective white noise. As stated in the main text, for a homogeneous population, the stationary rates are all the same so that $\langle I_{\mathrm{syn}}\rangle$ reduces to Gr. For the FBN, Eq. 5.8 must be self-consistently solved for r because it appears on both sides of the equation [40, 125].

5.6 Appendix B: Linear response theory[3]

We now obtain approximate expressions for single-neuron PSDs and cross-spectra for the various cases explored in the main text. We use the so-called linear response theory (LRT) [125, 209]. This theory assumes that the spontaneous activity of the network is dominant, such that external stimulation can be treated as perturbation of that activity. In [209], it is shown that the perturbed spectral density matrix (SDM), $\mathbf{S}(f)$, of a network with connectivity matrix $\mathbf{W}$ is given by

$$\mathbf{S} \approx [\mathbf{1} - \mathbf{A}\mathbf{W}_s]^{-1}(\mathbf{S}^{(0)} + \mathbf{A}\mathbf{S}_{\text{ext}}\mathbf{A}^\dagger)\{[\mathbf{1} - \mathbf{A}\mathbf{W}_s]^{-1}\}^\dagger, \tag{5.9}$$

where the diagonal elements of $\mathbf{S}$ are single-neuron PSDs, whereas its off-diagonal elements are cross-spectral densities. $\mathbf{S}^{(0)}$ is the 'unperturbed' SDM, $\mathbf{A}$ is a diagonal matrix containing susceptibilities, $\mathbf{S}_{\text{ext}}$ is the SDM of the external input, and $\mathbf{W}_s \equiv s\mathbf{W}$, with $s = s(f)$ the Fourier transform of the synaptic function. Dagger (†) denotes Hermitian conjugate, and $\mathbf{1}$ is the identity matrix.

For LIF neurons, analytical expressions for $\mathbf{A}$ and $\mathbf{S}^{(0)}$ are readily available [125], and this is what motivated the use of this particular model for the present study. However, Eq. 5.9 is quite general—as long as the linear approximation remains valid—so other unperturbed PSDs and susceptibilities could be considered as well, such as ones obtained numerically from simulations of more complex models or from experimental estimates. Also, it should be noted that the cross-spectra obtained using the above method may not fit perfectly with the true quantities (i.e. obtained from simulations) [125]. Below and in the main text only approximations for PSDs are considered.

A last note concerns the use of Gaussian white noise as external input. For analytical computations of PSDs using the LRT, white noises are problematic because their variance is infinite. One can circumvent this problem as in [125] by treating the external noise from the same standpoint as the individual noise of each neuron. This means that the susceptibility appearing in terms containing $\mathbf{1} - \mathbf{A}\mathbf{W}_s$ in Eq. 5.9, are to be evaluated using the effective noise $\sigma_{\text{eff}} = \sqrt{\sigma^2 + \sigma_{\text{ext}}^2}$. Also, the term $\mathbf{S}^{(0)} + \mathbf{A}\mathbf{S}_{\text{ext}}\mathbf{A}^\dagger$ in Eq. 5.9 must be replaced by $\mathbf{S}^{(0)}_{\sigma_{\text{eff}}}$, the unperturbed spectral matrix evaluated at σ_{eff}. All firing rates needed in the computations of analytical susceptibilities and unperturbed spectra are obtained from Eq. 5.8.

[3] The reader may skip the following paragraphs and continue reading at section 5.6.2.

5.6.1 Application to a recurrent network of identical inhibitory neurons with uniform synaptic strengths under correlated stimulation

We consider the situation of Fig. 5.1A. The connectivity matrix is $\mathbf{W}_s = (G/N)s\mathbf{U}$, where $\mathbf{U}$ is a matrix whose elements are all 1, G/N is the weight of every synapse and $s = s(f)$ is the Fourier transform of the (inhibitory) synaptic function. Since the network is homogeneous, $\mathbf{A} = A\mathbf{1}$ and $\mathbf{S}^{(0)} = S^{(0)}\mathbf{1}$. The neurons receive a common input $I_{\text{ext}}(t)$, so that $\mathbf{S}_{\text{ext}} = S_{\text{ext}}\mathbf{U}$. Under these conditions,

$$\mathbf{S} \approx [\mathbf{1} - A(G/N)s\mathbf{U}]^{-1} \left(\mathbf{S}^{(0)} + |A|^2 S_{\text{ext}}\mathbf{U}\right) \{[\mathbf{1} - A(G/N)s\mathbf{U}]^{-1}\}^{\dagger}.$$

Define $b(f) \equiv (G/N)A(f)s(f)$. Then

$$\begin{aligned}[\mathbf{1} - b\mathbf{U}]^{-1} &= \mathbf{1} + b\mathbf{U} + b^2\mathbf{U}^2 + b^3\mathbf{U}^3 + \dots \\ &= \mathbf{1} + b\mathbf{U} + b^2N\mathbf{U} + b^3N^2\mathbf{U} + \dots = \mathbf{1} + \frac{b}{1 - bN}\mathbf{U}.\end{aligned}$$

The largest eigenvalue of $b\mathbf{U}$ in absolute value is $|bN|$. Since $f(x) = [1 - x]^{-1} = \sum_{n=0}^{\infty} x^n$ converges for $|x| < 1$, the expansion of $[\mathbf{1} - b\mathbf{U}]^{-1}$ will converge if $|bN| < 1$, i.e. if $|GAs| < 1$. Assuming that this is true, we have, after some developments,

$$\begin{aligned}\mathbf{S} \approx \mathbf{S}^{(0)} + |A|^2 S_{\text{ext}}\mathbf{U} + \frac{2}{N}\,\text{Re}\left\{\frac{AGs}{1 - AGs}\right\} S^{(0)}\mathbf{U} + 2\,\text{Re}\left\{\frac{AGs}{1 - AGs}\right\} |A|^2 S_{\text{ext}}\mathbf{U} \\ + \frac{1}{N}\left|\frac{AGs}{1 - AGs}\right|^2 S^{(0)}\mathbf{U} + \left|\frac{AGs}{1 - AGs}\right|^2 |A|^2 S_{\text{ext}}\mathbf{U}.\end{aligned}$$

After simplification, the PSDs are given by

$$S \approx S^{(0)} + |A|^2 S_{\text{ext}} + \left(|A|^2 S_{\text{ext}} + \frac{S^{(0)}}{N}\right) \frac{2\,\text{Re}\,\{AGs\} - |AGs|^2}{|1 - AGs|^2},$$

which is exactly equation 22 in [125] with $c = 1$ and $F(f) = Gs(f)$. For a large neuron number N, the last expression can be further reduced to

$$S = S^{(0)} + |A_{\text{eff}}|^2 S_{\text{ext}}, \tag{5.10}$$

with the effective susceptibility

$$A_{\text{eff}} = \frac{A}{1 - AGs}. \tag{5.11}$$

5.6.2 Application to the feedforward network

For the FFN (Fig. 5.1B), we have $\mathbf{W}_s = Gs\mathbf{T}_R/N$, where

$$\mathbf{T}_R = \begin{pmatrix} \mathbf{0} & \mathbf{U}_N \\ \mathbf{0} & \mathbf{0} \end{pmatrix}. \tag{5.12}$$

Here, $\mathbf{U}_N$ denotes a matrix of ones of order N connecting IL to OL cells. Using the properties of block matrices, we know that all the eigenvalues of matrix $\mathbf{T}_R$ are 0. Hence, the criterion for the convergence of the Taylor expansion of $[\mathbf{1} - \mathbf{A}\mathbf{W}_s]^{-1}$ is fulfilled for all values of AGs. The correlated stimulus yields $\mathbf{S}_{\text{ext}} = S_{\text{ext}}\mathbf{U}_{2N}$, as before. Since $\mathbf{T}_R^n = \mathbf{0}$ for all $n > 1$, we have $[\mathbf{1} - \mathbf{A}\mathbf{W}_s]^{-1} = \mathbf{1} + \mathbf{A}\mathbf{W}_s = \mathbf{1} + GsA\mathbf{T}_R/N$. Hence,

$$\begin{aligned}\mathbf{S} \approx \mathbf{S}^{(0)} &+ S_{\text{ext}}\mathbf{A}\mathbf{U}_{2N}\mathbf{A}^\dagger + \frac{GsA}{N}(S_I^{(0)}\mathbf{T}_R + S_{\text{ext}}\mathbf{T}_R\mathbf{A}\mathbf{U}_{2N}\mathbf{A}^\dagger) + \frac{Gs^*A^*}{N}S_I^{(0)}\mathbf{T}_R^T \\ &+ S_{\text{ext}}\frac{Gs^*A^*}{N}\mathbf{A}\mathbf{U}_{2N}\mathbf{A}^\dagger\mathbf{T}_R^T + \left|\frac{GsA}{N}\right|^2 S_I^{(0)}\mathbf{T}_R\mathbf{T}_R^T + \left|\frac{GsA}{N}\right|^2 S_{\text{ext}}\mathbf{T}_R\mathbf{A}\mathbf{U}_{2N}\mathbf{A}^\dagger\mathbf{T}_R^T,\end{aligned}$$

where $S_I^{(0)}$ is the unperturbed spectrum of any of the IL cells. We define the matrices:

$$\mathbf{T}_L = \begin{pmatrix} \mathbf{U}_N & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}, \quad \mathbf{B}_R = \begin{pmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{U}_N \end{pmatrix}, \quad \mathbf{B}_L = \begin{pmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{U}_N & \mathbf{0} \end{pmatrix}.$$

With these, the matrix products give:

$$\begin{aligned}\mathbf{A}\mathbf{U}_{2N}\mathbf{A}^\dagger &= |A|^2(\mathbf{T}_L + \mathbf{B}_R + \mathbf{T}_R + \mathbf{B}_L), \\ \mathbf{T}_R\mathbf{A}\mathbf{U}_{2N}\mathbf{A}^\dagger &= N|A|^2(\mathbf{T}_L + \mathbf{T}_R), \\ \mathbf{A}\mathbf{U}_{2N}\mathbf{A}^\dagger\mathbf{T}_R^T &= N|A|^2(\mathbf{T}_L + \mathbf{B}_L), \\ \mathbf{T}_R\mathbf{T}_R^T &= N\mathbf{T}_L, \\ \mathbf{T}_R\mathbf{A}\mathbf{U}_{2N}\mathbf{A}^\dagger\mathbf{T}_R^T &= N^2|A|^2\mathbf{T}_L.\end{aligned}$$

The PSD of any of the postsynaptic neurons (i.e. any OL cell) is read off from the upper left hand block of $\mathbf{S}$ (i.e., terms multiplying $\mathbf{T}_L$ and upper part of $\mathbf{S}^{(0)}$):

$$S \approx S^{(0)} + |A|^2 S_{\text{ext}} + Gs|A|^2 S_{\text{ext}}A + Gs^*|A|^2 S_{\text{ext}}A^* + \frac{|GsA|^2}{N}S_I^{(0)} + |GsA|^2 S_{\text{ext}}$$

or, in the large N limit:

$$\begin{aligned} S &\approx S^{(0)} + |A|^2 S_{\text{ext}}[1 + 2\,\text{Re}\,\{GsA\} + |GsA|^2] \\ &= S^{(0)} + |A_{\text{eff}}|^2 S_{\text{ext}}, \end{aligned} \tag{5.13}$$

with the effective susceptibility

$$A_{\text{eff}} \equiv A(1 + GsA). \tag{5.14}$$

As in the feedback case, the PSD in the large N limit is composed of the unperturbed PSD and of the PSD of the external input, but now filtered by the effective susceptibility of Eq. 5.14. From Eq. 5.13, the filtering process can be thought of being composed of three filters: (1) OL cell filters the external input (giving the term $|A|^2 S_{\text{ext}}$); (2) OL cell filters the synaptic current, which is itself the result of a filtering process (term $|A|^2|Gs|^2|A|^2 S_{\text{ext}}$), and (3) OL cell filters the 'correlation' between the external input and the synaptic current (term $2|A|^2\,\text{Re}\,\{GsA\}S_{\text{ext}}$). The synaptic current is the result of a filtering process because $I_{\text{syn}}(t) = G(s * Y_I)(t)$, i.e., Gs is the filter and the IL population activity $Y_I = (1/N)\sum_{k=1}^{N} y_{I,k}$ is the input to the filter.

5.6.3 Multilayer feedforward network

Suppose now that the feedforward net contains M layers instead of just two. All layers receive the same correlated external input, and all neurons in layer $n-1$ project onto all neurons of layer n. For the second layer, the result of Eq. 5.14 applies, i.e. $A_{\text{eff},2} = A(1 + GsA)$. We remark that $A_{\text{eff},2}$ is just the susceptibility of layer 1 (A) times a modulation $(1 + GsA)$ which contains the filtering done by the previous layer and the synapses, i.e. the term GsA. By extension, then, we would have for layer M that

$$A_{\text{eff},M} = A(1 + GsA_{\text{eff},M-1}).$$

Applying this formula recursively with $A_{\text{eff},2} = A(1 + GsA)$ gives:

$$A_{\text{eff},M} = A\left(\sum_{k=0}^{M-1} (GsA)^k\right). \tag{5.15}$$

The sum is a geometric series. Hence, if $|Gs(f)A(f)| < 1$, $A_{\text{eff},M}$ converges to

$$A_{\text{eff},\infty} = \frac{A}{1 - GsA} \tag{5.16}$$

when $M \to \infty$. This last expression is just the effective susceptibility for the FBN. Therefore, as the number of layers increases, the multilayer FFN becomes equivalent to the FBN.

5.6.4 Heterogeneous delays

We now assume that the propagation delay from IL to OL neurons in the FFN is not the same for all neuron pairs. For a given inhibitory IL neuron k, the average delay is $\tau_{d,k}$, and the delay separating this neuron from any of its N targets is $\tau_{d,k} + x_{n,k}$, where $n = 1, \ldots, N$ and the $x_{n,k}$'s are sampled from a given probability distribution with zero mean and variance σ_d^2 (that could depend on k as well). The synaptic functions $s_{n,k}$ are then

$$s_{n,k}(\omega) = e^{i2\pi f(\tau_{d,k}+x_{n,k})}\alpha(\omega)$$

with

$$\alpha(\omega) = \frac{1}{(1 - i2\pi f\tau_s)^2}$$

the Fourier transform of the un-delayed alpha function, i.e. $\alpha(t) = (t/\tau_s^2)\exp(-t/\tau_s)\Theta(t)$, with $\Theta(t)$ the Heaviside function.

For heterogeneous delays, $\mathbf{W}_s = Gs\mathbf{T}/N$ with

$$\mathbf{T} = \begin{pmatrix} \mathbf{0} & \mathbf{E} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}. \tag{5.17}$$

where $\mathbf{E}$ is a matrix with elements $E_{nk} = e^{i2\pi f(\tau_{d,k}+x_{n,k})}$ (from IL neuron k to OL neuron n). The mathematical developments are similar to those above, with $\mathbf{E}$ replacing $\mathbf{U}_N$ in $\mathbf{W}_s$ (defined just before Eq. 5.12). We readily get the PSD of OL neuron n:

$$S_n = S^{(0)} + |A|^2 S_{\text{ext}}\left(1 + 2\,\text{Re}\left\{G\alpha A\frac{1}{N}\sum_{k=1}^{N} E_{nk}\right\} + \left|\frac{G\alpha A}{N}\right|^2 \sum_k \sum_{k'} E_{nk}E_{nk'}^*\right) + \left|\frac{G\alpha A}{N}\right|^2 S_I^{(0)} \sum_{k=1}^{N} |E_{nk}|^2.$$

A simplifying assumption is to set all average delays between any IL neuron and its N OL

targets equal, i.e. $\tau_{d,k} = \tau_d$. Then, in the large N limit

$$\frac{1}{N}\sum_{k=1}^{N} E_{nk} = \frac{1}{N} e^{i2\pi f \tau_d} \sum_{k=1}^{N} e^{i2\pi f x_{n,k}} \approx \psi_n(f) e^{i2\pi f \tau_d}$$
$$\frac{1}{N^2}\sum_{k}\sum_{k'} E_{nk} E^*_{nk'} \approx |\psi_n(\omega)|^2,$$

where $\psi_n(\omega) = \lim_{N\to\infty}(1/N)\sum_k e^{i2\pi f x_{n,k}}$ is the characteristic function of the random variable $x_{n,k}$. Finally, always in the large N limit,

$$S_n = S^{(0)} + |A_{\text{eff},n}|^2 S_{\text{ext}}, \tag{5.18}$$

with the effective susceptibility

$$A_{\text{eff},n} = A(1 + GsA\psi_n) \tag{5.19}$$

where, as before, $s = \alpha e^{i2\pi f \tau_d}$. Thus, introducing a distribution of delays between each IL neuron and its N target OL neurons amounts to a simple correction factor in the effective susceptibility (compare Eqs. 5.14 and 5.19). In the main text all distribution functions were the same (a Gaussian distribution of delays with standard deviation σ_d), removing the dependency on n.

Discussion and outlook

General discussion

In this thesis, we studied the oscillatory and gain control properties of recurrent and feedforward sensory networks. Using linear response theory and simulations, we found that neurons with opposite ON and OFF responses to inputs are detrimental to the occurrence of oscillations in recurrent nets, even when the external input is correlated across neurons. This is true when the intrinsic neuronal properties of the ON and OFF subpopulations are identical. When the subpopulations are not symmetric it is possible to recover the oscillations, as shown using theory and simulations. In the context of the electrosensory system this could mean that the physiologically characterized ON and OFF neurons should constitute two separate subnetworks. This would be a simple way to preserve gamma rhythms in the ELL in response to global random amplitude modulations of the EOD.

Next, we considered gain control in inhibitory feedforward networks. Various types of gain control have been shown to exist, depending on the neuronal noise level and the strength of inhibition. At low noise, subtractive gain control occurs; at higher noise and moderate synaptic strength, divisive gain control is obtained. Stronger inhibition results in a non-monotonic gain control. We used these results and the known properties of electroreceptor afferents to explain the divisive gain control that have been measured in the ELL of weakly electric fish.

Finally, we found that, unexpectedly, inhibitory feedforward networks can display oscillatorylike properties when subjected to spatially correlated inputs. The delayed interaction between the input received by the output layer (possibly superficial pyramidal cells in the ELL) and a projection of that input *via* the input layer (possibly deep cells) caused burst-like events in the spike trains of the output layer neurons. With comparatively high noise and

short axonal delays, the power spectral density of these latter neurons displayed a salient peak, a known characteristic of oscillating neurons.

We now discuss the significance of these results in a general context, and we formulate a few hypotheses that emerge from these results. We finish with a short discussion of future projects in direct continuation of the work presented herein.

Segregation of ON and OFF pathways

Neurons with ON and OFF response types are found in visual [188] auditory [36] and olfactory [56] systems. In chapter 3 we concluded that the ON and OFF pathways in a fully recurrent network must be segregated to preserve gamma oscillations. Such a pathway splitting has recently been shown to promote a more efficient encoding of stimuli, while also being energetically beneficial [93]. We had to use ON and OFF populations with asymmetric properties to recover oscillations when all neurons were connected together. This asymmetry could still be useful once the pathways are segregated, as has been studied in the visual system [60]. In weakly electric fish the ON and OFF subnetworks cooperate to build a distributed code capable of detecting moving objects in the animal's surrounding [61]. In this context of local signal detection, it has been suggested that gamma oscillations could enhance the directional selectivity, i.e. the preferential response to stimuli moving in a certain direction, of midbrain neurons efferent to ELL SP cells [172]. Hence, weakly electric fish could use gamma rhythms caused by global EOD signals resulting from interactions with conspecifics to efficiently detect preys in their surroundings.

Equivalence between recurrent and feedforward networks

The results of chapters 4 and 5 support the hypothesis that feedforward networks may provide the same computational flexibility as recurrent networks. Rhythms and memory [79] are two phenomena usually pertaining to networks with a high level of interconnections. Oscillations in recurrent nets have been discussed at length in sections 1.3 and 2.7. Oscillatory instabilities develop naturally in these networks, helped by axonal delays. However, in chapter 5 we found that feedforward nets are capable of displaying oscillatorylike properties—without oscillatory instability—when subjected to correlated inputs, a situation that is enhanced by heterogeneities. This echoes the recent proposal of memory without feedback by Mark Goldman [94]. He found that positive feedback loops are not strictly required for generating persistent mnemonic activity, and that purely feedforward nets show comparable performances for this task. Furthermore, we showed in chapter 4 that

gain control in feedforward nets is even more efficient than in a corresponding recurrent net, and that various gain control types are easily obtainable by tuning noise and adjusting the synaptic strength through, e.g., short- or long-term plasticity.

Possible origin of gamma oscillations in the ELL

Is the ELL network feedback or feedforward (refer to section 1.4.4 and tables 1.1 and 1.2 for this discussion)? Our results suggest that a feedforward circuitry could sustain oscillatorylike network behavior. However, single-neuron statistics like autocorrelation functions unequivocally indicate that superficial cells display real oscillatory activity [75]. Pharmalogical blocking of the StF pathway kills oscillations [75], and Pd's stellate cells are quite unresponsive to stimulation above 16 Hz (see section 1.4.4). Hence, Pd's BP cells would cause the observed rhythms when broadband, spatially correlated random stimuli are used. The DP cells' apical dendrites reach the ventral molecular layer of the ELL, so it can get some of the inhibition from the BP cells [18]. We suggest that the loop DP cells→BP→DP cells generate the basic oscillatory pattern (PING-type gamma). If the BP cells are part of the feedback loop, they should oscillate as well and entrain the SP cells. Correlated spatial inputs are required to provide enough inhibition to the DP cells *via* the BP cells. The stellate cells—being entrained by the DP cells—could also participate moderately in the entrainment of the SP cells. To our knowledge however, no direct evidence exists that demonstrates oscillations in DP or BP cells. If experiments show that the DP cells do not oscillate under correlated stimuli, then local synaptic connections in the Pd could be responsible for the oscillation. Moreover, we lack the evidence to completely rule out the feedback loop composed of the SP cells, the midbrain structure *torus semicircularis* and the Pd nucleus. To summarize, ON and OFF DP cells could constitute two segregated pathways capable of generating rhythms in the SP cells by an entrainment of the latter *via* the BP cells. An analogous entrainment has recently been measured in the hippocampus, wherein the oscillating CA3 region inhibits CA1 pyramidal cells [225].

Aspects of the modeling

Some aspects of our modeling are obviously oversimplified when put in the specific context of the weakly electric fish. First, neurons were driven directly by Gaussian white noise, and in so doing we totally ignored the preprocessing performed by electroreceptor afferents. Even at this stage in the processing line complex operations are performed. For instance, the adaptation displayed by these afferents have been shown to help encode the distance of preys,

invariantly of their speed [62]. The afferents' spike-train statistics together with the synaptic filtering at the afferent-pyramidal cell synapses are not expected to produce Gaussian-like postsynaptic currents. In its defence, Gaussian white noise is routinely used for system identification, as it probes the system at all frequencies at once, with equal weights [70]. Second, the response and filtering properties of ELL pyramidal cells are a direct consequence of their morphology (e.g., presence or absence of a basilar dendrite), dorsoventral position (deep, intermediate, superficial) and intrinsic biophysical properties. Simple point neuron models are clearly not sufficient to encompass all of these aspects. In particular, the bursting dynamics of these cells [89, 20, 76] are not taken into account by the simple LIF model that we used. In the end, drawing conclusions from simplified models can be treacherous, but adding system-specific details has the obvious drawback of limiting the generality of the results.

Outlook

Several interesting lines of research stem from our work. One would be to implement the sensory searchlight hypothesis using what is known about the direct feedback pathway in the ELL. This would have applications outside the scope of the electrosensory system, as this hypothesis was originally formulated in the context of the thalamocortical system [67]. Another important issue is the function of gamma rhythms in the ELL *per se*, for instance how it might structure the bursting of ELL pyramidal cells and help prey detection. Burst shaping by descending inputs has already received some attention from the experimental side [20], but the role of gamma oscillations has not yet been addressed properly [140]. The functional aspects of gamma rhythms are a big issue in systems neuroscience [49], so delineating the mechanisms behind a functional role accomplished by them would further promotes these rhythms as a real computational tool and not a mere epiphenomenon.

Appendix A

Numerical methods

In this appendix, I briefly discuss the implementation of the neural network dynamics and the numerical evaluation of power spectra. These numerical methods have been used to obtain most of the simulation results appearing in the thesis.

A.1 Neural network dynamics

After discretizing time, the dynamics of neuron i within its network is (*cf.* Eq. 2.15)

$$v_i(t+\Delta t) = \mu\Delta t + (1-\Delta t)v_i(t) + \sigma\sqrt{\Delta t}\mathcal{N} + I_i^{\text{syn}}(t)\Delta t + I_i^{\text{ext}}(t)\Delta t, \tag{A.1}$$

with $I_i^{\text{syn}}(t) = \sum_{j=1}^{N} W_{ij}(\alpha * y_j)(t)$ and $\mathcal{N}$ a normal random variable with mean 0 and variance 1. This is because the process $\xi(t)dt = dW(t)$ is a Wiener process increment with distribution $p_{dW}(x) = \frac{1}{\sqrt{2\pi dt}}\exp\left(-x^2/(2dt)\right)$, hence it is normally distributed with zero mean and variance dt: $dW(t) \sim \mathcal{N}(0;dt) \sim \sqrt{dt}\mathcal{N}(0;1)$, where $\sim$ here means "distributed like". The nature of I^{ext} depends on the context; in the main text, I mostly use $I_i^{\text{ext}}(t) = \sigma_{\text{ext}}\eta_i(t)$, a GWN process which may be spatially correlated or uncorrelated or something in between.

The most efficient way to compute $I_i^{\text{syn}}(t)$ is to find the ordinary differential equation (ODE) solved by it. An easy way to find the ODE associated with a function is to compute its Laplace transform. For instance, the Laplace transform of the delayed α synaptic function (Eq. 2.6) is

$$\hat{s}(\lambda) = \frac{e^{-\lambda\tau_d}}{(1+\lambda\tau_s)^2}. \tag{A.2}$$

Rearranging this equation gives

$$\lambda^2\tau_s^2\hat{s}(\lambda) + 2\lambda\tau_s\hat{s}(\lambda) + \hat{s}(\lambda) = e^{-\lambda\tau_d}. \tag{A.3}$$

Taking the inverse Laplace transform on both sides of this equation with $s(0) = s'(0) = 0$ yields the ODE

$$\tau_s^2 s''(t) + 2\tau_s s'(t) + s(t) = \delta(t - \tau_d). \tag{A.4}$$

Since this is a linear ODE, $I_i^{\text{syn}}(t)$ with alpha synapses must be a solution of

$$\tau_s^2 u_i''(t) + 2\tau_s u_i'(t) + u_i(t) = \sum_{j=1}^{N}\sum_f W_{ij}\delta(t - t_j^{(f)} - \tau_d), \tag{A.5}$$

with initial conditions $u_i(0) = u_i'(0) = 0$. With definition $z_i = \tau_s u_i' + u_i$ the ODE becomes the first-order system

$$\begin{aligned}\tau_s u_i' &= -u_i + z_i \\ \tau_s z_i' &= -z_i + \sum_{j=1}^{N}\sum_f W_{ij}\delta(t - t_j^{(f)} - \tau_d).\end{aligned}$$

Discretizing time, we obtain

$$\begin{aligned}u_i(t + \Delta t) &= u_i(t)(1 - \frac{\Delta t}{\tau_s}) + z_i(t)\frac{\Delta t}{\tau_s} \\ z_i(t + \Delta t) &= z_i(t)(1 - \frac{\Delta t}{\tau_s}) + \frac{1}{\tau_s}\sum_{j=1}^{N} W_{ij} \sum_{\{f:t_j^{(f)} = t - \tau_d\}} 1.\end{aligned}$$

The equality $t_j^{(f)} = t - \tau_d$—defining the spikes taken into account in the summation—is meaningful because all time quantities are now discretized. A straightforward way to implement this system is to define a vector $\mathbf{s}_j$ of length τ_d for every neuron j, and N doublets (u_i, z_i). The vector $\mathbf{s}_j$ contains the discretized spike train of neuron j, a sequence of zeros and ones; $\mathbf{s}_j[0]$ contains the spike or absence of spike in the last time bin (corresponding to t). With each time increment, the last element of $\mathbf{s}_j$, $\mathbf{s}_j[\tau_d]$, is discarded, and all other elements are shifted to the right. Then,

$$\begin{aligned}u_i(t + \Delta t) &= u_i(t)(1 - a\Delta t) + z_i(t)\Delta t \\ z_i(t + \Delta t) &= z_i(t)(1 - a\Delta t) + a^2\sum_{j=1}^{N} W_{ij}\mathbf{s}_j[\tau_d].\end{aligned}$$

A.2 Power spectra

Computing PSDs of point processes is not a trivial task [112]. I now describe a simple protocol to yield a numerical estimate of a spiking process' PSD. Assume that M realizations[1] of a spiking process $X(t)$ have been obtained. Estimating the PSD of $X(t)$, $S(\omega)$, from numerical simulations involves various sources of errors. First, spike times are given in integer multiples of Δt, the time step of the numerical integration (typically $\sim 10^{-3}$). Second, a single realization has a finite length T, and third, a finite number of realizations is used.

The effect of having a finite recording is the following. Let us define

$$S_T(\omega) \equiv \frac{\langle |\tilde{X}_T(\omega)|^2 \rangle}{T}$$

with $\tilde{X}_T(\omega) = \mathscr{F}\{X_T\}(\omega) = \mathscr{F}\{X(t)\Theta(t)\Theta(T-t)\}(\omega)$, the Fourier transform of the restriction of $X(t)$ to interval $[0, T)$. This is just the definition of the PSD (Eq. 2.27) without the limit $T \to \infty$. $S_T(\omega)$ and $S(\omega)$ are related by [88]

$$S_T(\omega) = (\tilde{\mathcal{B}} * S)(\omega) = \int_{-\infty}^{\infty} \tilde{\mathcal{B}}(\omega - \Omega) S(\Omega) d\Omega, \tag{A.6}$$

where $\tilde{\mathcal{B}}$ is the Fourier transform of

$$\mathcal{B}(t) = \begin{cases} 1 - \frac{|t|}{T}, \text{ for } -T \leq t \leq T \\ 0, \text{ otherwise} \end{cases}, \tag{A.7}$$

i.e.

$$\tilde{\mathcal{B}}(\omega) = \frac{4}{T} \frac{\sin^2(\omega T/2)}{\omega^2} = T \text{sinc}^2(\omega T/2) \tag{A.8}$$

with $\text{sinc}(x) \equiv \sin(x)/x$. Therefore, for any ω, $S_T(\omega)$ is a summation (integration) of nearby values $S(\Omega)$ with weight $\tilde{\mathcal{B}}(\omega - \Omega)$. For larger values of T, $S_T(\omega)$ is less 'polluted' by nearby S_Ω; the width of the central peak of $\tilde{\mathcal{B}}(\omega)$ is $4\pi/T$. Another method—used in conjunction with a large T—is to use a different filter than the box filter $\mathcal{B}(t) = \Theta(t)\Theta(T-t)$. Non-rectangular windows, like Hann[2] or Bartlett windows, are used to reduced the size of the side lobes of the filter's Fourier transform [101].

The next source of error is that, instead of $S_T(\omega) = \frac{\langle |\tilde{X}_T(\omega)|^2 \rangle}{T}$, we use a finite number of

[1] It could be a long realization subdivided into M portions, thus exploiting ergodicity.

[2] Hanning window should actually be called Hann window. 'Hanning' is a verbification of the name Hann, coming from Julius Von Hann. Confusion often arises as well with Hamming windows, named after Richard W. Hamming. Hamming and Hann function are quite similar [101]. Why make things simple?

realizations of length T:

$$\hat{S}_T(\omega) = \frac{1}{MT} \sum_{k=1}^{M} |\tilde{x}_k(\omega)|^2. \tag{A.9}$$

We would like $\hat{S}_T(\omega)$ to be an unbiased estimate of $S_T(\omega)$, i.e. that $\lim_{M\to\infty} \hat{S}_T(\omega) = S_T(\omega)$. Our experiences with estimations of power spectra suggest that this is the case (see also [179], figure 3.8).

Finally, spike times are only obtained at multiples of $\Delta t \sim 10^{-3}$, the integration time step. If $v(t + \Delta t) > v_T$ and $v(t) < v_T$, we say that a spike occurs at $t + \Delta t$. Hence, the numerical spike train is a discrete function of the discretized time; values of the function are either $1/\Delta t$ (as an approximation of a Dirac delta) or 0. A discrete approximation of the Fourier transform of $x_k(t)$ is required. We write

$$\tilde{x}_k(\omega) = \int_0^T x_k(t) e^{i\omega t} dt \approx \Delta t \sum_{n=0}^{N-1} x_k(n\Delta t) e^{i\omega n \Delta t} \tag{A.10}$$

with $T = N\Delta t$. The maximum representable frequency is the Nyquist frequency, $\pi/\Delta t$. The frequency range is then $(-\pi/\Delta t, \pi/\Delta t)$, of width $2\pi/\Delta t$. The Fourier transform consists of N components uniformly distributed on this domain. Hence, the frequency resolution is $\Delta\omega = 2\pi/(N\Delta t) = 2\pi/T$. The final expression for the estimate of the PSD is

$$S_{\text{est}}(k\Delta\omega) = \frac{\Delta t}{MN} \sum_{k=1}^{M} \left| \sum_{n=0}^{N-1} x_k(n\Delta t) e^{i2\pi kn/N} \right|^2. \tag{A.11}$$

The discrete Fourier transform—the expression in the absolute value—can readily and efficiently be computed with the help of the Fast Fourier Transform (FFT) algorithm (see for instance [170]).

A.2.1 Aliasing and spike trains

The approximation of Eq. A.10 can be viewed as the result of two subsequent operations applied to the spike train realizations $x_k(t) = \sum_f \delta(t - t^{(f)})$, $0 < t^{(f)} < T$ [14]. First, every x_k is convolved with a normalized boxcar filter, $b(t) = (1/D)\Theta(t + \frac{D}{2})\Theta(\frac{D}{2} - t)$, where D is the window's width; the integral of $b(t)$ over $\mathbb{R}$ is 1. The result, $(b * x_k)(t)$, is then multiplied by the sampling function $f_s(t) \equiv D \sum_{n=-\infty}^{\infty} \delta(t - nD)$, which picks values of $(b * x_k)(t)$ at sampling interval D. The sampling must be at interval D because a larger interval could

miss spikes and a smaller interval could overcount spikes. We recover Eq. A.10 if $D = \Delta t$:

$$\begin{aligned}\mathscr{F}\{f_s(t)(b * x_k)(t)\}(\omega) &= \int_0^T f_s(t)(b * x_k)(t)e^{i\omega t}dt \\ &= D\sum_{n=-\infty}^{\infty}\int_0^T \delta(t-nD)\sum_f b(t-t^{(f)})e^{i\omega t}dt \\ &= D\sum_{\{n|nD\in[0,T]\}}\sum_f b(nD-t^{(f)})e^{i\omega nD}\end{aligned}$$

and identifying $\sum_f b(nD - t^{(f)})$ with the sequence $x_k(n\Delta t)$ used above. From the theory of sampled band-limited causal signals[3], choosing $D < \pi/\omega_M$ prevents the aliasing of Fourier components, where ω_M is the bandwidth of the signal [109]. That is, the sampling rate is fast enough for all frequencies contained in the signal to be fully discriminated (Nyquist frequency > ω_M). So, in our case, $\mathscr{F}\{f_s(t)(b * x_k)(t)\}(m\Delta\omega)$ would correspond to $\mathscr{F}\{(b * x_k)(t)\}(m\Delta\omega)$—note the absence of f_s in the brackets—provided that $(b * x_k)(t)$ is band-limited. However, since $|\mathscr{F}\{(b * x_k)(t)\}(\omega)| = |\tilde{b}(\omega)||\tilde{x}_k(\omega)|$, with $\tilde{b}(\omega) = \mathrm{sinc}(\omega D/2)$ and $\tilde{x}_k(\omega) = \sum_f e^{i\omega t^{(f)}}$, we have

$$|\tilde{b}(\omega)||\tilde{x}_k(\omega)| = |\mathrm{sinc}(\omega D/2)|\sqrt{N_k + 2\sum_{f>f'}\cos(\omega(t^{(f)} - t^{(f')}))}, \tag{A.12}$$

with N_k the spike count for realization k. Hence, $(b * x_k)(t)$ is not band-limited. The Nyquist frequency π/D is half the value of the first zero of $\tilde{b}(\omega)$, namely $2\pi/D$. However, for $D \sim 10^{-2}$, say, $\pi/D \sim 100\pi$. Since an adequate range for ω is $[-5\pi, 5\pi]$ (see Fig. 2.5), aliasing is not such a troublesome issue.

More sophisticated methods, like the multitaper [88] and Welch methods, are also possible. The function pwelch in Matlab™ (The MathWorks, Inc., Natick, Massachusetts) is useful for this purpose.

[3] A causal signal $x(t)$ has $x(t) = 0$ if $t < 0$. $x(t)$ is band-limited if $|\tilde{x}(\omega)| = 0$ for $|\omega| > \omega_M$, where ω_M is the bandwidth [109].

References

[1] L. F. Abbott, J. A. Varela, K. Sen, and S. B. Nelson. Synaptic depression and cortical gain control. *Science*, 275(5297):220–224, 1997.

[2] M. Abeles. *Corticonics: Neural circuits of the cerebral cortex*. Cambridge University Press, Cambridge, 1991.

[3] M. Abramovich and I. A. Stegun. *Handbook of Mathematical functions*. New York, Dover, 1968.

[4] R. Albert and A.-L. Barabási. Statistical mechanics of complex networks. *Rev. Mod. Phys.*, 74:47–97, 2002.

[5] H. J. Alitto and W. M. Usrey. Influence of contrast on orientation and temporal frequency tuning in ferret primary visual cortex. *J. Neurophysiol.*, 91(6):2797–2808, 2004.

[6] U. Alon. Network motifs: theory and experimental approaches. *Nat. Rev. Genet.*, 8(6):450–461, 2007.

[7] C. A. Anastassiou, R. Perin, H. Markram, and C. Koch. Ephaptic coupling of cortical neurons. *Nat. Neurosci.*, 14(2):217–223, 2011.

[8] R. A. Andersen and V. B. Mountcastle. The influence of the angle of gaze upon the excitability of the light-sensitive neurons of the posterior parietal cortex. *J. Neurosci.*, 3(3):532–548, 1983.

[9] O. Avila Akerberg and M. J. Chacron. Noise shaping in neural populations. *Phys. Rev. E*, 79:011914, 2009.

[10] O. Avila-Akerberg and M. J. Chacron. Noise shaping in neural populations with global delayed feedback. *Math. Model. Nat. Phenom.*, 5(2):100–124, 2010.

[11] A. Ayaz and F. S. Chance. Gain modulation of neuronal responses by subtractive and divisive mechanisms of inhibition. *J. Neurophysiol.*, 101(2):958–968, 2009.

[12] D. Babineau, A. Longtin, and J. E. Lewis. Modeling the electric field of weakly electric fish. *J. Exp. Biol.*, 209:3636–3651, 2006.

[13] M. Bacher. A new method for the simulation of electric fields, generated by electric fish, and their distorsions by objects. *Biol. Cybern.*, 47(1):51–58, 1983.

[14] W. Bair, C. Koch, W. Newsome, and K. Britten. Power spectrum analysis of bursting cells in area mt in the behaving monkey. *J. Neurosci.*, 14(5):2870–2892, 1994.

[15] J. Bastian. Gain control in the electrosensory system: a role for the descending projections to the electrosensory lateral line lobe. *J. Comp. Physiol. A*, 158(4):505–515, 1986.

[16] J. Bastian. Gain control in the electrosensory system mediated by descending inputs to the electrosensory lateral line lobe. *J. Neurosci.*, 6(2):553–562, 1986.

[17] J. Bastian, M. J. Chacron, and L. Maler. Receptive field organization determines pyramidal cell stimulus-encoding capability and spatial stimulus selectivity. *J. Neurosci.*, 22(11):4577–4590, 2002.

[18] J. Bastian, M. J. Chacron, and L. Maler. Plastic and nonplastic pyramidal cells perform unique roles in a network capable of adaptive redundancy reduction. *Neuron*, 41(5):767–779, 2004.

[19] J. Bastian and J. Courtright. Morphological correlates of pyramidal cell adaptation rate in the electrosensory lateral line lobe of weakly electric fish. *J. Comp. Physiol. A*, 168(4):393–407, 1991.

[20] J. Bastian and J. Nguyenkim. Dendritic modulation of burst-like firing in sensory neurons. *J. Neurophysiol.*, 85(1):10–22, 2001.

[21] D. Battaglia, N. Brunel, and D. Hansel. Temporal decorrelation of collective oscillations in neural networks with local inhibition and long-range excitation. *Phys. Rev. Lett.*, 99:238106, 2007.

[22] D. Battaglia and D. Hansel. Synchronous chaos and broad band gamma rhythm in a minimal multi-layer model of primary visual cortex. *PLoS Comput. Biol.*, 7(10):e1002176, 2011.

[23] M. Bazhenov and I. Timofeev. Thalamocortical oscillations. *Scholarpedia*, 1(6):1319, 2006.

[24] C. C. Bell and L. Maler. Central neuroanatomy of electrosensory systems in fish. In R. R. F. Theodore Holmes Bullock, Carl D. Hopkins, editor, *Electroreception*, pages 68–111. Springer, New York, 2005.

[25] J. Benda and A. V. Herz. A universal model for spike-frequency adaptation. *Neural Comput.*, 15(11):2523–2564, 2003.

[26] J. Benda, A. Longtin, and L. Maler. Spike-frequency adaptation separates transient communication signals from background oscillations. *J. Neurosci.*, 25(9):2312–2321, 2005.

[27] H. Berger. Über das elektrenkephalogramm des menschen. *European Archives of Psychiatry and Clinical Neuroscience*, 87(1):527–570, 1929.

[28] N. J. Berman and L. Maler. Interaction of gabab-mediated inhibition with voltage-gated currents of pyramidal cells: computational mechanism of a sensory searchlight. *J. Neurophysiol.*, 80(6):3197–3213, 1998.

[29] N. J. Berman and L. Maler. Neural architecture of the electrosensory lateral line lobe: adaptations for coincidence detection, a sensory searchlight and frequency-dependent adaptive filtering. *J. Exp. Biol.*, 202(10):1243–1253, 1999.

[30] N. J. Berman, J. Plant, R. W. Turner, and L. Maler. Excitatory amino acid receptors at a feedback pathway in the electrosensory system: implications for the searchlight hypothesis. *J. Neurophysiol.*, 78(4):1869–1881, 1997.

[31] S. Boccaletti, V. Latora, Y. Moreno, M. Chavez, and D.-U. Hwang. Complex networks: Structure and dynamics. *Physics Reports*, 424(4–5):175 – 308, 2006.

[32] K. Bol, G. Marsat, E. Harvey-Girard, A. Longtin, and L. Maler. Frequency-tuned cerebellar channels and burst-induced ltd lead to the cancellation of redundant sensory inputs. *J. Neurosci.*, 31(30):11028–11038, 2011.

[33] K. Bol, G. Marsat, J. F. Mejías, L. Maler, and A. Longtin. Modeling cancelation of periodic inputs with burst-stdp and feedback. *Neural Networks*, 47:120–133, 2013.

[34] C. Börgers, S. Epstein, and N. J. Kopell. Background gamma rhythmicity and attention in cortical local circuits: a computational study. *Proc. Natl. Acad. Sci. U.S.A.*, 102(19):7002–7007, 2005.

[35] C. Börgers and N. Kopell. Synchronization in networks of excitatory and inhibitory neurons with sparse, random connectivity. *Neural Comput.*, 15(3):509–538, 2003.

[36] C. Borges-Merjane and L. O. Trussell. On and off unipolar brush cells transform multisensory inputs to the auditory system. *Neuron*, 85(5):1029–1042, 2015.

[37] A. Borst and F. E. Theunissen. Information theory and neural coding. *Nat. Neurosci.*, 2:947–957, 1999.

[38] B. Bratton and J. Bastian. Descending control of electroreception. ii. properties of nucleus praeeminentialis neurons projecting directly to the electrosensory lateral line lobe. *J. Neurosci.*, 10(4):1241–1253, 1990.

[39] R. Brette and W. Gerstner. Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *J. Neurophysiol.*, 94(5):3637–3642, 2005.

[40] N. Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *J. Comput. Neurosci.*, 8(3):183–208, 2000.

[41] N. Brunel, F. S. Chance, N. Fourcaud, and L. Abbott. Effects of synaptic noise and filtering on the frequency response of spiking neurons. *Phys. Rev. Lett.*, 86(10):2186–2189, 2001.

[42] N. Brunel and V. Hakim. Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural Comput.*, 11(7):1621–1671, 1999.

[43] N. Brunel and D. Hansel. How noise affects the synchronization properties of recurrent networks of inhibitory neurons. *Neural Comput.*, 18(5):1066–1110, 2006.

[44] 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. *J. Neurophysiol.*, 90(1):415–430, 2003.

[45] G. Buzsàki. *Rhythms of the Brain*. Oxford University Press, Oxford, 2006.

[46] G. Buzsáki, C. A. Anastassiou, and C. Koch. The origin of extracellular fields and currents —eeg, ecog, lfp and spikes. *Nat. Rev. Neurosci.*, 13(6):407–420, 2012.

[47] G. Buzsàki and A. Draguhn. Neuronal oscillations in cortical networks. *Science*, 304:1926–1929, 2004.

[48] G. Buzsáki, L. Lai-Wo S., and C. H. Vanderwolf. Cellular bases of hippocampal eeg in the behaving rat. *Brain Research Reviews*, 6(2):139–171, 1983.

[49] G. Buzsàki and X.-J. Wang. Mechanisms of gamma oscillations. *Annu. Rev. Neurosci.*, 35:203–225, 2012.

[50] R. M. Capocelli and L. M. Ricciardi. Diffusion approximation and first passage time problem for a model neuron. *Kybernetik*, 8(6):214–223, 1971.

[51] M. Carandini and D. J. Heeger. Summation and division by neurons in primate visual cortex. *Science*, 264(5163):1333–1336, 1994.

[52] J. A. Cardin, M. Carlen, K. Meletis, U. Knoblich, F. Zhang, K. Deisseroth, L.-H. Tsai, and C. I. Moore. Driving fast-spiking cells induces gamma rhythm and controls sensory responses. *Nature*, 459(7247):663–667, 2009.

[53] C. Carr and M. Konishi. A circuit for detection of interaural time differences in the brain stem of the barn owl. *J. Neurosci.*, 10(10):3227–3246, 1990.

[54] M. J. Chacron, B. Doiron, L. Maler, A. Longtin, and J. Bastian. Non-classical receptive field mediates switch in a sensory neuron's frequency tuning. *Nature*, 423(6935):77–81, 2003.

[55] M. J. Chacron, A. Longtin, and L. Maler. Delayed excitatory and inhibitory feedback shape neural information transmission. *Phys. Rev. E*, 72(5):051917, 2005.

[56] S. H. Chalasani, N. Chronis, M. Tsunozaki, J. M. Gray, D. Ramot, M. B. Goodman, and C. I. Bargmann. Dissecting a circuit for olfactory behaviour in caenorhabditis elegans. *Nature*, 450(7166):63–70, 2007.

[57] F. S. Chance and L. Abbott. Divisive inhibition in recurrent networks. *Network-Comp. Neural.*, 11(2):119–129, 2000.

[58] F. S. Chance, L. F. Abbott, and A. D. Reyes. Gain modulation from background synaptic input. *Neuron*, 35(4):773–782, 2002.

[59] L. Chen, J. L. House, R. Krahe, and M. E. Nelson. Modeling signal and background components of electrosensory scenes. *J. Comp. Physiol. A*, 191(4):331–345, 2005.

[60] E. Chichilnisky and R. S. Kalmar. Functional asymmetries in on and off ganglion cells of primate retina. *J. Neurosci.*, 22(7):2737–2747, 2002.

[61] S. E. Clarke, A. Longtin, and L. Maler. A neural code for looming and receding motion is distributed over a population of electrosensory on and off contrast cells. *J. Neurosci.*, 34(16):5583–5594, 2014.

[62] S. E. Clarke, R. Naud, A. Longtin, and L. Maler. Speed-invariant encoding of looming object distance requires power law spike rate adaptation. *Proc. Natl. Acad. Sci. U.S.A.*, 110(33):13624–13629, 2013.

[63] J. C. Coggshall. Linear models for biological transducers and impulse train spectra: General formulation and review. *Kybernetik*, 13:30-37, 1973.

[64] C. E. Connor, J. L. Gallant, D. C. Preddie, and D. C. Van Essen. Responses in area V4 depend on the spatial relationship between stimulus and attention. *J. Neurophysiol.*, 75(3):1306–1308, 1996.

[65] C. E. Connor, D. C. Preddie, J. L. Gallant, and D. C. Van Essen. Spatial attention effects in macaque area V4. *J. Neurosci.*, 17(9):3201–3214, 1997.

[66] A. K. Corsi, B. Wightman, and M. Chalfie. A transparent window into biology: A primer on caenorhabditis elegans. *Genetics*, 200(2):387–407, 2015.

[67] F. Crick. Function of the thalamic reticular complex: the searchlight hypothesis. *Proc. Natl. Acad. Sci. U.S.A.*, 81(14):4586–4590, 1984.

[68] J. Csicsvari, B. Jamieson, K. D. Wise, and G. Buzsáki. Mechanisms of gamma oscillations in the hippocampus of the behaving rat. *Neuron*, 37(2):311 – 322, 2003.

[69] D. J. Daley and D. Vere-Jones. *An Introduction to the Theory of Point Processes*. Springer-Verlag, New York, 1998.

[70] P. Dayan and L. F. Abbott. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. The MIT Press, Cambridge, MA, 2001.

[71] J. de la Rocha, B. Doiron, E. Shea-Brown, K. Josić, and A. Reyes. Correlation between neural spike trains increases with firing rate. *Nature*, 448(7155):802–807, 2001.

[72] J. de la Rocha, R. Moreno, and N. Parga. Correlations modulate the non-monotonic response of a neuron with short-term plasticity. *Neurocomputing*, 58:313–319, 2004.

[73] J. de la Rocha and N. Parga. Short-term synaptic depression causes a non-monotonic response to correlated stimuli. *J. Neurosci.*, 25(37):8416–8431, 2005.

[74] A. A. Disney, C. Aoki, and M. J. Hawken. Gain modulation by nicotine in macaque V1. *Neuron*, 56(4):701–713, 2007.

[75] B. Doiron, M. J. Chacron, L. Maler, A. Longtin, and J. Bastian. Inhibitory feedback required for network oscillatory responses to communication but not prey stimuli. *Nature*, 421(6922):539–543, 2003.

[76] B. Doiron, C. Laing, A. Longtin, and L. Maler. Ghostbursting: a novel neuronal burst mechanism. *J. Comput. Neurosci.*, 12(1):5–25, 2002.

[77] B. Doiron, B. Lindner, A. Longtin, L. Maler, and J. Bastian. Oscillatory activity in electrosensory neurons increases with the spatial correlation of the stochastic input stimulus. *Phys. Rev. Lett.*, 93(4):048101, 2004.

[78] B. Doiron, A. Longtin, N. Berman, and L. Maler. Subtractive and divisive inhibition: effect of voltage-dependent inhibitory conductances and noise. *Neural Comput.*, 13(1):227–248, 2001.

[79] D. Durstewitz, J. K. Seamans, and T. J. Sejnowski. Neurocomputational models of working memory. *Nat. Neurosci.*, 3:1184–1191, 2000.

[80] L. D. Ellis, W. H. Mehaffey, E. Harvey-Girard, R. W. Turner, L. Maler, and R. J. Dunn. Sk channels provide a novel mechanism for the control of frequency tuning in electrosensory neurons. *J. Neurosci.*, 27(35):9491–9502, 2007.

[81] B. Ermentrout. Type I membranes, phase resetting curves, and synchrony. *Neural Comput.*, 8(5):979–1001, 1996.

[82] G. Ermentrout and N. Kopell. Oscillator death in systems of coupled neural oscillators. *SIAM J. Appl. Math.*, 50(1):125–146, 1990.

[83] G. B. Ermentrout and D. H. Terman. *Mathematical Foundations of Neuroscience*. Springer, New York, 2010.

[84] A. A. Faisal, L. P. Selen, and D. M. Wolpert. Noise in the nervous system. *Nat. Rev. Neurosci.*, 9(4):292–303, 2008.

[85] J.-M. Fellous, M. Rudolph, A. Destexhe, and T. J. Sejnowski. Synaptic background noise controls the input/output characteristics of single cells in an in vitro model of in vivo activity. *Neuroscience*, 122(3):811–829, 2003.

[86] D. Ferster and K. D. Miller. Neural mechanisms of orientation selectivity in the visual cortex. *Annu. Rev. Neurosci.*, 23(1):441–471, 2000.

[87] N. Fourcaud and N. Brunel. Dynamics of the firing probability of noisy integrate-and-fire neurons. *Neural Comput.*, 14(9):2057–2110, 2002.

[88] F. Gabbiani, S. J. Cox, and S. J. Cox. *Mathematics for Neuroscientists*. Academic Press, London, 2010.

[89] F. Gabbiani, W. Metzner, R. Wessel, and C. Koch. From stimulus encoding to feature extraction in weakly electric fish. *Nature*, 384(6609):564–567, 1996.

[90] R. F. Galán, N. Fourcaud-Trocmé, G. B. Ermentrout, and N. N. Urban. Correlation-induced synchronization of oscillations in olfactory bulb neurons. *J. Neurosci.*, 26(14):3646–3655, 2006.

[91] W. Gerstner and W. Kistler. *Spiking Neuron Models: An Introduction*. Cambridge University Press, New York, NY, USA, 2002.

[92] W. Gerstner, W. M. Kistler, R. Naud, and L. Paninski. *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press, Cambridge, 2014.

[93] J. Gjorgjieva, H. Sompolinsky, and M. Meister. Benefits of pathway splitting in sensory coding. *J. Neurosci.*, 34(36):12127–12144, 2014.

[94] M. S. Goldman. Memory without feedback in a neural network. *Neuron*, 61(4):621 – 634, 2009.

[95] D. Golomb and D. Hansel. The number of synaptic inputs and the synchrony of large, sparse neuronal networks. *Neural Comput.*, 12(5):1095–1139, 2000.

[96] C. M. Gray, P. Konig, A. K. Engel, and W. Singer. Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, 338(6213):334–337, 1989.

[97] C. M. Gray and W. Singer. Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proc. Natl. Acad. Sci. U.S.A.*, 86(5):1698–1702, 1989.

[98] D. Gussin, J. Benda, and L. Maler. Limits of linear rate coding of dynamic stimuli by electroreceptor afferents. *J. Neurophysiol.*, 97(4):2917–2929, 2007.

[99] V. Z. Han, K. Grant, and C. C. Bell. Reversible associative depression and nonassociative potentiation at a parallel fiber synapse. *Neuron*, 27(3):611–622, 2000.

[100] D. Hansel, G. Mato, and C. Meunier. Synchrony in excitatory neural networks. *Neural Comput.*, 7(2):307–337, 1995.

[101] F. Harris. On the use of windows for harmonic analysis with the discrete fourier transform. *Proceedings of the IEEE*, 66(1):51–83, 1978.

[102] E. Harvey-Girard, J. Lewis, and L. Maler. Burst-induced anti-hebbian depression acts through short-term synaptic dynamics to cancel redundant sensory signals. *J. Neurosci.*, 30(17):6152–6169, 2010.

[103] 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. *J. Neurosci.*, 26(34):8787–8799, 2006.

[104] K. J. Hildebrandt, J. Benda, and R. M. Hennig. Multiple arithmetic operations in a single neuron: the recruitment of adaptation processes in the cricket auditory pathway depends on sensory context. *J. Neurosci.*, 31(40):14142–14150, 2011.

[105] N. Hô and A. Destexhe. Synaptic background activity enhances the responsiveness of neocortical pyramidal neurons. *J. Neurophysiol.*, 84(3):1488–1496, 2000.

[106] A. L. Hodgkin and A. F. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol.*, 117(4):500–544, 1952.

[107] G. R. Holt and C. Koch. Shunting inhibition does not have a divisive effect on firing rates*. *Neural Comput.*, 9(5):1001–1013, 1997.

[108] S. Hong, B. N. Lundstrom, and A. L. Fairhall. Intrinsic gain modulation and adaptive neural coding. *PLoS Comput. Biol.*, 4(7):e1000119, 2008.

[109] H. Hsu. *Signals and Systems*. Schaum's outlines. McGraw-Hill, New York, 2011.

[110] D. H. Hubel and T. N. Wiesel. Receptive fields of single neurones in the cat's striate cortex. *J. Physiol.*, 148(3):574–591, 1959.

[111] D. Johnson. Point process models of single-neuron discharges. *J. Comput. Neurosci.*, 3(4):275–299, 1996.

[112] R. E. Kass, E. N. Brown, and U. T. Eden. *Analysis of Neural Data*. Springer Series in Statistics. Springer-Verlag, New York, 2014.

[113] B. W. Knight. Dynamics of encoding in a population of neurons. *J. Gen. Physiol.*, 59(6):734–766, 1972.

[114] R. Krahe, J. Bastian, and M. J. Chacron. Temporal processing across multiple topographic maps in the electrosensory system. *J. Neurophysiol.*, 100(2):852–867, 2008.

[115] R. Krahe and L. Maler. Neural maps in the electrosensory system of weakly electric fish. *Curr. Opin. Neurobiol.*, 24(1):13–21, 2014.

[116] G. Kreiman, R. Krahe, W. Metzner, C. Koch, and F. Gabbiani. Robustness and variability of neuronal coding by amplitude-sensitive afferents in the weakly electric fish *Eigenmannia*. *J. Neurophysiol.*, 84(1):189–204, 2000.

[117] A. Kreiter and W. Singer. Stimulus-dependent synchronization of neuronal responses in the visual cortex of the awake macaque monkey. *J. Neurosci.*, 16(7):2381–2396, 1996.

[118] M. J. Lannoo, L. Maler, and B. Tinner. Ganglion cell arrangement and axonal trajectories in the anterior lateral line nerve of the weakly electric fish Apteronotus leptorhynchus (gymnotiformes). *J. Comp. Neurol.*, 280(3):331–342, 1989.

[119] P. Latham, B. Richmond, P. Nelson, and S. Nirenberg. Intrinsic dynamics in neuronal networks. i. theory. *J. Neurophysiol.*, 83(2):808–827, 2000.

[120] J. Lefebvre, A. Longtin, and V. G. LeBlanc. Dynamics of driven recurrent networks of on and off cells. *Phys. Rev. E*, 80(4):041912, 2009.

[121] J. Lefebvre, A. Longtin, and V. G. LeBlanc. Responses of recurrent nets of asymmetric on and off cells. *J. Biol. Phys.*, 37(2):189–212, 2011.

[122] J. E. Lewis, B. Lindner, B. Laliberté, and S. Groothuis. Control of neuronal firing by dynamic parallel fiber feedback: implications for electrosensory reafference suppression. *J. Exp. Biol.*, 210(24):4437–4447, 2007.

[123] B. Lindner. *Coherence and Stochastic Resonance in Nonlinear Dynamical Systems*. PhD thesis, Humboldt-Universität zu Berlin, 2002.

[124] B. Lindner, M. J. Chacron, and A. Longtin. Integrate-and-fire neurons with threshold noise: A tractable model of how interspike interval correlations affect neuronal signal transmission. *Phys. Rev. E*, 72:021911, 2005.

[125] B. Lindner, B. Doiron, and A. Longtin. Theory of oscillatory firing induced by spatially correlated noise and delayed inhibitory feedback. *Phys. Rev. E*, 72(6):061919, 2005.

[126] B. Lindner, L. Schimansky-Geier, and A. Longtin. Maximizing spike train coherence or incoherence in the leaky integrate-and-fire model. *Phys. Rev. E*, 66(3):031916, 2002.

[127] A. Litwin-Kumar, M. J. Chacron, and B. Doiron. The spatial structure of stimuli shapes the timescale of correlations in population spiking activity. *PLoS Comput. Biol.*, 8(9):e1002667, 2012.

[128] A. Longtin. Neuronal noise. *Scholarpedia*, 8(9):1618, 2013.

[129] A. Longtin, B. Doiron, and A. R. Bulsara. Noise-induced divisive gain control in neuron models. *Biosystems*, 67(1):147–156, 2002.

[130] C. Ly and B. Doiron. Divisive gain modulation with dynamic stimuli in integrate-and-fire neurons. *PLoS Comput. Biol.*, 5(4):e1000365, 2009.

[131] Z. F. Mainen and T. J. Sejnowski. Reliability of spike timing in neocortical neurons. *Science*, 268(5216):1503–1506, 1995.

[132] L. Maler. The posterior lateral line lobe of certain gymnotoid fish: Quantitative light microscopy. *J. Neur. Comp.*, 183:323–364, 1979.

[133] L. Maler. Neural strategies for optimal processing of sensory signals. *Prog. Brain. Res.*, 165:135–154, 2007.

[134] L. Maler. Receptive field organization across multiple electrosensory maps. i. columnar organization and estimation of receptive field size. *J. Comp. Neurol.*, 516(5):376–393, 2009.

[135] L. Maler, E. Sas, S. Johnston, and W. Ellis. An atlas of the brain of the electric fish Apteronotus leptorhynchus. *J. Chem. Neuroanat.*, 4(1):1–38, 1991.

[136] L. Maler, E. K. B. Sas, and J. Rogers. The cytology of the posterior lateral line lobe of high-frequency weakly electric fish (gymnotidae): Dendritic differentiation and synaptic specificity in a simple cortex. *J. Comp. Neurol.*, 195(1):87–139, 1981.

[137] H. A. Mallot. *Computational Neuroscience: A First Course*, volume 2 of *Springer Series in Bio-/Neuroinformatics*. Springer Cham, 2013.

[138] A. Manwani and C. Koch. Detecting and estimating signals in noisy cable structures, I: Neuronal noise sources. *Neural Comput.*, 11(8):1797–1829, 1999.

[139] D. Marinazzo, H. J. Kappen, and S. C. A. M. Gielen. Input-driven oscillations in networks with excitatory and inhibitory neurons with dynamic synapses. *Neural Comput.*, 19(7):1739–1765, 2007.

[140] N. Masuda and B. Doiron. Gamma oscillations of spiking neural populations enhance signal discrimination. *PLoS Comput. Biol.*, 3(11):e236, 2007.

[141] M. Mattia and P. Del Giudice. Population dynamics of interacting spiking neurons. *Phys. Rev. E*, 66(5):051917, 2002.

[142] A. Mazzoni, S. Panzeri, N. K. Logothetis, and N. Brunel. Encoding of naturalistic stimuli by local field potential spectra in networks of excitatory and inhibitory neurons. *PLoS Comput. Biol.*, 4(12):e1000239, 2008.

[143] J. F. Mejias, G. Marsat, K. Bol, L. Maler, and A. Longtin. Learning contrast-invariant cancellation of redundant signals in neural systems. *PLoS Comput. Biol.*, 9(9):e1003180, 2013.

[144] J. F. Mejias, A. Payeur, E. Selin, L. Maler, and A. Longtin. Subtractive, divisive and non-monotonic gain control in feedforward nets linearized by noise and delays. *Front. Comput. Neurosci.* , 8(19), 2014.

[145] J. W. Middleton, C. Omar, B. Doiron, and D. J. Simons. Neural correlation is stimulus modulated by feedforward inhibitory circuitry. *J. Neurosci.*, 32(2):506–518, 2012.

[146] K. J. Miller, L. B. Sorensen, J. G. Ojemann, and M. den Nijs. Power-law scaling in the brain surface electric potential. *PLoS Comput. Biol.*, 5(12):e1000609, 2009.

[147] R. E. Mirollo and S. H. Strogatz. Synchronization of pulse-coupled biological oscillators. *SIAM J. Appl. Math.*, 50(6):1645–1662, 1990.

[148] S. J. Mitchell and R. A. Silver. Shunting inhibition modulates neuronal gain during synaptic excitation. *Neuron*, 38(3):433–445, 2003.

[149] B. K. Murphy and K. D. Miller. Multiplicative gain changes are induced by excitation or inhibition alone. *J. Neurosci.*, 23(31):10040–10051, 2003.

[150] R. Naud and W. Gerstner. Coding and decoding with adapting neurons: A population approach to the peri-stimulus time histogram. *PLoS Comput. Biol.*, 8(10):e1002711, 2012.

[151] R. Naud, N. Marcille, C. Clopath, and W. Gerstner. Firing patterns in the adaptive exponential integrate-and-fire model. *Biol. Cybern.*, 99(4-5):335–347, 2008.

[152] J. O'Keefe and M. L. Recce. Phase relationship between hippocampal place units and the eeg theta rhythm. *Hippocampus*, 3(3):317–330, 1993.

[153] B. Oksendal. *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, Heidelberg, 2013.

[154] F. W. J. Olver, D. W. Lozier, R. F. Boisvert, and C. W. Clark, editors. *NIST Handbook of Mathematical Functions*. Cambridge University Press, New York, NY, 2010.

[155] S. Ostojic. Interspike interval distributions of spiking neurons driven by fluctuating inputs. *J. Neurophysiol.*, 106(1):361–373, 2011.

[156] S. Ostojic, N. Brunel, and V. Hakim. How connectivity, background activity, and synaptic properties shape the cross-correlation between spike trains. *J. Neurosci.*, 29(33):10234–10253, 2009.

[157] E. Parzen. *Stochastic processes*. SIAM, Philadelphia, USA, 1999.

[158] A. X. Patel and D. Burdakov. Mechanisms of gain control by voltage-gated channels in intrinsically-firing neurons. *PLOS ONE*, 10(3):e0115431, 2015.

[159] A. Payeur, J. Lefebvre, L. Maler, and A. Longtin. Linear response theory for two neural populations applied to gamma oscillation generation. *Phys. Rev. E*, 87:032703, 2013.

[160] A. Payeur, L. Maler, and A. Longtin. Oscillatorylike behavior in feedforward neuronal networks. *Phys. Rev. E*, 92:012703, 2015.

[161] M. Penttonen, A. Kamondi, L. Acsády, and G. Buzsáki. Gamma frequency oscillation in the hippocampus of the rat: intracellular analysis in vivo. *Eur. J. Neurosci.*, 10(2):718–728, 1998.

[162] D. H. Perkel, J. H. Schulman, T. H. Bullock, G. P. Moore, and J. P. Segundo. Pacemaker neurons: effects of regularly spaced synaptic input. *Science*, 145(3627):61–63, 1964.

[163] V. Pernice, B. Staude, S. Cardanobile, and S. Rotter. How structure determines correlations in neuronal networks. *PLoS Comput. Biol.*, 7(5):e1002059, 2011.

[164] J. W. Pillow, J. Shlens, L. Paninski, A. Sher, A. M. Litke, E. J. Chichilnisky, and E. P. Simoncelli. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*, 454(7207):995–999, 2008.

[165] P.-O. Polack, J. Friedman, and P. Golshani. Cellular mechanisms of brain state-dependent gain modulation in visual cortex. *Nat. Neurosci.*, 16(9):1331–1339, 2013.

[166] A. Pouget, P. Dayan, and R. Zemel. Information processing with population codes. *Nat. Rev. Neurosci.*, 1(2):125–132, 2000.

[167] F. Pouille, A. Marin-Burgin, H. Adesnik, B. V. Atallah, and M. Scanziani. Input normalization by global feedforward inhibition expands cortical dynamic range. *Nat Neurosci*, 12(12):1577–1585, 2009.

[168] F. Pouille and M. Scanziani. Enforcement of temporal fidelity in pyramidal cells by somatic feed-forward inhibition. *Science*, 293(5532):1159–1163, 2001.

[169] S. A. Prescott and Y. De Koninck. Gain control of firing rate by shunting inhibition: roles of synaptic noise and dendritic saturation. *Proc. Natl. Acad. Sci. U.S.A.*, 100(4):2076–2081, 2003.

[170] W. H. Press, S. A. Teulkosky, W. T. Vetterling, and B. P. Flannery. *Numerical recipes in C++: the art of scientific computing*. Cambridge University Press, Cambridge, 2002.

[171] N. J. Priebe and D. Ferster. A new mechanism for neuronal gain control (or how the gain in brains has mainly been explained). *Neuron*, 35(4):602604, 2004.

[172] J. U. Ramcharitar, E. W. Tan, and E. S. Fortune. Global electrosensory oscillations enhance directional responses of midbrain neurons in eigenmannia. *J. Neurophysiol.*, 96(5):2319–2326, 2006.

[173] B. Rasnow. The effects of simple objects on the electric field of apteronotus. *J. Comp. Physiol. A*, 178(3):397–411, 1996.

[174] B. Rasnow, C. Assad, M. E. Nelson, and J. M. Bower. Simulation and measurement of the electric fields generated by weakly electric fish. In D. S. Touretzky, editor, *Advances in neural information processing systems*, pages 436–443, San Mateo, CA, 1989.

[175] 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. *J. Neurophysiol.*, 90(3):1598–1612, 2003.

[176] A. Renart, N. Brunel, and X.-J. Wang. *Computational Neuroscience: A Comprehensive Approach*, chapter Mean-field theory of irregularly spiking neuronal populations and working memory in recurrent cortical networks. Boca Raton, CRC Press, 2004.

[177] A. Reyes. Influence of dendritic conductances on the input-output properties of neurons. *Annu. Rev. Neurosci.*, 24(1):653–675, 2001.

[178] L. M. Ricciardi. *Diffusion processes and related topics in biology*. Springer-Verlag, Berlin, 1977.

[179] F. Rieke, D. Warland, R. de Ruyter van Steveninck, and W. Bialek. *Spikes: Exploring the Neural Code*. MIT Press, Cambridge, MA, USA, 1997.

[180] H. Risken. *Fokker-Planck Equation*. Springer-Verlag, Berlin, 1984.

[181] M. T. Roberts, S. C. Seeman, and N. L. Golding. A mechanistic understanding of the role of feedforward inhibition in the mammalian sound localization circuitry. *Neuron*, 78(5):923–935, 2013.

[182] J. S. Rothman, L. Cathala, V. Steuber, and R. A. Silver. Synaptic depression enables neuronal gain control. *Nature*, 457(7232):1015–1018, 2009.

[183] E. Salinas and L. Abbott. A model of multiplicative neural responses in parietal cortex. *Proc. Natl. Acad. Sci. U.S.A.*, 93(21):11956–11961, 1996.

[184] E. Salinas and T. Sejnowski. Gain modulation in the central nervous system: Where behavior, neurophysiology, and computation meet. *Neuroscientist*, 7:430–440, 2001.

[185] E. Salinas and T. J. Sejnowski. Impact of correlated synaptic input on output firing rate and variability in simple neuronal models. *J. Neurosci.*, 20(16):6193–6209, 2000.

[186] E. Salinas and P. Thier. Gain modulation: a major computational principle of the central nervous system. *Neuron*, 27(1):15–21, 2000.

[187] J. Saunders and J. Bastian. The physiology and morphology of two types of electrosensory neurons in the weakly electric fishApteronotus leptorhynchus. *J. Comp. Physiol. A*, 154(2):199–209, 1983.

[188] P. H. Schiller. The on and off channels of the visual system. *Trends Neurosci.*, 15(3):86–92, 1992.

[189] J. Segundo. Communication and coding by nerve cells. *The neurosciences, second study program. Rockefeller University Press, New York*, pages 569–586, 1970.

[190] E. Serrano, T. Nowotny, R. Levi, B. H. Smith, and R. Huerta. Gain control network conditions in early sensory coding. *PLoS Comput. Biol*, 9:e1003133, 2013.

[191] E. Shea-Brown, K. Josić, J. de la Rocha, and B. Doiron. Correlation and synchrony transfer in integrate-and-fire neurons: Basic properties and consequences for coding. *Phys. Rev. Lett.*, 100:108102, 2008.

[192] O. Shriki, D. Hansel, and H. Sompolinsky. Rate models for conductance-based cortical neuronal networks. *Neural Comput.*, 15(8):1809–1841, 2003.

[193] C. A. Shumway. Multiple electrosensory maps in the medulla of weakly electric gymnotiform fish. I. Physiological differences. *J. Neurosci.*, 9(12):4388–4399, 1989.

[194] C. A. Shumway. Multiple electrosensory maps in the medulla of weakly electric gymnotiform fish. II. Anatomical differences. *J. Neurosci.*, 9(12):4400–4415, 1989.

[195] A. J. Siegert. On the first passage time probability problem. *Phys. Rev.*, 81(4):617–623, 1951.

[196] A. M. Sillito, H. E. Jones, G. L. Gerstein, and D. C. West. Feature-linked synchronization of thalamic relay cell firing induced by feedback from the visual cortex. *Nature*, 369(6480):479–482, 1994.

[197] W. R. Softky and C. Koch. The highly irregular firing of cortical cells is inconsistent with temporal integration of random epsps. *J. Neurosci.*, 13(1):334–350, 1993.

[198] V. S. Sohal, F. Zhang, O. Yizhar, and K. Deisseroth. Parvalbumin neurons and gamma rhythms enhance cortical circuit performance. *Nature*, 459(7247):698–702, 2009.

[199] H. Spekreijse and H. Oosting. Linearizing: A method for analysing and synthesizing nonlinear systems. *Kybernetik*, 7(1):22–31, 1970.

[200] O. Sporns, G. Tononi, and G. Edelman. Connectivity and complexity: the relationship between neuroanatomy and brain dynamics. *Neural Networks*, 13(8–9):909–922, 2000.

[201] M. Stemmler. A single spike suffices: the simplest form of stochastic resonance in model neurons. *Network*, 7(4):687–716, 1996.

[202] K. M. Stiefel, B. Englitz, and T. J. Sejnowski. Origin of intrinsic irregular firing in cortical interneurons. *Proc. Natl. Acad. Sci. U.S.A.*, 110(19):7886–7891, 2013.

[203] C. Sutherland, B. Doiron, and A. Longtin. Feedback-induced gain control in stochastic spiking networks. *Biol. Cybern.*, 100(6):475–489, 2009.

[204] K. Thurley, W. Senn, and H.-R. Lüscher. Dopamine increases the gain of the input-output response of rat prefrontal pyramidal neurons. *J. Neurophysiol.*, 99(6):2985–2997, 2008.

[205] P. Tiesinga and J. V. José. Robust gamma oscillations in networks of inhibitory hippocampal interneurons. *Network*, 11(1):1–23, 2000.

[206] P. Tiesinga and T. J. Sejnowski. Cortical enlightenment: are attentional gamma oscillations driven by ING or PING? *Neuron*, 63(6):727–732, 2009.

[207] R. D. Traub, A. Bibbig, F. E. N. LeBeau, E. H. Buhl, and M. A. Whittington. Cellular mechanisms of neuronal population oscillations in the hippocampus in vitro. *Annu. Rev. Neurosci.*, 27:247–278, 2004.

[208] Y. Trotter and S. Celebrini. Gaze direction controls response gain in primary visual-cortex neurons. *Nature*, 398(6724):239–242, 1999.

[209] J. Trousdale, Y. Hu, E. Shea-Brown, and K. Josić. Impact of network structure and cellular response on spike time correlations. *PLoS Comput. Biol.*, 8(3):e1002408, 2012.

[210] 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. *J. Neurophysiol.*, 93(2):1074–1089, 2005.

[211] M. V. Tsodyks and H. Markram. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proc. Natl. Acad. Sci. U.S.A.*, 94(2):719–723, 1997.

[212] H. C. Tuckwell. *Introduction to Theoretical Neurobiology Vol. 1: Linear Cable Theory and Dendritic Structure*. Cambridge University Press, Cambridge, 1988.

[213] H. C. Tuckwell. *Introduction to Theoretical Neurobiology Vol. 2: Nonlinear and stochastic theories*. Cambridge University Press, Cambridge, 1989.

[214] R. W. Turner, J. R. Plant, and L. Maler. Oscillatory and burst discharge across electrosensory topographic maps. *J. Neurophysiol.*, 76(4):2364–2382, 1996.

[215] G. E. Uhlenbeck and L. S. Ornstein. On the theory of the brownian motion. *Phys. Rev.*, 36(5):823–841, 1930.

[216] N. Ulanovsky and C. F. Moss. What the bat's voice tells the bat's brain. *Proc. Natl. Acad. Sci. U.S.A.*, 105(25):8491–8498, 2008.

[217] C. Van Vreeswijk, L. Abbott, and G. Bard Ermentrout. When inhibition not excitation synchronizes neural firing. *J. Comput. Neurosci.*, 1(4):313–321, 1994.

[218] L. M. W. Hamish Mehaffey and R. W. Turner. Intrinsic frequency tuning in ell pyramidal cells varies across electrosensory maps. *J. Neurophysiol.*, 99(5):2641–2655, 2008.

[219] X.-J. Wang. Neurophysiological and computational principles of cortical rhythms in cognition. *Physiol. Rev.*, 90(3):1195–1268, 2010.

[220] X.-J. Wang and G. Buzsáki. Gamma oscillation by synaptic inhibition in a hippocampal interneuronal network model. *J. Neurosci.*, 16(20):6402–6413, 1996.

[221] X.-J. Wang and J. Rinzel. Alternating and synchronous rhythms in reciprocally inhibitory model neurons. *Neural Comput.*, 4(1):84–97, 1992.

[222] J. A. White, J. T. Rubinstein, and A. R. Kay. Channel noise in neurons. *Trends Neurosci.*, 23(3):131–137, 2000.

[223] M. A. Whittington, R. D. Traub, and J. G. Jefferys. Synchronized oscillations in interneuron networks driven by metabotropic glutamate receptor activation. *Nature*, 373:612–615, 1995.

[224] H. R. Wilson and J. D. Cowan. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophys. J.*, 12(1):1–24, 1972.

[225] R. Zemankovics, J. M. Veres, I. Oren, and N. Hájos. Feedforward inhibition underlies the propagation of cholinergically induced gamma oscillations from hippocampal ca3 to ca1. *J. Neurosci.*, 33(30):12337–12351, 2013.

[226] Z.-w. Zhang and D. Arsenault. Gain modulation by serotonin in pyramidal neurones of the rat prefrontal cortex. *J. Physiol.*, 566(2):379–394, 2005.

[227] L. Zhao, B. n. Beverlin, T. Netoff, and D. Q. Nykamp. Synchronization from second order network connectivity statistics. *Front. Comput. Neurosci.*, 5:28, 2011.

[228] Z. Zhou, M. R. Bernard, and A. Bonds. Deconstruction of spatial integrity in visual stimulus detected by modulation of synchronized activity in cat visual cortex. *J. Neurosci.*, 28(14):3759–3768, 2008.