

ALMA MATER STUDIORUM
UNIVERSITÀ DEGLI STUDI DI BOLOGNA

Dottorato di Ricerca in Fisica
Ciclo XXVII

**Bayesian Computations
in Noisy
Spiking Neurons**

Settore Concorsuale di Afferenza: 02/D1
Settore Scientifico Disciplinare: FIS/07
Coordinatore: Prof. Gastone Castellani

Presentata da:
Alessandro Ticchi

Relatore:
Prof. Roberto Soldati

Esame Finale Anno Accademico 2015-2016

Abstract

Understanding an information processing system requires three distinct, complementary levels of analysis. First, we need to understand what the system is doing, and why is it doing it (computational level). Moreover, we need to comprehend what representations and processes are used to achieve the desired goal (algorithmic level) and, finally, we need to discern how these computations are physically realized (implementation level). Guided by this approach, here we investigate brain's functioning and behavior starting from its fundamental constituents (neurons) and from the molecular processes underlying neural dynamics (e.g. synaptic vesicle release, ions flow and action potentials).

The world is stochastic and chaotic, and organisms have access to limited information to take decisions and to interact with the environment. For this reason, brains are continuously required to deal with probability distributions, and experimental evidence confirms that they are dealing with these distributions optimally or close to optimally, according to the rules of Bayesian probability theory. Yet, a complete understanding of how these computations are implemented at the neural level is still missing. We assume that the "computational" goal of neurons is to perform Bayesian inference and to represent the state of the world efficiently. Starting from this assumption, we derive from first Bayesian principles two distinct models of neural functioning, one in single neuron and one in neural populations, and we show how known biophysics and molecular processes of neurons can be explained by the proposed models.

Neural processes and functioning are intrinsically stochastic and noisy. Molecular sources of variability have historically been considered as a problem for information processing in the brain, which affects reliability and efficiency of neural functioning. Yet, in the models we propose noise plays an original constructive role. In fact, molecular variability is exploited by the system as a precious resource to get in correspondence with the stochastic properties of the external world, and to actively implement inference and learning.

One common problem in mathematical models of biological systems is fine tuning of model parameters. In the models we propose, the parameters controlling neural functioning depend on the statistics and the dynamical properties of the encoded stimuli. In both models, we propose a mechanism that allows to learn these parameters locally from the experienced sensory input, and to adapt them on-line to changes in the stimulus properties.

The models we propose suggest a new original interpretation for various biophysical quantities and processes. Action potentials, which are usually considered the paramount form of communication between neurons, in our model of single neuron dynamics are reinterpreted as an internal communication channel, which uses an efficient (predictive) coding scheme to synchronize and coordinate local calcium concentration in different regions of the neuron. Following this new interpretation, intracellular calcium concentration is interpreted as the most explicit representation of the external world inside the neuron. Specifically, we propose that calcium level represents the log-odds probability ratio of a particular hidden state in the world. Furthermore, we reinterpret synaptic vesicle release as a sampling process, which simulates the external world given all the available information (encoded in local calcium concentration). Finally, the neural population dynamics we propose interpret spontaneous neural activity as a process of sampling from the prior world statistics. This interpretation is not entirely new in the literature, but in our model we show how the generated samples can be used to implement a Markov Chain Monte Carlo algorithm that produces inference by sampling.

The proposed models generate various observable predictions, which match experimental results about synaptic vesicle release, short-term synaptic potentiation, ions channels open probability, intracellular calcium dynamics and propagation, spike rate adaptation and neural receptive fields.

Acknowledgements

Almost ten years ago I became part of the University of Bologna as a Bachelor student in Physics. I had no idea of where I was going. Now, ten years later, I still have no idea of where this journey will lead me in the end but, turning back, I can see how I have been truly lucky. I met people that supported me, step by step, and guided me on an unknown path, which was full of surprises and wonder.

I realize how much I owe to the University of Bologna, who instilled in me the thirst for knowledge. In particular, I want to thank my supervisor Roberto Soldati, who accepted to support my research unconditionally. He gave me the freedom to pursue my interests, and he has always been ready to advise me with his experience and positive attitude.

I am deeply thankful to Aldo Faisal, the person that guided me, day after day, in an unfamiliar territory, and that more than anyone else supported me in my research. He accepted me at Imperial College as a part of his group, he believed in my abilities and he introduced me to the world of Theoretical Neuroscience. He has been a real reference for me in these years, both on a professional and personal level, and his contribute to this thesis can not be overestimated.

I want to thank Prof. Vasilis Marmarelis, who welcomed me in his group at the University of Southern California. He has been for me a model of the authentic scientist. His attitude has been inspiring, and I have been honoured to work under his guide for one year. I am also grateful to Tracy, Jennifer and Suzanna for taking care of me in every situation, and to the Viterbi School of Engineering who, together with the University of Bologna, financially supported my research.

I am thankful to all the members of my lab in London. Sharing this experience with them made it so precious and rich. Ali, Anastasia, Andreas, Ariadne, Chin-Hsuan, Constantinos, Ekaterina, Feryal, Luke, Raj, Scott, William, Yesna. The complicity connecting us is rare. Together we will conquer the world one day, we know it's just about finding the right idea.

This thesis exists today for one only reason, and the reason is my parents, Tiziana and Giampiero. Their unconditional love pushed me to follow my passion. Their presence gave me the strength to find my own route, and the resilience to chase it. Their example guided me on this journey.

I also want to thank my grandmother Maria Pia, my godfather Paolo and all my uncles and cousins: Andrea, Arturo, Cristina, Emma, Francesco, Franco, Giulia, Laura, Luca, Marco, Marina, Michele, Nicola and Valeria. I am lucky to have such a beautiful (very Italian) family.

In these years, I have been supported by precious friends. Despite living separate for many years, they represent like a second family for me. Anna, Cesco, Daniele, Elena, Elisa, Eloisa, Francesco, Giacomo, Giulia, Luca, Ludovica, Maria Elena, Riccardo, Roberto, Samuele, Simone, Salima, Theorangehat, Vittoria and all the others, they are too many to fit on this page.

I want to thank Anthony and Dylan, who have truly become like brothers to me in the last year. I also want to thank Cesco for always pushing me to be ambitious.

I undertook this journey with Ilenia. She helped me improve, day after day, and she immensely contributed to shaping the person I am now. We shared the good moments, and she supported me in the difficult ones. For these reasons (and for many others) I thank her parents, Eleonora and Antonio, for making her the girl she is.

Declaration

This is to certify that to the best of my knowledge, the content of this thesis is my own work.

This thesis has not been submitted for any degree or other purposes.

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

Alessandro Ticchi



Contents

Abstract	i
Acknowledgements	iii
1 Introduction	1
1.1 Motivation and objectives	1
1.2 Contributions	5
1.3 Structure of this dissertation	7
2 Background theory	8
2.1 Introduction	8
2.2 Bayes and the logic of science	9
2.2.1 The Bayes' theorem	10
2.2.2 A case study	11

2.2.3	Bayes' theorem, model comparison and Occam's razor	12
2.2.4	Applications of the Bayes' theorem	14
2.3	Physics of the brain	16
2.3.1	Neurons	16
2.3.2	Synapses	21
2.3.3	Noise in the brain	24
3	Adaptation in single neuron	27
3.1	Introduction	27
3.2	Methods	29
3.2.1	The model	31
3.2.2	Inferring the input drive from neural spiking activity	33
3.2.3	Exploiting prior knowledge about the temporal structure of the input	35
3.2.4	From inference to neural adaptation	37
3.2.5	Approximate inference: Steady state solution of the posterior variance	38
3.2.6	Local, on-line learning of optimal model parameters	44
3.2.7	Sanity check: Testing approximated model predictions	47
3.2.8	Decoding	48

3.3	Testable predictions	51
3.3.1	Predicting the time course of neural spiking rate	52
3.3.2	Linear transformation in the input space	55
3.3.3	Linear transformation in the output space	56
3.4	Results	56
3.4.1	Experimental setting	57
3.4.2	Data analysis	57
3.4.3	Validating model predictions on experimental results	60
3.5	Biophysical implementation	63
3.6	Discussion	67
4	Calcium Coding Theory	70
4.1	Introduction	70
4.2	Representing information inside the neuron	74
4.2.1	Calcium encodes log-odds probability ratio of hidden state	74
4.2.2	Why the log-odds probability ratio is a good choice	76
4.3	Synaptic vesicles release	79
4.3.1	Sampling	81

4.3.2	Short term synaptic plasticity	85
4.4	Dendritic integration	94
4.4.1	Decoding synaptic vesicles release	95
4.4.2	Combining different synapses	102
4.5	Action potentials	106
4.5.1	Gain rate normalization	107
4.5.2	Calcium dependent potassium current	109
4.5.3	Action potentials initiation	113
4.5.4	Integrating computational and algorithmic description	119
4.6	Calcium dynamics	124
4.6.1	Inferring the instantaneous neural spiking rate	125
4.6.2	Calcium ions flow	127
4.6.3	Local learning of calcium (hyper) parameters and model predictions . .	130
4.7	Discussion	132
5	Inference by sampling from noisy neuronal circuits	138
5.1	Introduction	138
5.2	Learning the statistics of the world	142

5.2.1	The model	142
5.2.2	Synaptic plasticity	143
5.2.3	Lateral Inhibition	145
5.2.4	Results	146
5.3	Sampling from the prior	148
5.3.1	Spontaneous activity	148
5.3.2	Interpreting spontaneous activity	149
5.4	Inference by sampling (Markov Chain Monte Carlo)	149
5.4.1	Noise driven inference	150
5.4.2	Alternative algorithms	152
5.4.3	Results	152
5.5	Discussion	154
6	Conclusion	158
6.1	Summary of thesis achievements	159
6.1.1	Single neuron	159
6.1.2	Neural population	161
6.2	Common themes shared by both models	163

6.2.1	Noise	163
6.2.2	Learning	164
6.2.3	Sampling	165
6.3	Applications	165
6.3.1	Intra-cellular calcium concentration as a gateway to brain functioning .	166
6.3.2	Synaptic plasticity and computations	167
6.3.3	From neural parameters to the structure of the world	167
6.3.4	A detailed, biophysical model of specific brain areas	167
6.3.5	Spontaneous activity	168
6.3.6	Brain inspired architectures and prosthesis	168
6.4	Future work	169
6.4.1	Single neuron	169
6.4.2	Neural population	170
6.4.3	Combining single neuron and neural population models	172
	List of symbols	176
	Bibliography	178

List of Figures

2.1	Graphical representation of a generative model of causes and consequences, and how it can be used (according to the Bayes' theorem) to infer the probability of hidden states given observations.	11
2.2	A. Visual representation of neural anatomy. B. Diagram representing the time-course of neural membrane potential during a spike event and its different phases.	21
2.3	A. Visual representation of a synapse and its anatomy. B. Diagram of an example of Spike-Timing-Dependent Plasticity (STDP) rule.	23

- 3.1 (A) Neuronal responses of a non-adapting (orange) and an ideally adapting (green) neuron driven by the same input. The adapting neuron represents un-predicted changes in the input instead of the input itself. (B) The σ_e^2 approximations of equation (3.21) (green line), (3.22) (blue line) and (3.48) (orange line) are compared with the exact solution in (3.20) (black dashed line). Parameters $\sigma_i^2 = 100s^{-2}$, $\mu_o = 10s^{-1}$, $\mu_i = 10s^{-1}$, $r_i = 0.01s^{-1}$. (C) The approximated expected value (continuous line) for the term $\langle \sigma_t^2 / c_t \rangle_{c_t} = 9/8 \sqrt{r_i \sigma_i^2 / \mu_o}$ (as in (3.48)), is compared with the simulation results (error bars), computed by simulating the stochastic system in silico and inferring the exact Bayesian posterior $D(\nu)$ (without any approximation). For large values of r_i , condition (3.5) is violated and, as expected, the simulation results diverge from the model predictions. Parameters used in the simulation are $\mu_o = 10s^{-1}$ and $\mu_i = 10s^{-1}$ 43

3.2 The neuron adapts internal parameters locally and on-line to the temporal structure of the input. Simulations results for evolution of hyperparameters $\hat{\mu}$ (A), \hat{r} (B) and \hat{k} (C), following equations (3.55), (3.57) and (3.59). The parameters (solid lines) consistently converge to the optimal values (dashed lines). Simulation parameters: $\mu_i = 10s^{-1}$, $r_i = 0.1s^{-1}$, $\sigma_i^2 = 10s^{-2}$, $\mu_0 = 10s^{-1}$. The time-scale of hyperparameters adaptation is compatible with experimental observations (Sernagor et al., 2001). (D) Raster plot of a non-adapting neuron and (E) of our ideal adapting neuron, encoding the input in fig. 3.2F. (F) The plot represents the encoded input drive (black dashed line), the optimally decoded non-adapting neuron (orange line)(MSE = $11s^{-2}$) and the optimally decoded adapting neuron (blue line)(MSE = $5s^{-2}$). It can be noted that the adapting neuron represents the input drive more accurately than the non-adapting one using the same number of total spikes. In addition, the approximated value c (green line) inferred from the activity of the adapting neuron using equation 3.49 ($R^2 = 0.96$) is compared with the exact Bayesian decoding (blue line), this further validating the used approximations. Simulation parameters $\mu_i = 10s^{-1}$, $r_i = 0.1s^{-1}$, $\sigma_i^2 = 10s^{-2}$, $\mu_0 = 10s^{-1}$. It can be noted that adaptation takes place on the time-scale of individual inter-spike intervals. 49

3.3 The plot visualizes solutions of equation (3.78) for different initial conditions ($f_{t=0} = \tau_o D_b/D_a$). Parameters used are $\frac{9}{8} \frac{1}{D_a} \frac{\sigma_i}{\mu_i} \sqrt{r_i \mu_o} = 2$ 54

-
- 3.4 Visualization of the data we used to test our model, which was generously shared with us by the authors of (Benucci et al., 2013). Color code represents activation level of different neural pools from cat V1 in response to the stimulus provided. Responses are scaled to a value of 1 at the preferred orientation and 0 at the orthogonal orientation. Neural pools are ordered on the y-axes according to the center of their receptive fields (see fig.3.5). The value of the stimulus provided (grid orientation) is represented by a black line. 57
- 3.5 The figure represent the receptive field of cat V1 neural pools extracted from data in fig.3.4. Red dots represent the center of the respective tuning curves. The grids used as stimulus in the experiment are non-directional and, as a consequence, neural tuning curves are invariant for 180 rotations. Neural pools are ordered according to the position of their receptive field centers. 58
- 3.6 Visual representation of data processing used to extract relevant information from the given dataset. For each neuron, we considered a relevant window (centered either at the tuning curve center or at the orthogonal orientation). We then recorded all the neural activation trajectories, starting when the stimulus entered in the relevant window and ending when it first exited the given region. We repeated the same process to collect trajectories for all the neural pools. We superimposed all the trajectories by synchronizing the initial times (when the stimulus entered the relevant window). We then calculated mean and standard error of all the trajectories in each moment in time, to extract the average time-course of adaptation. 59

- 3.7 Time-course of adaptation (normalized neural activity) from in vivo recordings of cat V1 (blue error bars) (Benucci et al., 2013) is compared with analytical predictions from our model (green line) (equation (3.78), fig. 3.3). Red lines represent exponential fits, in which the shared decay constant was chosen as to best predict the set of 6 different conditions. Different plots refer to different relevant window sizes and locations (as described in the top-right diagrams). All the parameters in the analytical model were fixed using dynamical properties (timescale of change, distribution and dynamical range) of the stimulus used in the experimental setting. Plots A, B and C represent time-course of adaptation relative to a relevant window centered at the tuning curve center of the neural pools. Plots D, E and F refer to a relevant window centered at the orthogonal orientation. We can note that in plots A, B, C, E and F analytical predictions match both timescale and shape of adaptation (generalized R^2 A=0.57, B=0.87, C=0.64, E=0.61, F=0.79). In addition model predictions explain experimental data consistently better than the exponential fit (generalized R^2 A=0.24, B=0.77, C=0.64, E=0.54, F=0.71), which is often used in the literature to fit adaptation (Dean et al., 2008; Westerman and Smith, 1984). Experimental results in plot D show a non-adaptive behavior, which differs from the predictions of our model. In the discussed plot, the relevant window is very small and centered at the orthogonal orientation. As a consequence, the neuron receives a very low input drive and resilient network effects become dominant, introducing an inertial component to neural activity decay. 62

- 4.1 The figure represents the result of Bayesian inference in a binary classification problem, where the class-conditional distributions (black continuous lines) are mono-dimensional Gaussians with unit variance and means, respectively, $\mu_0 = -1$ and $\mu_1 = 1$. The posterior probability (blue dashed line) is given by the logistic function $1/(1 + \exp(-2x))$ and the log-odds ratio (orange dashed line) is given by a linear function of the input ($2x$). Figure adapted from (Jordan, 1995). 78
- 4.2 Functional (A) and biophysical (B) illustration of a chemical synapse. When an action potential approaches the synapse, the positive charge causes the opening of voltage-gated calcium channels (green). Calcium pours into the synaptic button and binds to several proteins, changing their shape. The activated proteins dynamically rearrange the blue cytoskeleton to transport green vesicles filled with yellow neurotransmitters to the synaptic cleft, which is filled with red adhesion proteins. Calcium-activated (SNARE) proteins bind to both the vesicles and the synaptic membrane, causing the vesicles to fuse with the membrane, turning them inside out and spilling neurotransmitters into the synaptic cleft. Neurotransmitters then bind to proteins on the receiving cell. There are several types of receptor proteins. Sodium (Na^+) channels (excitatory) respond to the neurotransmitter Glutamate. Chloride (Cl^-) channels (inhibitory) respond to the neurotransmitter GABA. Dopamine, Serotonin, and Opioids bind to G-Protein Coupled Receptors (GPCRs) which cause complicated phosphorylation cascades that change the metabolism of the cell. . . . 80

4.3 Predicted and experimental dependence of release rate of neurotransmitter vesicles on intra-cellular calcium concentration at the axonal termination. (A) Our model predicts that vesicle release rate is proportional to the release probability $P([c])$, which is given by the logistic function of calcium concentration. The graph represents the curve predicted by our model (equations (4.14)-(4.3)), which describe the dependence of release probability on calcium concentration. Parameters $h_o = -3$, $h_1 = 5$ and $[Ca_{ref}^{2+}] = 1\mu M$. (B) Experimental observations about vesicle release rate as a function intra-cellular calcium concentration. Figure adapted from (Schneggenburger and Rosenmund, 2015). 86

4.4 Model predictions about short-term potentiation in synaptic transmission of action potentials (A-C) are compared with experimental observations (B-D). All the graphs represent the ratio between the EPSP evoked by the second and by the first pre-synaptic spike, as a function of the inter-spike interval. (A) Curve predicted by our model of synaptic communication (equation (4.25)). (B) Experimental paired-pulse ratio from mouse hippocampal area CA1 (in-vitro). Schaffer collaterals axons were stimulated using bipolar tungsten electrodes (FHC) with enough current (50-s pulses) to reliably elicit synaptic responses. Figure adapted from (Moresco, 2002). (C) Curve predicted by our model (equation (4.25)) plotted in a semi-log plot. (D) Ratio between first and second field excitatory postsynaptic potential (fEPSP) amplitudes in response to a pair of 175 A stimulations in mice hippocampal Tc1. Figure adapted from (Witton et al., 2015). 91

- 4.5 Experimental results show correlation between the timescale of calcium dynamics and synaptic potentiation decay, validating our predictions. (A) Our model predicts that facilitation is controlled by calcium concentration and, as a consequence, the timescale of facilitation decay should match calcium recovery. (B) Experimental observations in the crayfish claw opener neuromuscular junction, showing a linear correlation between the timescales of calcium and synaptic potentiation. Figure adapted from (Delaney and Tank, 1994). 92
- 4.6 The probability of release predicted by our model (equation (4.14), blue line) is well approximated by the fourth power of calcium concentration (equation (4.30), red line) for small values of calcium (A). (B) For large values of calcium, neurotransmitter release diverges in the case of the power approximation, while it converges to a finite value according to our model. Parameters used in the plots are $h_0 = -7$, $h_1 = 4$ and $n = 0.05$ 93

- 4.7 Open probability of a calcium dependent potassium channel as a function of local calcium concentration. Model predictions are compared with experimental observations. (A) Potassium channel open probability P , as predicted both by our normative model of potassium ions flow (equation (4.77)) and by the proposed statistical mechanics description of ions channels dynamics (equation (4.82)-(4.84)). Parameters used are $h_0 = 0.2$, $h_1 = 0.4$ and $[Ca_{ref}^{2+}] = 1\mu M$ (referring to equation (4.3)). (B) Single channel recordings of BK channels in Chinese hamster neurons. Different lines represent open probability for different mutations of BK channels at different Ca^{2+} concentrations (0.4, 0.7, 10, 20, 70, and 100 μM) measured at +40 mV . Figure adapted from (Díez-Sampedro, 2006). (C) Open probability for different mutations of mSlo potassium channels as a function of calcium concentrations. Values were measured at +20 mV . Figure adapted from (Sullivan and Holmqvist, 1997). 113
- 4.8 Graphical comparison of the expressions in approximation (4.102). (A) The logistic function (blue line) is a good approximation of the normal cumulative distribution function (NCDF - red line) for an appropriate choice of the parameter $q = 1.7$. (B) The difference between the two equations is plotted for $q = 1.7$ 118

- 4.9 The desired gain normalization rule (equation 4.117) is compared to the behaviour of our model of spike initiation (equations (4.108) and (4.112)). As predicted, the detailed model matches the abstract rule for small values of the probability P (A, $P = 0.1$), while for larger values of the probability (B, $P = 0.3$)-(C, $P = 0.6$) the model is a reasonably good approximation of the desired behaviour. $k = \tau_r \mu_o = 3 \text{ ms} \times 0.03 \text{ ms}^{-1} = 0.01$. It can be noted that in the point of interest $P_c = P$ the 2 equations match in all the scenarios, while in the region $P_C \approx P$ the model is a good approximation of the abstract rule also for high values of P 123
- 4.10 Spike evoked calcium concentration is consistent in proximal and distal regions of the neuron. (A) Qualitative model predictions about the time-course of calcium concentration. We predict that spike evoked calcium concentration in different parts of the neuron should be linearly proportional to a baseline value (equation (4.145)). (B) Calcium action potential, propagating along the apical axon of a cortical neuron after synaptic stimulation. Results are expressed as relative change in fluorescence ($\Delta F/F$). Figure adapted from (Schiller et al., 1997). (C-D) Calcium influx caused by strong (C) and weak (D) back-propagation of an action potential in the dendrites of different pyramidal neurons. Figure adapted from (Golding et al., 2001). 130
- 4.11 Calcium Coding Theory. Computational description of neural functioning. . . 135
- 4.12 Calcium Coding Theory. Algorithmic description of neural functioning. . . . 136
- 4.13 Calcium Coding Theory. Implementation description of neural functioning. . . 137

- 5.1 Model of neural population dynamics and functioning. A) Abstract representation of the proposed system: The sensory input, conditional to the state of the world, drives the activity of a cortical population, which encodes the state of the world. B) Probabilistic description of the model, which highlights the relations between different system's constituents. C) Neural implementation of the abstract model. D) Dynamics used to simulate neural populations in-silico. Successive values of the external variable are simulated using a Markov Chain Monte Carlo process, which generates correlated samples from a certain underlying statistics. The spiking activity of the sensory population is modeled using independent Poisson processes, whose instantaneous rate is conditional to the state of the world according to a provided receptive field. Synapses dynamics is simulated using an exponential decay of the synaptic activation variable, which increases of a fixed quantity after each pre-synaptic spike. Cortical neurons are modeled using a Leaky integrate-and-fire conductance-based model, with the addition of white noise and lateral inhibition. E) Simulated behavior of the system. The time-dependent value of the encoded world feature (green line) conditions the spiking activity of the sensory population (200 neurons, orange dots), which in turn drives spikes in the cortical population (20 neurons, blue dots). The position of every sensory/cortical neuron in the ordinate axis corresponds to the center of its tuning curve in the feature space. F) Stochastic leaky integrate-and-Fire dynamics of cortical neurons: In blue, the membrane potential of a cortical neuron is plotted over repeated stochastic simulations driven by the same pre-synaptic input. The dotted-green line represents the leakage potential (-54 mV) while the dotted-orange line represents the threshold for action potential initiation (-40 mV). 144

5.2 Learning the optimal tuning curves: (A) The optimal configuration of tuning curves is plotted for a given stimulus prior (blue line). (B) Mathematical and (C) visual representation of the plasticity rule (TDSP) used in the simulation, which provides the neural population to adapt its tuning curves to the experienced input statistics. W_i^n represents the weight of the synapse connecting pre-synaptic neuron i to post-synaptic neuron n . α is a learning rate parameter and r is a regulatory term. ρ^+ and ρ^- are constants. (D) The cortical population adapts its tuning curves to the input statistics. Different colored lines represent the tuning curves of 20 cortical neurons and their evolution in the feature space, as simulated by our model of neural activity and plasticity. In blue, the prior distribution of the input used in the simulations is plotted, which represents the statistics of the encoded stimulus in the world. The activity of 20 cortical neurons receiving input from 200 sensory neurons was simulated for few hundred minutes (simulation time) starting from a random connectivity matrix. Optimal encoding would predict the gain of the tuning curves to be constant, the tuning curves density to be directly proportional to the prior statistics and the tuning curves widths to be inversely proportional to the prior distribution. (E) The inverse of the distance between the centers of neighbor tuning curves (green dots) and the inverse of tuning curves widths (blue dots) are plotted as a function of the position of the tuning curves centers in the feature space, as produced by our simulations. Optimal encoding would predict that the points lie on the prior distribution (blue line). (F) The results of the simulations for the inverse distance between centers of neighbor tuning curves (green dots) and the inverse tuning curves widths (blue dots) are compared to the values predicted under the optimal encoding assumption. The comparisons provide R^2 values consistently approaching 1. 147

5.3 Inference by sampling. A) In the absence of sensory drive, the neural population generates spontaneous activity as a consequence of lateral inhibition and, in average, each neuron generates the same number of spike events. If decoded optimally (using a decoder adapted to the encoding scheme of the population in presence of sensory drive) each spike event can be interpreted as a sample from the prior distribution, which corresponds to the statistics of the stimulus in the world. The graph shows simulation results. The blue line represents the mathematical prior, while black dots represent the probability distribution extracted from the frequency of samples observed in-silico. Kullback-Leibler divergence $7 \cdot 10^{-5}$). B) A feedback system allows to evaluate the likelihood of each generated samples, and to use this information to control the generation of next samples, implementing a Markov Chain Monte Carlo (MCMC) algorithm. The graph represents simulation results using the MCMC algorithm based on noise discussed in the text. Blue line represents the prior distribution (encoded by spontaneous activity in absence of feedback) and the orange line represents the likelihood function (encoded by the feedback associated to each sample). The green line represents the mathematical posterior (calculated using the Bayes' equation (2.1)) while black dots represent the distribution extracted from the number of spike events generated in-silico by the neural population. The Kullback-Leibler divergence between the mathematical posterior distribution and the distribution inferred in-silico is $1 \cdot 10^{-4}$. Alternative MCMC mechanisms and their respective neural implementations (as discussed in the text) were also simulated in-silico, and they produced consistent results. 153

Chapter 1

Introduction

‘Theoretical physics is the search for simple and universal mathematical descriptions of the natural world. In contrast, much of modern biology is an exploration of the complexity and diversity of life. For many, this contrast is prima facie evidence that theory, in the sense that physicists use the word, is impossible in a biological context. For others, this contrast serves to highlight a grand challenge. I’m an optimist, and believe (along with many colleagues) that the time is ripe for the emergence of a more united theoretical physics of biological systems, building on successes in thinking about particular phenomena.’

William Bialek, “Perspectives on theory at the interface of physics and biology”

1.1 Motivation and objectives

Biological systems are extremely complex and diversified. For this reason, research in living organisms and, specifically, in Neuroscience, is usually driven by a descriptive approach, and experiments represent the primary investigation tool in this research. In the last years, thanks to technological development, more and more areas of the living world are becoming accessible to quantitative investigation. The increasing amount of data requires adopting tools from

mathematics and physics, which become fundamental to make sense of the available information, and to create models that capture system's functioning and dynamics. This approach to biological sciences tries to answer the question of "how" do living systems function, and what are interactions and relations between different components of the systems.

Yet, a different approach is possible, which investigates the question of "why" these systems do what they do, and what are the simple principles underlying this behavior. This alternative approach is not limited to the application of physics to the problem of biology. On the contrary, it requires to extend the current boundaries of theoretical physics, and to "ask physicists' questions about the phenomena of life, looking for the kinds of compelling answers that we expect in the traditional core of physics" (Bialek, 2015). This approach drove various renowned physicist to investigate the question of life and intelligence in the past (Schrodinger and Lewin, 1967; Rayleigh, 1907; Penrose, 2001) and, if successful, it would "reconcile the physicists' desire for concise, unifying theoretical principles with the obvious complexity and diversity of life." This would "generate theories that engage meaningfully with the myriad experimental details of particular systems, yet still are derivable from principles that transcend these details" (Bialek, 2015). The human brain is one of the most complex structures existing in nature, and understanding the principles underlying its dynamics would set light to how we function and what we are.

Brain activity can be modeled as a highly non-linear dynamical system, where each fundamental computational unit (neuron) follows some specific stochastic equations coupled with the rest of the system. The purpose of the nervous system in every living organism is to infer the state of the world (both internal and external to its body) given some available observations and to control muscle contraction and sweating accordingly, to maximize fitness.

In doing this, the nervous system faces several sources of stochasticity. First, the amount of information available is limited, this affecting reliability of the inference process. Moreover, world dynamics is stochastic and often chaotic, making it challenging to predict future events

from the current state of the world. Finally, the outcomes of our interactions with the world are uncertain, and this uncertainty affects our ability to evaluate and select actions in different circumstances. For all these reasons, brains need to be able to deal with probability distributions, and it is reasonable to think that evolutionary pressure drove nervous systems to do it in an efficient way.

Ample experimental evidence shows that brains are optimal or close to optimal when dealing with probability distributions in various situations (Ernst and Banks, 2002a). The Bayesian framework has proven to be very powerful in predicting human behavior in several tasks involving perception (Wozny et al., 2010), decision making (Bogacz, 2007) and motor control (Wolpert, 1997). However, a complete understanding of how these computations are implemented at the neural level is still missing (yet, see (Ma et al., 2006)). Understanding the neural mechanisms underlying Bayesian computations in the brain would provide a principled interpretation of brain dynamics and neural interactions. Furthermore, comprehending how these computations are performed by neurons would help diagnose and treat pathological conditions that affect our perception and decision-making capabilities (Moscovitch, 1982). Finally, learning how the nervous system deals with this complex task would help us design brain-inspired algorithms and circuits, which can be more efficient than existing architectures. For example, the human central nervous system can perform around $3 \cdot 10^{14}$ operations per Joule (Sarpeshkar, 1998). This makes the brain approximately 100 million times more energy efficient than an Intel i7-860 processor, and it would allow performing all computations currently running in Google data centers on a 2.2 Watts laptop battery.

One of the fundamental differences between hardware and wetware architectures is the contribute of noise and variability to the system's functioning. In fact, silicon circuits are designed to implement a deterministic function of the received inputs, and to consistently produce reliable outputs. On the contrary, noise and stochasticity are a dominant component of brain's dynamics (Faisal et al., 2008). Sensory noise affects how the brain receives information from the external world. Cellular, electrical and synaptic variability influence

information processing and communication while motor noise interferes with the execution of movements and actions. All these different sources of variability add up and propagate in the nervous system, producing stochastic responses to the received inputs (Tolhurst et al., 1983; Neishabouri and Faisal, 2014b). This trial-to-trial variability has always been considered as a limiting factor of neural functioning (Shadlen and Newsome, 1994; Abbott and Dayan, 1999; Bialek and Setayeshgar, 2005), and it has been addressed in the literature as something that can not be avoided (as a consequence of size-efficiency-reliability trade-offs implicit in brain architecture (Neishabouri and Faisal)) and that should be limited and neutralized as much as possible (Bialek, 1987; Aho et al., 1988). However, there are various scenarios in which stochasticity can improve computations and representation. For example, intentionally applying noise to continuous variables before digitalizing them allows randomizing quantization error. This procedure, usually referred to as “dithering” (Ulichney, 1988), is often used in audio (Pohlmann, 1985) and image (Roberts, 1962) processing to prevent the emergence of large-scale patterns. Additionally, in the process known as “Stochastic Resonance” noise has been shown to boost detection of weak signals (Longtin, 1993; Faisal et al., 2008; McDonnel and Abbott, 2008), which is useful when false negatives affect animal fitness more than false positives. Moreover, in various reinforcement learning algorithms random variables are used to produce exploration and learning (Abbeel and Ng, 2005). Finally, a broad class of simulation algorithms (usually referred to as Monte Carlo methods) exploits random number generation to probe stochastic real world problems that can not be solved analytically. Variability represents a precious resource for these algorithms, to the extent that when they are implemented in-silico most of the computational time is often spent in the generation of random variables. For this reason, we argue that the brain, driven by millions of years of evolutionary pressure, might have found a way to exploit intrinsic neural variability as a resource.

Specifically, we propose that the nervous system uses neural noise to get in correspondence with the stochastic nature of the external world. This enables us to represent and simu-

late external events, to calculate expected values of possible actions and to select optimal behaviors. Furthermore, noise could be used by the brain to avoid predictable (systematic) non-optimal behaviors caused by limited available information, which could be exploited by other living organisms. In this dissertation, we propose a model neural functioning both on the single neuron and on a neural population level. In both models noise plays a fundamental role, allowing to implement Bayesian computations and to deal efficiently with a stochastic world.

1.2 Contributions

In this dissertation, we investigate the principles governing neural dynamics, and we probe the computational role of neural constituents in brain's functioning. In the model we develop, molecular sources of noise play a crucial role, as they enable the system to generate the desired macroscopic behavior both at the level of single neuron and of neural population.

On a single neuron level, we start from some assumptions about what neurons are doing, and what is their goal. We use Bayesian theory to derive from first principles a normative model of neural computations and, following a top-down approach, we investigate how the resulting algorithms can be implemented by a biophysical neuron. Specifically, we derive a neural model of synaptic vesicle release, dendritic integration, membrane potential fluctuations, action potential initiation, intracellular calcium dynamics and spike-rate adaptation. We describe all the sub-processes of the neuron's processing pipeline at different levels of detail (computational, algorithmic and implementation level), and we show how distinct levels of description are consistent and equivalent, providing a principled interpretation of experimental observations about neural internal biophysics and dynamics. Our model predictions match experimental in-vivo recordings, as well as the known dynamics of neural biophysical quantities.

The model of neural functioning we develop provides a new interpretation of known neural quantities and processes, such as intracellular calcium concentration, synaptic communication and action potentials. Specifically, we propose that calcium level encodes inside the neuron the log-odds probability ratio of a specific hidden event happening in the world. In addition, we suggest that vesicle release takes the form of a sampling process from the given encoded probability and that action potentials serve to coordinate and synchronize calcium dynamics in different regions of the neuron using an efficient (predictive) coding scheme. This new interpretation of neural functioning, if validated, can lead to a deeper understanding of neural dynamics, and it can set a new light on brains computations on a population level.

Moreover, we show how parameters controlling neural dynamics depend on the dynamical properties (e.g. timescale of changes) of the encoded stimulus, and we propose a model of how these parameters can be learned and adapted on-line using the experienced sensory history. Our predictions create a link between neural functioning and the natural properties of the world, which is compatible with experimental results.

On the population level, we propose a model of network dynamics and learning. An Hebbian synaptic plasticity rule allows the population of neurons to adapt to the statistics of the encoded stimulus in order to represent information optimally. Furthermore, in the absence of sensory drive, the proposed model autonomously generates spontaneous activity driven by molecular noise. The given activity takes the form of a sampling process from the prior statistics of the stimulus. In the presence of additional (feedback) sources of information, the population is able to implement Bayesian inference by sampling and, specifically, a Markov Chain Monte Carlo (MCMC) algorithm that infers the posterior distribution over the encoded stimulus given all the available information. Such a process exploits neural sources of variability as the generative engine supporting active perception. Additionally, we propose that the MCMC method is implemented by selectively tuning the level of noise in the population according to the received feedback information. If this will be verified experimentally, it implies that the brain does not only use noise as a resource, but that the local level of noise

is also actively controlled and tuned by the brain to achieve the desired behavior in different situations. The proposed model of neural dynamics is able to autonomously switch between different operational regimes, and it can be extended to solve different tasks (e.g. decision making or action selection) that require computing expected values to select the optimal behavior. The given model is compatible with results from the literature, and the interpretation of spontaneous activity as sampling from the prior is supported by experimental evidence.

1.3 Structure of this dissertation

This dissertation is organized as follows. Chapter 2 introduces some background concepts, useful to better understand the rest of the exposition. No original work is presented in this Chapter, and all the novel results are collected in the following Chapters. Chapter 3 and Chapter 4 are conceptually related. Chapter 3 presents an abstract (computational) model of single neuron spike-rate adaptation. The results discussed in this Chapter inspired a detailed (algorithmic-implementation) model of how these computations can be performed in a biophysical neuron, which is explored in Chapter 4. The given model of neural functioning proposes a new interpretation of known biophysical quantities (e.g. intracellular calcium concentration) and, following this new interpretation, various stages of the neural processing pipeline are reinterpreted and discussed. Chapter 5 proposes a model of neural dynamics, learning and computations on a population level. This Chapter is structured independently. Some of the concepts and approaches used in this Chapter are related to the previous Chapters, yet its development is autonomous and doesn't rely on any of the results or assumptions made in the rest of the dissertation.

Chapter 2

Background theory

‘The actual science of logic is conversant at present only with things either certain, impossible or entirely doubtful, none of which (fortunately) we have to reason on. Therefore the true logic for this world is the Calculus of Probabilities, which takes account of the magnitude of the probability (which is, or which ought to be in a reasonable man’s mind).’

James Clerk Maxwell, Letter to Lewis Campbell, July 1850

2.1 Introduction

This Chapter introduces some background (non-original) information, which will help the reader to get familiar with some of the concepts we will be referring to in the rest of the dissertation.

We will first introduce Bayesian theory of probability, which describes how subjective beliefs should be rationally updated to account for additional sources of evidence. We will define the Bayes’ theorem and we will give a simple case study. We will then explore how the Bayesian approach to probability and model evaluation naturally accounts for model com-

plexity (compatibly with the Occam's razor), and we will give a brief overview of applications of the Bayesian framework to different problems and fields.

Moreover, we will give a brief introduction to the fundamental constituents of the brain from a physics perspective. We will investigate the dynamics of neurons, which are the elementary processing units of the brain, and of synapses, which control neural communication. Finally, we will discuss molecular sources of variability in neural dynamics, and the role of neuronal noise in brain functioning.

2.2 Bayes and the logic of science

Physics rules governing the evolution of nature are intrinsically stochastic. This makes it impossible for us to predict the future of physical systems from their current and past state, especially when dealing with chaotic dynamics. Moreover, we usually have access to limited available information, this limiting our ability to infer the current configuration of the environment and to take decisions accordingly. For these reasons, statistics plays a crucial role in the study of nature, and each prediction we can make about the future is going to be formulated in the form of probabilities and probability distributions.

There are two schools of statistics: frequentist and Bayesian. The frequentist approach relies on minimal assumptions, and it is based on the hypothesis that any given experiment can be considered as one of an infinite sequence of possible repetitions of the same experiment, each capable of producing statistically independent results (Everitt and Skron dal, 2002). This approach can be used to test a null hypothesis, and its outcome will be in the form of an either "true or false" conclusion from a significance test, which measures the probability of the null hypothesis being correct given the observed data.

On the other side, the Bayesian approach takes into account all available information about

the system of interest, and about the generative process underlying observations. This approach can be seen as an extension of propositional logic, and it allows reasoning with propositions whose truth or falsity is uncertain. Bayesian interpretation of probability enables us to represent a state of knowledge about a system, and to update this knowledge whenever additional sources of evidence get available. This approach requires more assumptions than the frequentist one, but it provides a more robust and adaptable framework to work with probabilities. Furthermore, the frequentist approach can be derived as a limit of a Bayesian procedure (Wald, 1950). For this reason, in the following we will focus on the Bayesian approach to statistics, and we will introduce its fundamentals and applications.

2.2.1 The Bayes' theorem

Bayesian probability theory is based on the Bayes' theorem, which describes how a subjective degree of belief should rationally change to account for additional sources of evidence. This approach to belief updating has been introduced by Rev. Thomas Bayes (1701 - 1761), and Pierre-Simon Laplace first proposed its current mathematical formulation (Laplace, 1812). According to the Bayes' theorem (also known as Bayes' rule), whenever a new observation (B) is available, the belief $P(A)$ over the hidden event A is updated as follows:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)} \quad (2.1)$$

where $P(B)$ is the marginal probability of B , $P(A)$ is the initial belief in A (usually referred as 'prior') and $P(B|A)$ is the conditional probability of B given that A is true (usually referred to as 'likelihood') which depends on a generative model of causes and consequences in the world. Finally $P(A|B)$ is the updated belief in A after taking into account the observation that B is true (usually referred as 'posterior').

The Bayes' rule can be easily derived from the chain rule (also called product rule) of prob-

abilities (Schum, 2001)

$$P(A|B)P(B) = P(A \cap B) = P(B|A)P(A) \quad (2.2)$$

In addition, it is possible to note that the denominator term $P(B)$ in the Bayes' rule can be interpreted as a normalization term. In fact, we can write $P(B) = \sum_j P(B|A_j)P(A_j)$, and this guarantees that $P(A|B)$ satisfies the normalization requirement for probabilities. Therefore, we can rewrite the Bayes Rule in a compact form as

$$\text{Posterior}_B(A) \propto \text{Likelihood}_B(A) \text{Prior}(A) \quad (2.3)$$

2.2.2 A case study

The Bayes' rule can be interpreted as a tool which allows to translate observations about the effects into probabilities about the causes generating them. In order to do so, what we need is a causal stochastic model of the world and a prior belief, which is the probability we attribute to each cause before observing any effect.

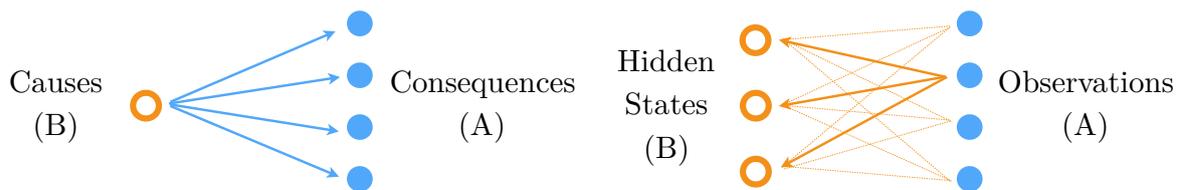


Figure 2.1: Graphical representation of a generative model of causes and consequences, and how it can be used (according to the Bayes' theorem) to infer the probability of hidden states given observations.

We can, for example, consider an application of the Bayes' theorem to virus testing. Suppose we want to test a patient for a certain virus, and we use an exam that is 98% sensitive and

99% specific. This means that, if a healthy person is tested, he will result negative 98% of the times, while if a person affected by the virus is tested he will have 99% probability of being diagnosed positive. In this case, being healthy or affected by the virus will be the set of possible causes, while resulting positive or negative will constitute the set of possible effects. The causal model will be given by the probability of the effects (being positive or negative) for each cause (respectively 0.02-0.98 if healthy and 0.99-0.01 if affected). We also know that, in average, only one person over 100 in the population is affected by the given virus. Let's assume we run the test, and the patient results positive. We can use the Bayes theorem to infer the probability that the given patient is actually affected by the virus:

$$P(\text{affected}|\text{positive}) = \frac{P(\text{positive}|\text{affected}) P(\text{affected})}{P(\text{positive})} = \frac{0.98 \cdot 0.01}{0.0297} = 0.\bar{3} \quad (2.4)$$

As a result, the subject who resulted positive at the test for the virus has about 0.67% probability of being actually healthy, and should be diagnosed negative in absence of additional convincing evidence. The reason for this is the existing prior we introduced in the equation: since only one person every hundred is affected by the virus ($P(\text{affected}) = 0.01$), the likelihood introduced by the outcome of the test is still not providing enough evidence that the patient is actually affected.

In a similar situation, a doctor with a frequentist approach to statistics would have discarded any prior information. He would have deduced that the null hypothesis (the patient being healthy) is rejected with a p-value of 0.02 and, therefore, the patient would have been erroneously diagnosed as affected by the virus.

2.2.3 Bayes' theorem, model comparison and Occam's razor

When fitting models to data, Bayes' theorem can be used at two different levels of inference. At a lower level, given a certain model M and some data to be fitted d , we can use a Bayesian

approach to infer the distribution over the model parameters g . We can write

$$P(g|d, M) = \frac{P(d|g, M)P(g|M)}{P(d|M)} \quad (2.5)$$

The term $P(d|M)$ is usually referred to as evidence for M and, at this level of inference, it can be ignored, since we are only interested in the overall parameter distribution shape and not in the normalization term. In most cases (especially when working in high dimensional spaces) the posterior distribution will be significantly different than zero only in a small region centered at the values of the best parameters g^* that maximizes it. For this reason, we can locally approximate this function by evaluating the Hessian with in g^* :

$$H = -\nabla\nabla\ln P(g|d, M^i)|_{g^*} \quad (2.6)$$

and Taylor expanding the log posterior with $\Delta g = g - g^*$

$$P(g|d, M) \approx P(g^*|d, M)e^{(-\frac{1}{2} \Delta g^t H \Delta g)} \quad (2.7)$$

At a higher level of inference, we can use the Bayes' theorem to infer which of two alternative models (M^1 and M^2) is more plausible given the data. We will assume that each model depends on a set of parameters g , which can have different dimensionality for the different models. Moreover, each model is associated with a prior probability of being the correct one, and it also comes with a prior distribution over the parameters' space $P(g)$. For each model, the likelihood probability of generating the data will be given by the average of the likelihood for different values of the parameters:

$$P(d|M^i) = \int dg P(d|g, M^i) P(g|M^i) \quad (2.8)$$

From (2.5), we can recognize that the integrated function is linearly proportional to the posterior distribution over the model parameters. Since the Hessian is defined as the gradient

of a logarithm, then it is invariant for scaling and, for this reason, we can approximate the integrated function as

$$P(d|g, M^i) P(g|M^i) \approx P(d|g^*, M^i) P(g^*|M^i) e^{(-\frac{1}{2} \Delta g^t H \Delta g)} \quad (2.9)$$

and, using the Laplace method (Azevedo-Filho and Shachter, 1994), we can write

$$P(d|M^i) \approx P(d|g^*, M^i) P(g^*|M^i) \det^{-\frac{1}{2}} \left(\frac{H}{2\pi} \right) \quad (2.10)$$

where, in the mono-dimensional case, $\det(H)$ will correspond to the standard deviation σ_g of the parameter posterior distribution $P(g|d, M)$. We can note that the likelihood of each model is expressed as the product of two terms. The first term $P(d|g^*, M^i)$ is usually called “best fit likelihood”, and it accounts for how well the model can explain the data given the best available choice of the parameters g^* . The second term $P(g^*|M^i) \det(H)$ depends both on the prior of the parameters evaluated at the best fit value g^* , and on how rapidly the accuracy of the fit decays if we deviate from the optimal parameters choice. This second term can be interpreted as an Occam’s Razor factor, in the sense that it advantages simpler models, with fewer parameters and a smaller parameters range. The presence of this term prevents over fitting, and it provides a principled way to compare two models with different complexity without the need to introduce external penalty terms.

2.2.4 Applications of the Bayes’ theorem

Bayesian theory can be invoked whenever we need to infer causes/structure from consequences and observations. In the last decades the Bayesian methods gained considerable popularity, mainly due to developments in computational methods and to the increased availability of big-data, and for this reason they have been successfully applied to different fields. For example, in Physics a Bayesian approach has been used to measure the appropriateness of

models (Hasson and Swithenby, 1999) and for hypothesis testing (Golchi and Lockhart, 2015). In Biochemistry, Bayes' theorem allowed to develop the GOR method (Garnier-Osguthorpe-Robson) for the prediction of secondary structures in proteins (Garnier et al., 1996) and the BATMAN (Bayesian tool for methylation analysis) procedure for analyzing methylated DNA immunoprecipitation (MeDIP) profiles (Down et al., 2008). Furthermore, Bayesian inference is used in marketing and in market research, to evaluate strategies under uncertainty and when dealing with limited data (Rossi and Allenby, 2003; Green and Frank, 1966). In Biology, Bayesian methods are used to infer the evolutionary history, development and relationships among groups of organisms (phylogeny) (Suchard and Redelings, 2006; Yang, 1993). In Medicine, the Bayes' theorem is used for drug testing (Berry, 2006) and for diagnosis (Saaty and Vargas, 1998). Bayesian methods are largely used in the design of intelligent learning systems (Lauria, 2005) and in data analysis (Carlin and Louis, 1997).

Bayesian Theory has also proven to be a very powerful tool for modeling the behavior of living organisms. Colonies of ants (Nonacs and Soriano, 1998) and bees (Biernaskie et al., 2009) have been shown to use a Bayesian approach to foraging and social interactions. Internal dynamics of cells and bacteria has been confirmed to implement Bayesian computations, and to take decisions (e.g. to express a particular gene) accordingly. The Bayesian framework has been successfully used to describe the unconscious decision-making process of Basketball players (e.g. the probability that Kobe Briant will decide to attempt a three vs two points shot, given the observed outcome of his recent attempts (Neiman and Loewenstein, 2011)). More in general, Bayesian models have been successfully used to describe human behavior under uncertainty (Körding and Wolpert, 2003). People has been shown to be Bayesian in perceiving information (Wozny et al., 2010; Ernst and Banks, 2002b), taking decisions (Bogacz, 2007; Körding and Wolpert, 2004) and controlling muscle contraction (Wolpert et al., 1995; Wolpert, 1997). Additionally, using Bayesian assumptions about how neural activity is decoded by other neurons allowed to explain experimental receptive fields of neurons (Ganguli and Simoncelli, 2010) and the structure of spontaneous brain activity in anesthetized

animals (Berkes et al., 2011). Therefore, there is evidence suggesting that the nervous system of living organisms is able to implement Bayesian computations. How these computations are physically implemented in the neural architecture remains an open question, and answering this question requires a deeper understanding of brain biophysics and dynamics.

2.3 Physics of the brain

The human brain contains, in average, 100 billion (10^{11}) neurons. Each neuron exhibits an internal dynamics that follows some non-linear evolution equations, and its internal state is coupled to the dynamics of thousand of other neurons by connections called synapses, giving a typical brain an average well over 100 trillion (10^{14}) total connections. Synapses communicate pieces of information, which are encoded by neurons in the form of action potentials (also called spikes or fires). In average, a single synapse transmits between 5 and 50 spikes every second, producing more than one quadrillion (10^{15}) neural interactions per second. Despite the large number of neurons and neural connections involved, it is difficult to study the brain using a statistical physics approach, since the system is built to continuously fluctuate between non-equilibrium states (Williams-García et al., 2014). Moreover, the main attractor of the system is a critical point (de Arcangelis et al., 2006). This self-organized criticality (SOC)(Bak et al., 1987) provides the system to exhibit macroscopic properties characteristic of critical regimes without the need of parameters' fine tuning. For all these reasons, the brain represents one of the most complex known physical systems, which is able to generate very rich macroscopic behaviors, such as learning, motor control and self-awareness.

2.3.1 Neurons

Nerve cells (usually referred as “neurons”) are the fundamental units of the brain. These cells are specialized to carry information exploiting an electrochemical process, and to com-

municate this information to other cells through neural connections (synapses). Each neuron is composed of three regions: 1) the dendrites, receiving and combining information from other neurons, 2) the soma, which contains the cell nucleus and that receives information integrated by the dendrites and 3) the axon, which propagates and communicates information to other neurons (see Figure 2.2A). In 1952, Alan Lloyd Hodgkin and Andrew Huxley first introduced a model of non-linear differential equations describing the electrochemical mechanism that allows propagating information through the squid giant axon (Hodgkin and Huxley, 1952). This model, which earned them the Nobel Prize in Physiology or Medicine in 1963, investigates the flow of ions (sodium Na and potassium K) through the lipid membrane of the neuron. According to the proposed model, the ions' flow is controlled by some voltage-gated ion channels, and it depends both on the electrical conductance of that specific ion channels (g_i) and on the reversal potential (V_i) of the considered ion type. Specifically, the model predicts that, for each ion, we can express the membrane current I_i as

$$I_i = g_i(V_i - V) \quad (2.11)$$

where V is the membrane potential of the neuron. Moreover, there is a flow of ions provided by the leakage of ions through the membrane potential, which can be express in a similar form given the leakage conductance (g_l) and the reverse leakage potential (V_l). Finally, we can add to these terms an external drive (I), which represents an external perturbation (e.g. the input received from other neurons or some artificially injected current). The sum of the given currents affects the membrane potential V_m according to the membrane capacitance (C_m)

$$C_m \frac{dV}{dt} = g_K(V_K - V) + g_{Na}(V_{Na} - V) + g_l(V_l - V) + I \quad (2.12)$$

In their model, Hodgkin and Huxley modeled the time and voltage dependent ion conductances g_K and g_{Na} as a combination of some dimensionless quantities between 0 and 1, which are associated with potassium channel activation (m), sodium channel activation (n), and

sodium channel inactivation (h). The dynamics of this quantities is described by ordinary linear differential equations and, as a consequence, the membrane potential can be modeled using a set of four coupled nonlinear differential equations as follows

$$C_m \frac{dV}{dt} = \bar{g}_K n^4 (V_K - V) + \bar{g}_{Na} m^3 h (V_{Na} - V) + \bar{g}_l (V_l - V) + I \quad (2.13)$$

$$\frac{dn}{dt} = \alpha_n(V)(1 - n) - \beta_n(V)n \quad (2.14)$$

$$\frac{dm}{dt} = \alpha_m(V)(1 - m) - \beta_m(V)m \quad (2.15)$$

$$\frac{dh}{dt} = \alpha_h(V)(1 - h) - \beta_h(V)h \quad (2.16)$$

Given an appropriate choice of the functions $\alpha_i(V)$ and $\beta_i(V)$, the state variables $V(t)$, $n(t)$, $m(t)$ and $h(t)$ describes how the neuron reacts to an external perturbation $I(t)$. If we treat the external perturbation I as a bifurcation parameter, it is possible to show that the given system presents a Hopf bifurcation. If the value of I is constant in time, and lower than a certain threshold I_{th} , then the system reaches a steady equilibrium. Otherwise, if the value of I exceeds the given threshold I_{th} , the system generates a limit cycle: the membrane potential suddenly undergoes a very large fluctuation (depolarization) driven by a constructive resonance in the fast dynamics of Calcium channels, which brings positive Calcium ions (Ca^{2+}) inside the cell. When the membrane potential reaches a very high value (around 50 mV), the slower dynamics of Potassium channels becomes dominant, a flow of Potassium ions (K^{+}) leaves the cell and it drives the membrane potential back to the original value (repolarization), ready to undergo another cycle. Each depolarization-repolarization event is called action potential (or spike). The duration of each cycle (and, therefore, the time interval between successive spikes) will depend on the external perturbation I . Yet, as a consequence of the Hopf bifurcation, there is a maximum allowed cycle duration, from which the system jumps directly to the stable attractor and the limit cycle disappears. In a similar fashion, if the neuron is perturbed using a single pulse instead of a continuous drive, there will be two possible outcomes: if the pulse amplitude is smaller than a certain threshold, then

the membrane potential will produce small oscillations around the equilibrium value and then converge again to the resting state. If, on the other side, the amplitude of the perturbation is big enough, an action potential will be generated. The action potential is then followed by an hyperpolarization (refractory) period, in which the membrane potential overshoots its resting value, and the neuron is not able (or it is less likely) to generate a new spike (see Figure 2.2B). Action potentials (or spikes) are traditionally considered to be the language used by neurons to encode and communicate information, while membrane fluctuations around the equilibrium value are usually discarded, and they do not influence synaptic communication.

The model of neural dynamics previously described (usually referred to as Hodgkin-Huxley model) captures the internal neural mechanisms responsible for the generation of action potentials, and it gives a detailed description of how the membrane potential dynamics of a neuron responds to and propagates external perturbations. Yet, in various situations we can decide to use a simpler and less computationally expensive model, that effectively captures the membrane response of a neuron to external inputs without getting to the level of details required by the Hodgkin-Huxley model. A popular example of this class of effective neural models is the leaky integrate-and-fire model (Burkitt, 2006). This model assumes that the membrane potential of a neuron follows the dynamics described by the following equation:

$$C_m \frac{dV}{dt} = I + \frac{E - V}{R} \quad (2.17)$$

$$\text{if } V > V_{th} \rightarrow \text{spike} \rightarrow V = E \quad (2.18)$$

where, again, C_m is the membrane capacitance of the neuron, an R is the membrane resistance. The given model proposes that the membrane potential of the neuron integrates the received inputs over time. Furthermore, the leakage component represents an exponential decay towards the resting value E . Finally, the model assumes that whenever the membrane potential exceeds a threshold value V_{th} , the neuron generates a spike, and the membrane potential of the neuron is reset to the resting value E . This model can be expanded with

the addition of various terms. For example, we can express the input drive I in form of conductance, assuming that the pre-synaptic drive will affect the neuron by modulating ions channels open probability, and the consequent flow of molecules

$$I = \sum w^i P^i (E^i - V) \quad (2.19)$$

where the sum over i represents the sum over different synapses, w^i represents the synaptic strength, P^i describes the time-dependent level of activation of the synapse (i.e. the conductance of the synaptic ion channels) and E^i is the equilibrium potential for the given synapse. In addition, we can add to the model a stochastic component, for example a white noise term dw . By adding these terms to our equation, we can write

$$C_m \frac{dV}{dt} = \sum w^i P^i (E^i - V) + \frac{E - V}{R} + dw \quad (2.20)$$

$$\text{if } V > V_{th} \rightarrow \text{spike} \rightarrow V = E \quad (2.21)$$

Finally, since we are dealing with an effective model, we can further combine the model parameters and write

$$\tau_m \frac{dV}{dt} = \sum w^i P^i (E^i - V) + (E - V) + dw \quad (2.22)$$

$$\text{if } V > V_{th} \rightarrow \text{spike} \rightarrow V = E \quad (2.23)$$

where τ_m is the timescale for the exponential decay of the membrane potential to its resting value. The leaky integrate-and-fire model is largely used in the literature because it is analytically tractable (Rudolph and Destexhe, 2006), it is computationally inexpensive to simulate (Hansel et al., 1998) and, despite being a very simple model, it is able to capture most variability of neural spiking activity (given an appropriate choice of the model parameters) (Jolivet, 2004; Burkitt and Clark, 1999; Mihalas and Niebur, 2009).

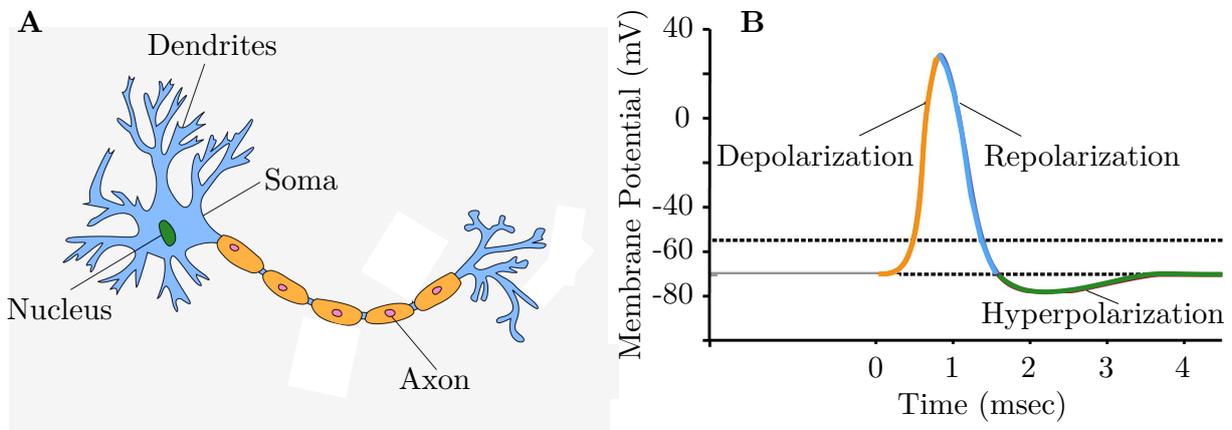


Figure 2.2: A. Visual representation of neural anatomy. B. Diagram representing the time-course of neural membrane potential during a spike event and its different phases.

2.3.2 Synapses

Neurons communicate information through structures called synapses (Eccles, 2013). Specifically, synapses usually convert electrical activity in the pre-synaptic axon (membrane fluctuations and spikes) into the release of chemicals called neurotransmitters, which bind to receptors in the dendrites of the post-synaptic cells. The binding of neurotransmitters can control the conductivity of ions channels in the receiving cell (and a consequent ionic current) or, alternatively, it can initiate a secondary messenger pathway that can either excite or inhibit the post-synaptic neuron. This mechanism enables different neurons to communicate and, given its complexity and the large variety of neurotransmitters, ions channels and receptors involved, it is considered by many to be the chief processing unit of the brain (see Figure 2.3A). It should be mentioned that another class of synapses exists, usually referred to as electrical synapses (in comparison to the previously described chemical synapses) (Connors and Long, 2004). These synapses exploit a very small gap junction between neurons, which is permeable to electric current, and they allow to propagate electric signals. The main advantage of these synapses is that they are very fast, but they enable much less processing compared to a chemical synapse and, for this reason, they are relatively rare in vertebrates'

nervous systems.

Synapses are highly plastic, and their evolution over time is largely responsible for memory formation and adaptation in the brain (Lynch, 2004). Synaptic plasticity is a very rich and complex behavior, it can manifest itself on different timescales (long or short term adaptation) and it is usually classified as either potentiation (when high activity strengthens the synapse) or depression (when repetitive spikes inhibit its activity). There are several mechanisms contributing to changes in synaptic strength over time (e.g. density of neurotransmitters, probability of neurotransmitter release and responsiveness of the post-synaptic cell to neurotransmitters), and synaptic adaptation strongly depends both on the activity history of the synapse, and on the joint dynamics of the pre-synaptic and post-synaptic neurons. One of the earliest models of synaptic adaptation was introduced by Donald Hebb in 1949, who stated that “When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A’s efficiency, as one of the cells firing B, is increased” (Hebb, 1961). This means that, whenever a neuron is efficiently driving the activity of another neuron, the connection between the two neurons gets strengthened. This concept has been reformulated by Carla Shatz’s in the famous statement “Cells that fire together, wire together” (Schatz, 1992; Doidge, 2007). On a first order approximation, this can be interpreted as a correlation learning: when the activity of pre and post-synaptic neurons is correlated (spikes are close in time), then the synapse strength is increased. A simple example of such a rule is given by the following equation:

$$\Delta w_{ij} = \eta x_i x_j \quad (2.24)$$

where w represents the strength of the synapse connecting neuron i with neuron j , η is a learning rate and x_i is the level of activation of neuron i . In this case, the activity of neurons is binned using a certain time window, usually of the order of few milliseconds. Yet, the problem with this rule is that it can not distinguish causality: according to this equation, if the post-synaptic spike proceeds the pre-synaptic one by a very short time there will still

be learning, and the connection strength will increase. Yet, in this case it would be wrong to assume that the pre-synaptic neuron successfully drove the post-synaptic dendrite, and potentiation should not take place. To overcome this problem, a new class of plasticity rules has been introduced, referred to as Spike-Timing-Dependent Plasticity (STDP) (Caporale and Dan, 2008; Roberts and Bell, 2002). According to these rules, the synaptic change depends on the relative time between pre-synaptic and post-synaptic spike, which is referred to as Inter-Spike-Interval (ISI). The learning is usually asymmetric for positive and negative Inter-Spike-Intervals: When a post-synaptic spike follows a pre-synaptic one, there will be potentiation of the synapse (according to the relative time difference), while if the order of the spikes is inverted synaptic depression will take place (see Figure 2.3B). Synapses, together with neurons, represent the fundamental units constituting the nervous system, and having an accurate dynamical model describing their behavior is fundamental in order to understand the functioning of the brain on a larger scale. On the other side, it is very difficult to deeply understand neurons' and synapses' dynamics without understanding their computational role in brain's functioning and, therefore, without putting them in the big picture of a working nervous system.

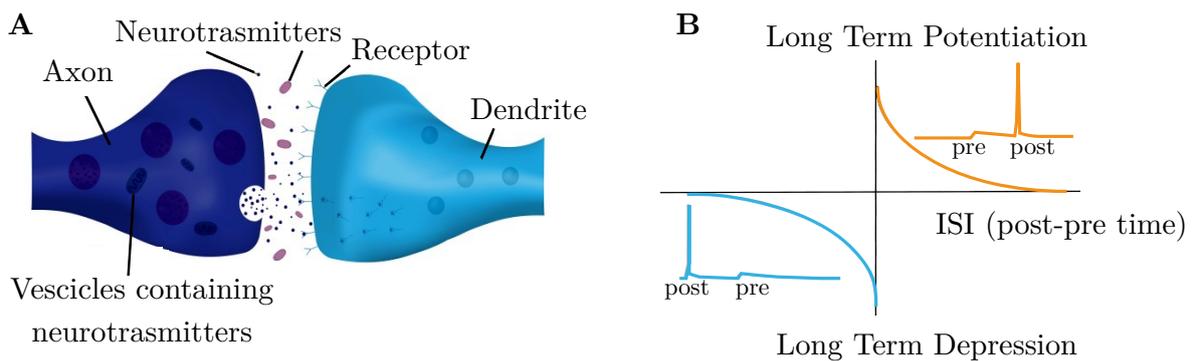


Figure 2.3: A. Visual representation of a synapse and its anatomy. B. Diagram of an example of Spike-Timing-Dependent Plasticity (STDP) rule.

2.3.3 Noise in the brain

Human behavior is intrinsically variable. During repeated trials, our response to a constant stimulus/task is different each time and unpredictable. This behavioral variability (observed in perception, decision making and motor control tasks) can be attributed to three different causes: a) Even if we try to create a controlled environment and to keep the stimulus/task constant, in each trial there will be some environmental differences. This can affect our behavior in a significant way, given the chaotic nature of neural dynamics. b) The initial state of the nervous system will be different at the beginning of each trial, and this will affect the whole trajectory of brain's evolution during the task. c) Brain dynamics is intrinsically noisy, referring to the Oxford English Dictionary definition of noise as 'random or irregular fluctuations or disturbances which are not part of a signal ... or which interfere with or obscure a signal or more generally any distortions or additions which interfere with the transfer of information.'

There are various sources of noise affecting neural dynamics, which can be categorized into five main groups: sensory, cellular, electrical, synaptic and motor noise (Faisal et al., 2008). 1) Sensory noise refers to the stochastic nature of external stimuli. For example, in chemical sensing (smell/taste) thermodynamic fluctuations affect the binding of molecules to the corresponding receptors (Bialek and Setayeshgar, 2005) or, in vision, quantum properties of photons determines their absorption and counting by photoreceptors, which takes the form of a Poisson process (Bialek, 1987). 2) Cellular noise underlies variability of internal neural functioning. Given the same input current, the stochastic nonlinear dynamics of the neuron will produce a variable response to the received stimulation. This affects how different synaptic inputs are integrated, and it generates trial-to-trial variability in the timing of actions potentials on the order of milliseconds (Tolhurst et al., 1983). This variability is captured, for example, by the stochastic term we introduced in the integrate-and-fire neural model in equation (2.21). 3) Electrical noise refers to the fluctuations in the membrane potential of

the neuron, caused by the stochastic nature of ions channels and pumps responsible for controlling the flow of ions through the neural membrane (Derksen and Verveen, 1966). These fluctuations will affect the threshold for initiating action potentials, as well as the propagation of action potentials through the axon. 4) Synaptic noise relates to the variability in how signals are transmitted from one neuron to the other (Kleppe and Robinson, 2006). As described in 2.3.2, electrical signals in the pre-synaptic neuron are communicated to post-synaptic neurons using neurotransmitters. These neurotransmitters are usually contained in synaptic vesicles (see Figure 2.3), and in some cortical neurons there can be very few of these vesicles (up to only one). The presence of a pre-synaptic spike will affect the probability of releasing a vesicle, and the actual release of one or more vesicles will be stochastic and unpredictable. Moreover, the response of the post-synaptic neuron to the release of neurotransmitters will also be stochastic, since the binding of neurotransmitters to receptors is controlled by thermodynamics. 5) Motor noise identifies all the sources of variability introduced by the motor neurons (which carry information from the central nervous system to the muscles), and by the conversion of electrical impulses into forces in the muscles (Frank et al., 2006). Since our only way to interact with the world is through muscles and sweating, motor variability constitutes a fundamental component of behavioral variability. All these sources of noise combined make neural dynamics highly stochastic and unpredictable. Furthermore, this stochasticity propagates and accumulates through different stages of neural processing.

Noise in the brain represents a limit, since it affects information transfer and reliability of neural computations. The brain uses various approaches to minimize negative effects of noise on neural processing, the main ones being redundancy and prior knowledge. Using redundant encoding of information enables the brain to average signals encoding the same information. This, thanks to the central limit theorem, allows to increase the signal-to-noise ratio and to improve information encoding (Kandel et al., 2000). Additionally, redundant muscles are often used to control one single joint. In these cases, it has been shown that the brain co-activates different muscles in such a way that minimizes variability of the movement

execution given specific muscles variabilities (Haruno and Wolpert, 2005). On the other side, having a prior knowledge about the structure of both noise and the encoded stimuli allows to separate one from the others, and to extract relevant information from highly noisy signals. Various sensory signals are highly structured (Barlow, 1961) and, for this reason, filters (neurons receptive fields) that match these structures enables the brain to extract more accurate information about the state of the world from the observed neural activity (Adelman et al., 2003).

Noise has been historically always treated as a problem for neural processing (Shadlen and Newsome, 1994). Yet, neural variability can actually have some beneficial effects on brain dynamics. In other fields, many ways have been discovered in which noise can be used to improve the performance of a system. For example, “stochastic resonance” is a mechanism that allows enhancing the performance of a system in detecting some target events. By adding noise to a weak signal, the probability of exceeding the detection threshold is increased, and such a mechanism has been discovered in cat visual neurons (Longtin et al., 1991) and in other living systems. A further example of noise’s benefits is the process called “dithering”, which allows randomizing quantization errors generated by the digitalization of continuous variables. This method is widely used in image (Roberts, 1962) and audio (Pohlmann, 1985) processing, since it prevents the rise of unwanted large-scale patterns in compressed files. Another possible function of noise is to produce exploration, and to improve learning of an unknown territory. For example, most reinforcement learning algorithms exploit random variables to explore alternative states (Abbeel and Ng, 2005) and to avoid getting stuck in local minima. For this reason, we argue that noise and neural variability can play a crucial role in brain functioning and, in the next Chapters, we will try to show how neural noise can be exploited as a resource by the nervous system.

Chapter 3

Adaptation in single neuron

‘It is hard to see how one could begin to develop a quantum-theoretical description of brain action when one might well have to regard the brain as “observing itself” all the time!’... ‘Beneath all this technicality is the feeling that it is indeed “obvious” that the conscious mind cannot work like a computer, even though much of what is involved in mental activity might do so.’

Roger Penrose, “The Emperor’s New Mind”, 1989

3.1 Introduction

Our sensory systems are continuously recalibrating their sensitivity, to account for changes in the environment (e.g. luminance of a scene) or in the observer (e.g. development or disease). This recalibration process is often referred to as adaptation, and involves different stages of the perception pipeline, from low-level sensing (e.g. sensitivity to brightness of photoreceptors (Stockman et al., 2006)) to abstract representations (e.g. identification of facial features in the cortex (Webster et al., 2004)). This suggests that adaptation operates at different levels of neural coding as a general neural mechanism (Laughlin, 1989), and an abstract model of

neural adaptation could be able to explain how it affects neural computations in different regions of the brain.

Neural adaptation mechanisms have been phenomenologically well described in the past (Kim and Rieke, 2001; Baccus and Meister, 2002; Chander and Chichilnisky, 2001; Zaghloul et al., 2005; Benucci et al., 2013), yet no complete understanding of its computational role and biophysical implementation is known. A key challenge is that biological neurons perform their computations locally as a dynamical system, without any external supervision (using only information directly available to the neuron) and in real-time “on-the-go”. Here, we investigate the computational role of adaptation in neural information encoding. Specifically, we interpret adaptation as an inference problem and, using a generative model for the stochastic dynamics of the encoded stimulus, we derive a normative model of neural adaptation which can be implemented by a biophysical neuron.

Barlows efficient coding hypothesis (Barlow, 1961) states that neural coding should be adapted to the natural properties of the stimulus encoded, in order to represent it as efficiently as possible (in terms of Shannon information) under biophysical (Faisal et al., 2005) and metabolic (Laughlin et al., 1998; Sengupta et al., 2013) constraints. This can be achieved by minimizing correlation and redundancy in neural activity, which limits information capacity. Under this hypothesis, the information theory framework has been successfully exploited to link global statistics of natural stimuli (e.g. statistics of natural images) to properties of neural coding (e.g. receptive fields) (Adelman et al., 2003; Ganguli and Simoncelli, 2010). However, sensory information often presents a correlation structure (both in space and time), and global statistics of natural stimuli (e.g. luminance levels) typically differ from their local distributions. For this reason, according to Barlows efficient coding hypothesis, we can assume that neurons try to adapt their encoding scheme online to the local statistics of the stimuli. In this sense, we can interpret adaptation as an ongoing inference process: a neuron infers on-line the current value of the encoded stimulus, and it uses this information to adapt its response function accordingly in order to maximize encoding efficiency. Under this assumption, we

can exploit a Bayesian theory framework to link dynamical properties of natural stimuli (e.g. timescale on which stimulus change) to properties of neural adaptation (e.g. amplitude and timescale of adaptation).

Starting from a generative model of the stimulus stochastic dynamics and exploiting Bayesian probability theory, we derive a normative model of neural adaptation that links properties of adaptation to the time structure of the encoded stimulus (Ticchi and Faisal, 2014). We assume that, to infer the current value of the stimulus, neurons exploit both observations about their own spiking history and the prior knowledge about the stimulus dynamical properties. We mathematically solve the inference problem and we show that, under some biophysically reasonable assumptions, the optimal solution can be approximated by a first-order linear differential equation, which can be implemented by neurons ion-channels dynamics. Our model provides various quantitative predictions, which are in good agreement with experimental results about spike rate adaptation in V1.

3.2 Methods

We take a normative approach, and we assume that the computational goal of neural adaptation is to maximize information encoded by a neuron (per spike event) about a certain hidden variable. This hidden variable follows a stochastic dynamics (e.g. random walk or stochastic jump process) with some specific dynamical properties (e.g. timescale on which the value changes, prior distribution mean, prior distribution standard deviation). In order to maximize encoding efficiency, the neuron exploits its own spiking history to infer on-line the instantaneous mean of the encoded hidden variable, and it uses the inferred value as divisive gain normalization to control its response function. This provides the neuron to encode only deviations of the hidden variable from the predicted value (predictive coding (Boerlin et al., 2013)) instead of representing the value itself. By doing this, neural adaptation reduces cor-

relation in the spiking activity of the neuron (coding redundancy), this maximizing entropy of the information channel and, therefore, information capacity. Moreover, this mechanism guarantees the average firing rate of the neuron to remain close to the desired value (which depends on neural biophysical and metabolic constraints), while the value of the encoded variable can span various orders of magnitude while dynamically changing.

In order to infer the instantaneous value of the encoded variable, the neuron will exploit both the recent observations about its own spiking activity and some prior knowledge about the dynamical properties of the encoded stimulus. Whenever a new observation is available (a spike event is generated, preceded by a certain inter spike interval), this information is used to estimate the value of the hidden variable that is driving neural activity. Furthermore, as time elapses, prior knowledge about the time structure of the stimulus is used to update the current estimation (distribution) of the given variable, to account for the increase in uncertainty derived from the stochastic dynamics of the stimulus. Finally, the updated distribution (specifically its mean value) is used by the neuron to tune its response function (using divisive gain normalization), and to control the generation of next spike event.

It is interesting to note that, in our model, neural adaptation only relies on the spiking history of the given neuron to infer the current value of the hidden variable. Since the spiking activity of a pre-synaptic neuron is known by a post-synaptic neuron decoding it, then the proposed encoding scheme is nonambiguous, and can be decoded optimally by a post-synaptic neuron. The adapting neuron would potentially have access to additional sources of information regarding the current value of the hidden variable (e.g. its membrane potential or synaptic activation level), but using these quantities in the inference process would make the encoding scheme ambiguous and impossible to be decoded optimally.

The full mathematical solution to the problem is relatively complicated (see supplementary material), and its implementation is computationally expensive. In fact, it requires to encode the full posterior distribution in each moment in time and to update it in each iteration,

which represents a complicated task for a biophysical neuron. We show that, under the biophysically reasonable assumption that the average inter spike interval is shorter than the timescale on which the stimulus changes, the posterior distribution of the hidden variable can be approximated using a gamma distribution. In this case, the inference problem can be rewritten in the form of flow equations for the parameters of the gamma distribution, which can be implemented by the internal dynamics of the neuron.

3.2.1 The model

We model the adaptation process as composed of two sub-processes: **a.** the neuron uses its own activity to infer the instantaneous stimulus statistics and then **b.** it adapts its internal response properties accordingly to maximize information transfer to downstream neurons. In **a.**, the neuron infers the posterior distribution over the future expected synaptic input drive I , which is the variable encoded by the neuron. For convenience, we will express the input drive I in units of firing rate, assuming that in absence of adaptation the input drive would imply a particular firing rate of the neuron (or a certain probability density to generate a spike event at each moment in time). To infer the distribution of I , the neuron exploits both the information contained in its own spiking history and some prior knowledge about the dynamical properties of the stimulus. The neuron may have access to additional sources of information about I available within the cell, e.g. membrane potential fluctuations or individual synaptic drives. Yet, the goal of the proposed adaptation mechanism is to implement predictive coding and, for this reason, we need to avoid decoding ambiguities (Fairhall et al., 2001). Therefore, we will base the adaptation process only on quantities that can be observed by the external decoder, i.e. by the post-synaptic neuron.

We show that, for a broad class of input dynamics, the inferred posterior is well described by a gamma distribution, and we derive the evolution equations for the distribution parameters. Under the ecologically reasonable assumption that the input stimulus changes on a time-scale

longer than the average inter-spike interval, we find the dynamics of the inference process to be well described by first order differential equations. We argue that these equations can be implemented by the internal dynamics of the neuron, and that the inferred quantities can be used by the neuron to regulate its response behavior accordingly (Turrigiano, 2008). We also show how weighted temporal integration may enable postsynaptic neurons to decode the information encoded in the spiking activity of a presynaptic Bayesian adapting neuron.

Our model uses divisive gain normalization, which is a widely documented mechanism for gain control both on the single neuron and population level (Olsen et al., 2010; Wark et al., 2007). Specifically, we propose that neurons will use the inferred distribution of I to calculate its mean c (i.e. the expected value of the input in the immediate future) and use this value to scale (adapt) the output firing rate ν :

$$\nu = I \frac{\mu_o}{c} \quad (3.1)$$

where I is the input drive of the neuron (in units of firing rate) and μ_o is the ideal firing rate, which depends on neural biophysical constraints (e.g. metabolic cost of generating action potentials (Sengupta et al., 2013) and maximum peak firing rate, due to refractory period (Yeomans, 1979)). The resulting value ν will correspond to the instantaneous probability density of the neuron to generate a spike event or, equivalently, with the average firing rate of the neuron. Equation (3.1) provides the neuron to maintain its average firing rate close to the desired value μ_o and, therefore, to maximize information transfer under neural biophysical constraints. To maximize information capacity of the information channel between pre-synaptic and post-synaptic neurons, Information Theory prescribes maximizing the entropy of spiking activity: Each spike event should be ideally as independent as possible from the preceding history of spikes. In our model of neural adaptation, this condition translates into the mean of the future firing rate ν being independent of the spiking history (as inferred by an external observer), thanks to (3.1). Each spike event will therefore encode new non-

redundant information (Hosoya et al., 2005; Rao and Ballard, 1999; Boerlin et al., 2013), and spiking activity will efficiently reflect unpredicted changes in the stimulus, thus implementing predictive coding (Fig. 3.1A).

In the first part of the derivation we will assume that, ideally, the adapting neuron knows everything about the dynamical properties of the input; then, in Section 3.2.6, we will show how adaptation parameters can be learned on-line and locally.

3.2.2 Inferring the input drive from neural spiking activity

In order to infer the instantaneous input drive I we consider an ideal observer, who exploits both the likelihood encoded in the inter-spike intervals of the neuron and some prior knowledge about the input dynamics to infer the posterior distribution over I . We approximate the spiking behavior of the single neuron by an inhomogeneous Poisson process. This approximation provides a good description of the variability observed in neural spiking (Shadlen and Newsome, 1994; Faisal et al., 2008). However, more detailed spiking models (Kass and Ventura, 2001) could be used in the future, to investigate how specific features of neural behavior (such as limits set by absolute and relative refractory periods) affects model predictions.

In the first part of this section we assume that the neuron is not adapting, and that the instantaneous firing rate is directly proportional to the input drive $\nu = I$. Later, in Section 3.2.4, we will show how the derived inference mechanism can be extended to an adapting neuron.

Let $D(\nu)$ be the posterior belief about the input $\nu = I$. Whenever a new spike (and its inter-spike interval τ) is observed, the current posterior distribution becomes the prior for a new inference iteration. In particular, we know that the inter-events distribution of a Poisson process is exponentially distributed and, therefore, given the observation τ , the normalised likelihood distribution over the unknown rate $\nu = I$ will be gamma distributed,

with parameters $\alpha_L = 2$ and $\beta_L = \tau$:

$$L_\tau(\nu) = \tau^2 \nu e^{-\nu\tau} = g(\nu; \alpha_L = 2, \beta_L = \tau) ; \quad g(\nu; \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \nu^{\alpha-1} e^{-\beta\nu}$$

The mean and variance of the Gamma distribution can be calculated, respectively, as

$$\mu_L = \frac{\alpha_L}{\beta_L} = \frac{2}{\tau} ; \quad \sigma_L^2 = \frac{\alpha_L}{\beta_L^2} = \frac{2}{\tau^2} \quad (3.2)$$

The gamma distribution is the conjugate prior of the exponential distribution (Gelman et al., 2003) and, therefore, if the prior is gamma distributed, we can exploit the properties for the product of Gamma distribution to prove the following relation between the parameters of the prior D_t and of the posterior $D_{t+\tau}$:

$$D_{t+\tau}(\nu) = K L_\tau(\nu) D_t(\nu) = g(\nu; \alpha_L + \alpha_t - 1, \beta_L + \beta_t) = g(\nu; \alpha_{t+\tau}, \beta_{t+\tau}) \quad (3.3)$$

In terms of mean and variance, whenever a new inter-spike interval is observed, the parameters of the Gamma distribution are update as follows:

$$\mu_{t+\tau} = \frac{\alpha_t + \alpha_L}{\beta_t + \beta_L} = \frac{\frac{\mu_t^2}{\sigma_t^2} + 1}{\frac{\mu_t}{\sigma_t^2} + \tau} ; \quad \sigma_{t+\tau}^2 = \frac{\alpha_t + \alpha_L}{(\beta_t + \beta_L)^2} = \frac{\frac{\mu_t^2}{\sigma_t^2} + 1}{(\frac{\mu_t}{\sigma_t^2} + \tau)^2} \quad (3.4)$$

We can now consider the regime when the time-scale over which the input is changing is bigger than the average interspike interval

$$\langle \tau \rangle_t \ll \frac{1}{r_i} \quad (3.5)$$

where $\langle \dots \rangle_t$ denotes averaging over time and r_i is defined as the inverse of the time-scale on which the encoded stimulus changes. This assumption is compatible with experimental observations: the average inter-spike interval of cortical neurons is usually within the range

50ms $< \tau < 1$ s (Roxin et al., 2011), while the time scale on which the statistics of the stimuli changes (e.g. we enter and exit a dark tunnel) is usually on the time scale of seconds. In addition, the given assumption is justified by the fact that a neuron would evidently be ineffective in representing variables that changes faster than its own firing rate. It is possible to show that condition (3.5) implies $\sigma_t^2 \ll \sigma_L^2 = \frac{2}{\tau^2}$ or, equivalently, $\frac{\mu_t}{\sigma_t^2} = \beta_t \gg \beta_L = \tau$. In the given regime, we can Taylor expand (3.4):

$$\mu_{t+\tau} = \mu_\tau + \frac{\sigma_t^2}{\mu_t} (1 - \mu_t \tau) - \frac{\sigma_t^4 \tau}{\mu_t^2} (1 - \mu_t \tau) \left(1 + O\left(\frac{\sigma_t^2 \tau}{\mu_t}\right) \right) \quad (3.6)$$

$$\sigma_{t+\tau}^2 = \sigma_t^2 + \frac{\sigma_t^4}{\mu_t^2} (1 - 2\mu_t \tau) \left(1 + O\left(\frac{\sigma_t^2 \tau}{\mu_t}\right) \right) \quad (3.7)$$

From now on, we will only consider the first order term of the expansion. Equations (3.6) and (3.7) describe how the posterior distribution over $I = \nu$ is updated according to a new observed inter-spike interval. Yet, the dynamic of the input I is stochastic and unknown, this introducing additional uncertainty about the current value of I . If we do not receive any new information for a long time, the entropy of the posterior belief $D(I)$ will increase, and the distribution will eventually converge to the global prior input statistics $P(I)$. How this information loss takes place and how it affects the posterior distribution $D(\nu)$ depends on the underlying dynamics of the input.

3.2.3 Exploiting prior knowledge about the temporal structure of the input

We use a stochastic jump process as the generative model for the dynamics of the input drive $I = \nu$. Specifically, we consider a hidden Poisson process with rate r_i and assume that, whenever an event takes place, the value of the input I is set to a new random value, sampled from a Gamma distribution with parameters α_i and β_i .

The distribution $g(I ; \alpha_i, \beta_i)$ represents the prior statistics of I , and the value of $I = \nu_t$ remains constant as long as there is no event in the underlying Poisson process. We will use this model because of its mathematical simplicity and its flexibility in describing different processes. We also tried alternative generative models for the input dynamics (e.g. modified Wiener process) and we achieved consistent analytical results, which will not be reported here for brevity.

In the stochastic jump model, the probability of a jump taking place during an infinitesimal time interval τ is $r_i\tau$, and the updated belief $D_{t+\tau}(\nu)$ is a weighted mixture of the posterior at time t $D_t(\nu)$ and of the input statistics $P(\nu) = g(\nu; \alpha_i, \beta_i)$, with weight $r_i\tau$.

$$D_{t+\tau}(\nu) = (1 - r_i\tau)D_t(\nu) + r_i\tau P(\nu) = (1 - r_i\tau)g(\nu ; \alpha_t, \beta_t) + r_i\tau g(\nu ; \alpha_i, \beta_i) \quad (3.8)$$

We approximate the resulting distribution in equation (3.8) to a Gamma distribution with matching mean and variance. In general, a weighted sum of Gamma distributions is not necessarily Gamma distributed. However, as a consequence of condition (3.5), it always holds that $r_i\tau \ll 1$. Thus, making the approximation accurate with negligible consequences. Under this assumption, we can exploit the equations for mean and variance of a weighted mixture of distributions

$$E(X) = E[E(X|\theta)] \quad ; \quad \text{Var}(X) = E[\text{Var}(X|\theta)] + \text{Var}(E[X|\theta]) \quad (3.9)$$

in order to write the flow equations for the mean and variance of the approximated posterior Gamma distribution $D(\nu)$

$$\mu_{t+\tau} = (1 - r_i\tau)\mu_t + r_i\tau\mu_i = \mu_t + r_i\tau(\mu_i - \mu_t) \quad (3.10)$$

$$\sigma_{t+\tau}^2 = (1 - r_i\tau)(\mu_t^2 + \sigma_t^2) + r_i\tau(\mu_i^2 + \sigma_i^2) - ((1 - r_i\tau)\mu_t + r_i\tau\mu_i)^2 \quad (3.11)$$

In the considered regime (3.5) of a slow changing external input $r_i\tau \ll 1$, we can expand

equation (3.11) as follows:

$$\sigma_{t+\tau}^2 = \sigma_t^2 + (\sigma_i^2 - \sigma_t^2 + (\mu_t - \mu_i)^2)r_i\tau + O(r_i^2\tau^2) \quad (3.12)$$

Therefore, integrating both the effects of the information gain (3.6)(3.7), derived from the observation of the inter-spike interval τ , and the information loss (3.10)(3.12), caused by the unknown input dynamics during the time interval τ , whenever a new inter-spike interval τ is observed the belief about the input value ν is updated as follows:

$$\mu_{t+\tau} = \mu_t + \frac{\sigma_t^2}{\mu_t} (1 - \mu_t\tau) + r_i\tau(\mu_i - \mu_t) \quad (3.13)$$

$$\sigma_{t+\tau}^2 = \sigma_t^2 + \frac{\sigma_t^4}{\mu_t^2} (1 - 2\mu_t\tau) + r_i\tau(\sigma_i^2 - \sigma_t^2 + (\mu_t - \mu_i)^2) \quad (3.14)$$

In each moment in time, the posterior distribution over the input drive $I = \nu$ will be given by $D_t(\nu) = g(\nu; \alpha_t, \beta_t)$ where $\alpha_t = \frac{\mu_t^2}{\sigma_t^2}$ and $\beta_t = \frac{\mu_t}{\sigma_t^2}$.

3.2.4 From inference to neural adaptation

We can now introduce adaptation by rescaling the firing rate of the neuron as in (3.1):

$$\nu_* = \nu \frac{\mu_o}{c} \quad (3.15)$$

where c is the mean of the inferred posterior distribution over the input drive I . Consider two neurons which receive the same input I : the non-adapting neuron will respond with an instantaneous Poisson firing rate ν , the adapting neuron with a Poisson firing rate ν_* , according to the adaptation driven rescaling. In (3.13) and (3.14) τ represents a sample from an exponential distribution with mean ν^{-1} . Yet, in the adapting neuron we will observe the value τ_* , which is a sample from the scaled exponential distribution with mean $\nu_*^{-1} =$

$(\nu \frac{\mu_o}{c})^{-1}$. We can therefore replace τ with $\tau_* \frac{\mu_o}{c}$ in (3.13) and (3.14)

$$c_{t+\tau_*} = c_t + \frac{\sigma_t^2}{c_t} (1 - \mu_o \tau_*) + r_i \tau_* (\mu_i - c_t) \quad (3.16)$$

$$\sigma_{t+\tau_*}^2 = \sigma_t^2 + \frac{\sigma_t^4}{c_t^2} (1 - 2\mu_o \tau_*) + r_i \tau_* (\sigma_i^2 - \sigma_t^2 + (c_t - \mu_i)^2) \quad (3.17)$$

The conversion term $\frac{\mu_o}{c}$ has not been introduced in the substitution of the second τ with τ_* in both (3.16) and (3.17). The reason is that the information loss depends only on the absolute elapsed time and on the underlying input dynamics, and not on the process from which the inter-spike interval is sampled.

3.2.5 Approximate inference: Steady state solution of the posterior variance

To simplify the proposed inference mechanism, we will now reduce the system of coupled equations (3.16) and (3.17) to a single equation describing the evolution of c_t , which can be potentially implemented by biochemical or electrical processes inside the neuron.

To this end, we investigate the behaviour of σ_t^2 and approximate it as a function of c_t . We first derive the stationary solution σ_e^2 of equation (3.17), that satisfies

$$\left\langle \frac{\sigma_t^4}{c_t^2} (1 - 2\mu_o \tau_*) + r_i \tau_* (\sigma_i^2 - \sigma_t^2 + (c_t - \mu_i)^2) \right\rangle = 0. \quad (3.18)$$

Using Ito's Lemma we can exploit the relation $\langle \tau_* \rangle = \frac{1}{\mu_o}$ and, as τ_* appears only as a linear term inside equation (3.17), we can write

$$\frac{\sigma_t^4}{c_t^2} \left(1 - 2\mu_o \frac{1}{\mu_o} \right) + r_i \frac{1}{\mu_o} (\sigma_i^2 - \sigma_t^2 + (c_t - \mu_i)^2) = 0 \quad (3.19)$$

The only positive solution of this equation being

$$\sigma_e^2 = \frac{-c_t^2 r_i + c_t \sqrt{r_i} \sqrt{\mu_o 4(c_t - \mu_i)^2 + c_t^2 r_i + 4\mu_o \sigma_i^2}}{2\mu_o} \quad (3.20)$$

We can now recall condition (3.5) which, translated to the dynamics of the adapting neuron, becomes $\langle \tau_* \rangle \ll \frac{1}{r_i}$. BY exploiting the relation $\langle \tau_* \rangle = \frac{1}{\mu_o}$, we can write $r_i \ll \mu_o$. Using this relation we expand (3.20) to first order in $\sqrt{\frac{r_i}{\mu_o}}$

$$\sigma_e^2 \approx c_t \sqrt{\frac{r_i ((c_t - \mu_i)^2 + \sigma_i^2)}{\mu_o}} \quad (3.21) \quad \xrightarrow{\sigma_i^2 \rightarrow \infty} \quad \sigma_e^2 \approx c_t \sqrt{\frac{r_i \sigma_i^2}{\mu_o}} \quad (3.22)$$

The expansion for $\sigma_i^2 \gg (c_t - \mu_i)^2$ is not completely justified. In fact, by definition, c_t tracks ν and $P(c) \approx P(\nu)$. Therefore, $\langle (c_t - \mu_i)^2 \rangle \approx \sigma_i^2$.

In the following, we will show that the linear approximation $\sigma_e^2 \propto c_t$ is reasonable independently of the value of σ_i^2 . The reason is that for small σ_i^2 the parameter c_t will have small oscillations around its mean value μ_i and, therefore, a linear approximation will be good enough for our purposes.

First, we can consider the case of big σ_i^2 , e.g. $\sigma_i^2 = \mu_i^2$, that represent the situation of an exponentially distributed prior. In this regime, the function can be reasonably approximated by a linear function in c_t in all the region of interest.

On the contrary, if we consider small σ_i^2 , we lose the desired linear behaviour of the function in the given region. However, for small σ_i^2 , c_t will be mainly concentrated around its mean value μ_o . Moreover, we can see that the derivative

$$\partial_{c_t} \sigma_e^2 = \frac{c_t \sqrt{r} (2c_t - 2\mu_i)}{2\sqrt{\mu_o} (c_t^2 - 2c_t \mu_i + \mu_i^2 + \sigma_i^2)} + \frac{\sqrt{r_i} \sqrt{\mu_o} (c_t^2 - 2c_t \mu_i + \mu_i^2 + \sigma_i^2)}{\mu_o} \quad (3.23)$$

which, if evaluated in $c_t = \mu_i$, becomes

$$\partial_{c_t} \sigma_e^2 = \frac{\sqrt{r_i \sigma_i^2}}{\sqrt{\mu_o}} \quad (3.24)$$

This meaning that the linear approximation in c_t

$$\sigma_e^2 \approx \sigma_e^2|_{c_t=\mu_o} + \partial_{c_t} \sigma_e^2|_{c_t=\mu_o} (c_t - \mu_o) = \frac{\sqrt{r_i \sigma_i^2}}{\sqrt{\mu_o}} \mu_o + \frac{\sqrt{r_i \sigma_i^2}}{\sqrt{\mu_o}} (c_t - \mu_o) \quad (3.25)$$

$$= \frac{\sqrt{r_i \sigma_i^2}}{\sqrt{\mu_o}} c_t \quad (3.26)$$

is a reasonable around the mean value $c_t \approx \mu_i$. For this reason, from now on we will accept this linear approximation, and for our purposes we will define $\sigma_e^2 = k c_t$ where

$$k_{\mu_i, \sigma_i^2, \mu_o, r_i} = \left\langle \frac{\sigma_e^2}{c} \right\rangle_P \quad (3.27)$$

$$= \int_0^\infty P(c; \mu_i, \sigma_i^2) \frac{\sigma_e^2}{c} dc \quad (3.28)$$

to evaluate the given function we can Taylor Expand σ_e^2 around μ_i , which is the mean value of P , and get

$$k_{\mu_i, \sigma_i^2, \mu_o, r_i} = \sum_{j=0}^{\infty} \frac{1}{j!} \left. \partial_c^j \frac{\sigma_e^2}{c} \right|_{\mu_i} \langle (c - \mu_i)^j \rangle_P \quad (3.29)$$

which leads to

$$k_{\mu_i, \sigma_i^2, \mu_o, r_i} = \left. \frac{\sigma_e^2}{c} \right|_{\mu_i} + \frac{1}{2} \left. \partial_c^2 \frac{\sigma_e^2}{c} \right|_{\mu_i} \sigma_i^2 + \frac{1}{8} \left. \partial_c^4 \frac{\sigma_e^2}{c} \right|_{\mu_i} \sigma_i^4 + O(\sigma_i^6) \quad (3.30)$$

We can now use the moments generating function of the Gamma distribution

$$mgf(t) = \left(1 - \frac{t}{\beta}\right)^{-\alpha} \quad (3.31)$$

to calculate the n-order moment of the distribution

$$M_{\mu_i}^n = \langle (c - \mu_i)^n \rangle_P \quad (3.32)$$

given

$$M_0^n = \langle c^n \rangle_P = \partial_t^n \text{mgf}(t)|_{t=0} \quad (3.33)$$

Finally, we can write

$$M_{\mu_i}^n = (n-1)! \frac{\alpha}{\beta^n} \quad (3.34)$$

$$= (n-1)! \frac{\mu_i^2 \sigma_i^{2n}}{\sigma_i^2 \mu_i^n} \quad (3.35)$$

$$= (n-1)! \frac{\sigma_i^{2n-2}}{\mu_i^{n-2}} \quad (3.36)$$

At this point, we can explicitly calculate the terms $\partial_c^4 \sigma_e^2|_{\mu_i}$ using equation 3.21

$$\partial_c^2 \frac{\sigma_e^2}{c} |_{\mu_i} = \frac{1}{\sigma_i} \frac{\sqrt{r_i}}{\sqrt{\mu_o}} \quad (3.37)$$

$$\partial_c^3 \frac{\sigma_e^2}{c} |_{\mu_i} = 0 \quad (3.38)$$

$$\partial_c^4 \frac{\sigma_e^2}{c} |_{\mu_i} = -3 \frac{1}{\sigma_i^3} \frac{\sqrt{r_i}}{\sqrt{\mu_o}} \quad (3.39)$$

$$\partial_c^5 \frac{\sigma_e^2}{c} |_{\mu_i} = 0 \quad (3.40)$$

$$\partial_c^6 \frac{\sigma_e^2}{c} |_{\mu_i} = 45 \frac{1}{\sigma_i^5} \frac{\sqrt{r_i}}{\sqrt{\mu_o}} \quad (3.41)$$

$$(3.42)$$

Substituting each explicit term in equation 3.29, we can now write

$$k_{\mu_i, \sigma_i^2, \mu_o, r_i} \approx \frac{\sqrt{r_i \sigma_i^2}}{\sqrt{\mu_o}} + \frac{1}{2} \frac{1}{\sigma_i} \frac{\sqrt{r_i}}{\sqrt{\mu_o}} \sigma_i^2 \quad (3.43)$$

$$- \frac{1}{4!} \frac{1}{3} \frac{1}{\sigma_i^3} \frac{\sqrt{r_i}}{\sqrt{\mu_o}} \frac{3! \sigma_i^6}{\mu_i^2} \quad (3.44)$$

$$+ \frac{1}{6!} \frac{1}{45} \frac{1}{\sigma_i^5} \frac{\sqrt{r_i}}{\sqrt{\mu_o}} \frac{5! \sigma_i^{10}}{\mu_i^4} \quad (3.45)$$

and, simplifying,

$$k_{\mu_i, \sigma_i^2, \mu_o, r_i} \approx \frac{\sqrt{r_i}}{\sqrt{\mu_o}} \left(\sigma_i + \frac{1}{2} \sigma_i - \frac{3}{4} \frac{\sigma_i^3}{\mu_i^2} + \frac{15}{2} \frac{\sigma_i^5}{\mu_i^4} \right) \quad (3.46)$$

The full series expansion would be exact if dealing with a gaussian distribution, but for our purposes will be sufficient, as for small values of the variance regime the gamma distribution converges to the Gaussian. If we calculate the first three terms of the expansion explicitly, we get

$$k_{\mu_i, \sigma_i^2, \mu_o, r_i} \approx \frac{9}{8} \frac{\sqrt{r_i \sigma_i^2}}{\sqrt{\mu_o}} \quad (3.47)$$

and we can write

$$\sigma_e^2 \approx k_{\mu_i, \sigma_i^2, \mu_o, r_i} c_t \approx \frac{9}{8} \frac{\sqrt{r_i \sigma_i^2}}{\sqrt{\mu_o}} c_t \quad (3.48)$$

It is possible to show that using higher order terms in the Taylor expansion would only affect the multiplicative factor $\frac{9}{8}$, and not the structure of the term. However, for the purposes of our neuron model, the approximation (3.22) will be sufficient: Figure 3.1B shows how the expansion in (3.21) and the linear approximations in (3.22) and (3.48) compare with the full analytical stationary solution (3.20). Furthermore, In Figure 3.1C, the approximation in (3.48) is compared with the exact values, calculated simulating *in silico* the stochastic

process and inferring the exact posterior distribution from the Bayes' formula (without any approximation). The model accurately predicts the simulated values and the power law dependence on the input properties r_i and σ_i^2 for small values of r_i .

Our approximation relies on assumption (3.5), which imposed an upper limit on r_i . For this reason, the model predictions diverges from the exact values for large values of r_i , as can be observed in 3.1C. Given the overall low firing rates of neurons in cortex (0.1 – 20 Hz) compared to their peak firing rates (60 – 300 Hz), we consider this approximation to hold for most scenarios. The linear dependence $\sigma_e^2 \propto c_t$ in (3.22) greatly simplify our evolution equation for c_t . In fact, by combining (3.48) with (3.16), we can write a closed form evolution

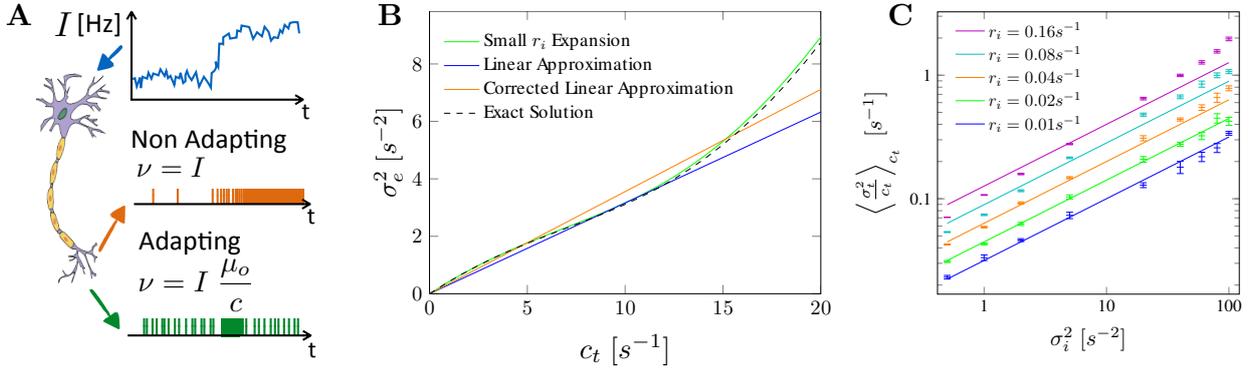


Figure 3.1: (A) Neuronal responses of a non-adapting (orange) and an ideally adapting (green) neuron driven by the same input. The adapting neuron represents unpredicted changes in the input instead of the input itself. (B) The σ_e^2 approximations of equation (3.21) (green line), (3.22) (blue line) and (3.48) (orange line) are compared with the exact solution in (3.20) (black dashed line). Parameters $\sigma_i^2 = 100s^{-2}$, $\mu_o = 10s^{-1}$, $\mu_i = 10s^{-1}$, $r_i = 0.01s^{-1}$. (C) The approximated expected value (continuous line) for the term $\langle \sigma_t^2 / c_t \rangle_{c_t} = 9/8 \sqrt{r_i \sigma_i^2 / \mu_o}$ (as in (3.48)), is compared with the simulation results (error bars), computed by simulating the stochastic system in silico and inferring the exact Bayesian posterior $D(\nu)$ (without any approximation). For large values of r_i , condition (3.5) is violated and, as expected, the simulation results diverge from the model predictions. Parameters used in the simulation are $\mu_o = 10s^{-1}$ and $\mu_i = 10s^{-1}$.

equation for the mean of the inferred posterior distribution about the instantaneous input c_t :

$$c_{t+\tau} = c_t + k_{\mu_i, \sigma_i^2, \mu_o, r_i} (1 - \mu_o \tau) + r_i \tau (\mu_i - c_t) \quad ; \quad k_{\mu_i, \sigma_i^2, \mu_o, r_i} = \frac{9}{8} \sqrt{\frac{r_i \sigma_i^2}{\mu_o}} \quad (3.49)$$

Equation (3.49) describes how the inference about the mean of the future input c (made by an ideal observer) evolves according to the observed inter-spike intervals τ_* of our adapting neuron. The given equation, expressed in (3.49) in form of finite-differences, can be also rewritten in form of differential equation as

$$\frac{dc_t}{dt} = k_{\mu_i, \sigma_i^2, \mu_o, r_i} (\delta(t - t_{\text{spike}}) - \mu_o) + r_i (\mu_i - c_t) \quad ; \quad k_{\mu_i, \sigma_i^2, \mu_o, r_i} = \frac{9}{8} \sqrt{\frac{r_i \sigma_i^2}{\mu_o}} \quad (3.50)$$

Crucially, this is a relatively simple first order differential equation, and could potentially be computed by simple biological mechanisms (linear chemical reaction kinetics) in neurons. In fact, c could potentially represent a biophysical quantity of the neuron internal dynamics, e.g. the concentration of some ion or neuro-transmitter, which acts on the neural response to the stimuli by scaling its spiking response rate. This action may take place both in the synaptic connections, the dendrite or where the spike initiation process takes place at the axon hillock. Additionally, we can predict that the concentration c of this putative physical quantity will follow the following physiologically testable dynamics: whenever perturbed it converges exponentially to its equilibrium value with time scale r_i . Moreover, whenever a spike takes places the concentration increases by value $\frac{K_1}{c}$, while a continuous leakage brings the concentration down with a rate $\frac{K_2}{c}$.

3.2.6 Local, on-line learning of optimal model parameters

As mentioned, all the parameters in equation (3.49) depend on the dynamical properties of the encoded stimulus. This means that, to produce optimal adaptation, the neuron needs to

fine-tune its adaptation properties to match the time structure of its input. Specifically, r_i corresponds to the inverse of the time scale on which the input ν_t changes, μ_i represents the average input value and σ_i^2 is the variability of the input ν_t , i.e. a measure of the range of variation of the input. Finally, μ_o is the desired neural firing rate that the adaptation process tries to maintain in order to maximize information transfer.

We propose a model that enables a neuron to learn the optimal adaptation parameters locally and online, without the need of any external supervision and exploiting neuron's spiking activity. This allows the neuron to simultaneously solve two different tasks: a) The neuron learns the optimal adaptation parameters from the time structure of the input (long time scale). b) It uses the learnt parameters to adapt optimally to the local statistics of the input (short time scale) in order to encode information efficiently.

In order to derive how the model parameters should evolve, we can write

$$c_{t+\tau} = c_t + \hat{k}_t (1 - \mu_o \tau) + \hat{r}_t \tau (\hat{\mu}_t - c_t) \quad (3.51)$$

where $(\hat{k}, \hat{r}, \hat{\mu})$ are now time dependent parameters. Their dynamics is regulated by evolution equations that, ideally, converges to optimal values $(k_{\mu_i, \sigma_i^2, \mu_o, r_i}, r_i, \mu_i)$. These optimal parameters result in the neural adaptation to be optimal given input dynamics: In order to maximize information transfer to optimally encode information, the aim of neural adaptation is to maximise the entropy of the inter spike interval distribution in every moment in time. This implies that each single spike event τ_i needs to be independent from the past spiking history of the neuron, encoded both in the variables c_t and $\tau_{j < i}$. In the following, we will show how we can exploit this property in order to derive internal dynamics for the parameters (k, r_i, μ_i) that results in them converging to the optimal values. In particular, we will exploit

the conditions

$$\langle c \rangle = \langle \nu_i \rangle = \mu_i = \hat{\mu} \quad (3.52)$$

$$\text{Corr}(\tau, c_{t+\tau} - \hat{\mu}) = 0 \quad (3.53)$$

$$\text{Corr}(\tau_i, \tau_{i-1}) = 0 \quad (3.54)$$

where Corr is the correlation function, $\tau = \tau_i$ and τ_{i-1} represent, respectively, the last and the next to last observed inter-spike intervals. First, it is easy to show that condition (3.52), is satisfied by the evolution equation

$$\hat{\mu}_{t+\tau} = \hat{\mu}_t + \alpha_\mu (c_t - \hat{\mu}_t) \quad (3.55)$$

where the parameter α_μ controls the velocity at which $\hat{\mu}$ is converging to the optimal value μ_i . Analogously, we can consider condition (3.53) and the properties

$$\hat{r}_t > r_i \quad \Rightarrow \quad \text{Corr}(\tau, c_t - \hat{\mu}_t) < 0 \quad \Rightarrow \quad \langle \tau(c_t - \hat{\mu}_t) \rangle_t < 0 \quad (3.56)$$

to derive the evolution equation for \hat{r}

$$\hat{r}_{t+\tau} = \hat{r}_t + \alpha_r (\tau(c_t - \hat{\mu}_t)) \quad (3.57)$$

where, again, α_r controls the convergence velocity of \hat{r} to r_i . Finally, we can consider condition (3.54) and the properties

$$\hat{k}_t > \frac{9}{8} \sqrt{\frac{r_i \sigma_i^2}{\mu_o}} \quad \Rightarrow \quad \text{Corr}(\tau_i, \tau_{i-1}) < 0 \quad \Rightarrow \quad \langle \tau_i \tau_{i-1} - \mu_o^2 \rangle_t < 0 \quad (3.58)$$

to derive the evolution equation for \hat{k}

$$\hat{k}_{t+\tau_i} = \hat{k}_t + \alpha_k (\tau_i \tau_{i-1} - \mu_o^2) \quad (3.59)$$

where α_k controls the convergence velocity of \hat{k} to $k_{\mu_i, \sigma_i^2, \mu_o, r_i}$. The evolution equations (3.55), (3.57) and (3.59) provide the internal parameters of the neuron ($\hat{k}, \hat{r}, \hat{\mu}$) to adapt to the dynamical properties of the input (r_i, μ_i, σ_i^2) and to converge to the optimal configuration. The proposed model, currently expressed in form of finite difference equations, can be rewritten as a linear system of 4 coupled first-order non-homogeneous differential equations with multiplicative delta function terms (Nedeljkov and Oberguggenberger, 2012)

$$\frac{dc_t}{dt} = \hat{k}_t (\delta(t - t_{\text{spike}}) - \mu_o) + \hat{r}_t (\hat{\mu}_t - c_t) \quad (3.60)$$

$$\frac{d\hat{\mu}_t}{dt} = \alpha_\mu \delta(t - t_{\text{spike}}) (c_t - \hat{\mu}_t) \quad (3.61)$$

$$\frac{d\hat{r}_t}{dt} = \alpha_r (c_t - \hat{\mu}_t) \quad (3.62)$$

$$\frac{d\hat{k}_t}{dt} = \alpha_k ((t_{\text{spike } i} - t_{\text{spike } i-1}) - \delta(t - t_{\text{spike}}) \mu_o^2) \quad (3.63)$$

where ‘‘spike i’’ is the last observed spike at time t , μ_o is the desired firing rate for the neuron, α_μ , α_r and α_k are parameters controlling parameters adaptation rates and c_t is the variable controlling neural adaptation (according to equation (3.1)). The given system can be potentially implemented by the internal dynamics of a biological neuron, where the parameters are stored in form of ions or neurotransmitters concentration.

3.2.7 Sanity check: Testing approximated model predictions

To test our model of neural adaptation, we compared the predictions of our adaptation dynamics with the results of optimal Bayesian computations, implemented in-silico by numerically computing the exact posterior distribution $D(\nu)$ and Bayes’ equation without any approximation. In particular, we implement a Markov Chain Monte Carlo process reproducing the stochastic dynamics of the input value ν_t and we simulated the spiking activity of the neuron by generating samples from a Poisson process with rate $\frac{\nu_t}{c_t}$, where c_t is the internal

adaptation parameter of the neuron. After every spike event, we adapted the parameter c_t according to equation (3.49). We used the evolution equations (3.55), (3.57) and (3.59) to adapt on-line the internal parameters $(\hat{\mu}, \hat{r}, \hat{k})$ to the input dynamical properties. Simultaneously, we computed the exact Bayesian posterior distribution $D(\nu)$ and we adapted this distribution after each spike event according to the exact Bayes' equation, assuming a complete knowledge of the input dynamical properties, of the neural spiking activity and of the time dependent parameter c_t regulating neural spiking activity. Finally, we compared the predictions of our adaptation model c_t with the average of the exact posterior distribution $D(\nu_t)$.

The model proved to be able to learn autonomously the optimal configuration for the internal parameters $(\hat{\mu}, \hat{r}, \hat{k})$ from the input dynamical properties (time-scale of weeks, see Fig. 3.2ABC) and to compute near-to-optimal inference of the posterior mean (time-scale of seconds, see Fig. 3.2F ; R^2 value between approximate and exact posterior mean consistently > 0.95). It is interesting to note that the approximated model, which captured most of the variability of the exact inference solution, necessitate much less computational power to be solved compared to the original inference problem, which requires to encode the full posterior distribution in each moment in time, and to update it in each iteration. Furthermore, the adapting neuron proved to be able to encode the value of the input drive I more accurately than the non adapting neuron, exploiting the same number of spikes (see Fig. 3.2F; MSE respectively $5s^{-2}$ and $11s^{-2}$).

3.2.8 Decoding

The proposed model of neural adaptation allows to efficiently encode information using a non-trivial coding scheme. The given coding scheme belongs to the category of predictive coding, since it uses spikes to encode only new non-redundant information. This increases entropy of the information channel (by reducing inter-spike-intervals correlation) and it maximizes

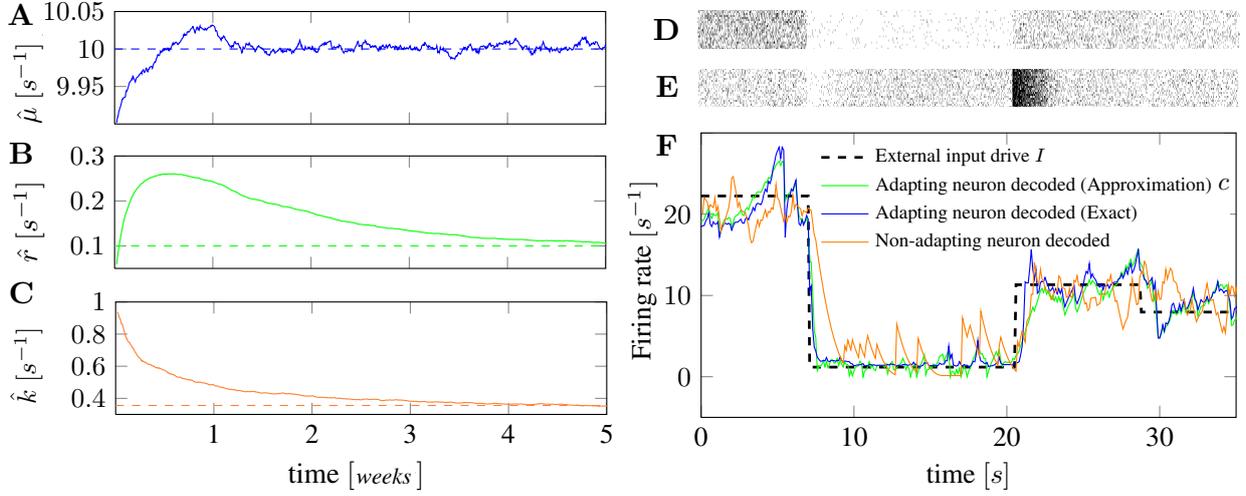


Figure 3.2: The neuron adapts internal parameters locally and on-line to the temporal structure of the input. Simulations results for evolution of hyperparameters $\hat{\mu}$ (A), \hat{r} (B) and \hat{k} (C), following equations (3.55), (3.57) and (3.59). The parameters (solid lines) consistently converge to the optimal values (dashed lines). Simulation parameters: $\mu_i = 10s^{-1}$, $r_i = 0.1s^{-1}$, $\sigma_i^2 = 10s^{-2}$, $\mu_0 = 10s^{-1}$. The time-scale of hyperparameters adaptation is compatible with experimental observations (Sernagor et al., 2001). (D) Raster plot of a non-adapting neuron and (E) of our ideal adapting neuron, encoding the input in fig. 3.2F. (F) The plot represents the encoded input drive (black dashed line), the optimally decoded non-adapting neuron (orange line)(MSE = $11s^{-2}$) and the optimally decoded adapting neuron (blue line)(MSE = $5s^{-2}$). It can be noted that the adapting neuron represents the input drive more accurately than the non-adapting one using the same number of total spikes. In addition, the approximated value c (green line) inferred from the activity of the adapting neuron using equation 3.49 ($R^2 = 0.96$) is compared with the exact Bayesian decoding (blue line), this further validating the used approximations. Simulation parameters $\mu_i = 10s^{-1}$, $r_i = 0.1s^{-1}$, $\sigma_i^2 = 10s^{-2}$, $\mu_0 = 10s^{-1}$. It can be noted that adaptation takes place on the time-scale of individual inter-spike intervals.

information encoded per spike event. But if, on one side, the efficiency of the proposed code has been tested against a non adapting encoding scheme (fig. 3.2), the practicability of its implementation in the brain depends on the ability of other neurons to decode its message efficiently. For this reason, we will here investigate how the information encoded by a pre-synaptic adapting neuron (according to the proposed model) can be efficiently decoded by a post-synaptic neuron.

Equation (3.49) describes how the message of neural spikes can be decoded. Coincidentally, the given equation defines a dynamics that matches the membrane potential of a Leaky Integrate and Fire neural model (Burkitt, 2006)

$$C_m \frac{dv(t)}{dt} = I_l(t) + I_\omega(t) ; \quad I_l = -\frac{C_m}{\tau_m}(v(t) - V_0) , \quad I_s = K_s \delta(t - t_s) \quad (3.64)$$

The equation describes the evolution of the membrane potential v , where I_ω represents the synaptic current driving the neuron, I_l is the leakage current and C_m is the membrane capacitance. If tuned to have the right parameters, the membrane potential of a leaky integrate and fire post-synaptic neuron will match the internal dynamics of the adaptation parameter c in the pre-synaptic neuron and, as a consequence, it will correspond to the hidden variable encoded by the spiking activity of the adapting neuron. Membrane fluctuations are then transformed into spikes by the Integrate and Fire spiking initiation mechanism: whenever the membrane potential reaches a certain threshold a spike event is generated, and the potential is re-set to the resting value. We can make this consistent with our model of neural adaptation by assuming that the adaptation variable c controls the probability of generating new spikes either by modulating the threshold for spike initiation, or by generating an additional flow of ions through the membrane, which affects the membrane potential (Smith et al., 2002). This would make the proposed model of neural adaptation consistent when implemented in series by successive neurons in the brain.

Being able to decode the adapting neuron enables us to compare its performance (in accurately encoding a hidden variable) with a non-adapting neuron. We encoded the same hidden variable using either an adapting (according to our model) or a non-adapting (inhomogeneous Poisson) neuron and, in both cases, we constrained the number of total spike events produced by the neuron (fig. 3.2). The adapting neuron, when decoded, revealed to be able to describe the hidden variable much more accurately than the non-adapting one by exploiting the same number of total spikes.

The Leaky Integrate and Fire model is widely used in the literature to describe neural dynamics (Moreno-Bote and Parga, 2006; Lindner et al., 2002). The fact that this simple model of membrane potential dynamics can optimally decode the efficient coding scheme we propose is a promising result. It supports the hypothesis that the proposed model of neural adaptation could be implemented in the brain to increase encoding efficiency, and it potentially provide a new interpretation for the computational role underlying the Leaky Integrate and Fire model.

3.3 Testable predictions

The model of neural adaptation we propose provides various quantitative predictions, that link properties of adaptation to the time structure of the encoded stimulus. For example, the model links the timescale of stimulus changes and the mean and variance of the prior distribution stimulus to the timescale adaptation, amplitude and equilibrium point of adaptation. Moreover, the proposed model quantitatively predicts neural level of adaptation following a certain spike train history, and how adaptation parameters are modified when the time structure of the stimulus changes (couplings between adaptation parameters generates non-trivial dynamics in converging to the new equilibrium solution).

Furthermore, the given model can be used to predict the average response (over various stochastic repetitions) of an adapting neuron encoding a particular trajectory of the hidden variable. In the following, we will mathematically derive the time course of the firing rate of a neuron when the encoded latent variable is a Heaviside theta, and we will compare our model predictions with experimental observations.

3.3.1 Predicting the time course of neural spiking rate

We can use equation (3.49) to derive some quantitative predictions about firing rate adaptation following a jump in the neural driving input. An example of this is the time course of the firing rate of a retinal neuron when we exit a dark gallery on a sunny day. The neuron, adapted to a low light regime, will first present a very high firing rate. However, after some seconds, the sustained high firing rate of the neuron will allow it to infer that the luminance of the environment has changed. The neuron will adapt to the new regime, and the spiking activity of the neuron will slowly decay to a lower, efficient and sustainable value. We can use our model of neural adaptation to predict properties (time scale and shape) of this time course (Ticchi and Faisal, 2015). Specifically, we will predict the time course of the firing rate of a neuron after a jump, averaged over different stochastic repetitions. We will do it using the temporal properties of the stimulus, without introducing any free parameter.

We first assume that, before the input change, the neuron is in a steady state (i.e. $\langle c_{t+\tau} \rangle = \langle c_t \rangle$). We will denote the equilibrium value of c before the input jump with c_b . Under the steady state hypothesis and using eq. (3.49), it is possible to show that the value c_b can be expressed as a function of the average inter spike interval at the equilibrium τ_b :

$$c_b = k \frac{1 - \mu_o \tau_b}{r_i \tau_b} + \mu_i \quad (3.65)$$

Additionally, by definition, if we identify the value of the drive pre-jump with the variable D_b , we can write

$$\tau_b = \frac{c_b}{D_b} \tau_o \quad (3.66)$$

where τ_o is the inverse of the desired neural firing rate μ_o , as determined by metabolic and biophysical constraints of the neuron. By combining equations 3.65 and 3.66, it is possible to show that the value of c_b can be expressed as a function of the input D_b by solving the

following equation:

$$c_b^2 \frac{\tau_o r_i}{k} + c_b \tau_o \left(\mu_o - \frac{r_i \mu_i}{k} \right) - D_b = 0 \quad (3.67)$$

It is interesting to note that, in the considered regime (equation (3.5)), we have

$$r_i \ll \mu_o = \tau_o^{-1} \quad (3.68)$$

In this situation, the solution of equation 3.67 becomes $c_b = D_b$ which, according to eq. 3.66, translates into $\tau_b = \tau_o$. The given approximation correspond to a stimulus which changes on a very long timescale. As a consequence, the neuron is able to accurately infer the value of the input before next jump, and the inference process is very accurate.

We can now focus on what happens after a jump in the input drive D , which set it to a new value D_a . Right after the jump, the instantaneous average interspike interval will be

$$\tau_a = \frac{c}{D_a} \tau_o = \frac{D_b}{D_a} \tau_o \quad (3.69)$$

In this situation, eq. 3.49 will no longer be at the equilibrium, and will start converging to a new value of c ($c \rightarrow c_a$). Do investigate the time course of c , we can simplify equation 3.49 in the regime of eq. 3.68

$$c_{t+\tau} = c_t + k_{\mu_i, \sigma_i^2, \mu_o, r_i} (1 - \mu_o \tau) \quad ; \quad k_{\mu_i, \sigma_i^2, \mu_o, r_i} = \frac{9}{8} \sqrt{\frac{r_i \sigma_i^2}{\mu_o}} \quad (3.70)$$

Exploiting the finite difference approximation for the time derivative, under the assumption of small average inter spike interval τ_t we can write

$$\tau_t \partial_t c_t = k (1 - \mu_o \tau_t) \quad (3.71)$$

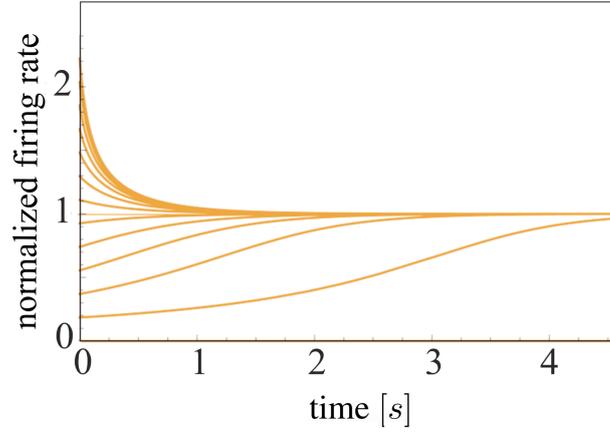


Figure 3.3: The plot visualizes solutions of equation (3.78) for different initial conditions ($f_{t=0} = \tau_o D_b/D_a$). Parameters used are $\frac{9}{8} \frac{1}{D_a} \frac{\sigma_i}{\mu_i} \sqrt{r_i \mu_o} = 2$.

and, using eq. 3.69, we get

$$\tau_t \partial_t \tau_t = \frac{\tau_o}{D_a} k (1 - \mu_o \tau_t) \quad (3.72)$$

In order to transform our equation for the average inter-spike interval τ_t into an equation for the firing rate $\nu_t = \frac{1}{\tau_t}$, we use the relation

$$\tau_t \partial_t \tau_t = -\frac{1}{\nu_t^3} \partial_t \nu_t \quad (3.73)$$

to write an equation describing the timecourse of ν_t

$$\partial_t \nu_t = \nu_t^2 \frac{k}{\mu_o D_a} (\mu_o - \nu_t) \quad (3.74)$$

The given equation describes the time course of the instantaneous firing rate after a jump (either positive or negative) in the input drive (see fig 3.3). The equation depends on some properties of the stimulus encoded k , on the desired firing rate of the neuron μ_o and on D_a , which is the input drive currently experienced by the neuron.

3.3.2 Linear transformation in the input space

We expect the considered model to be invariant under linear coordinate transformations in the space of the input drive D (as long as the internal neural variable c is changed accordingly). To make this property more explicit, we can express D_a as follows

$$D_a = \mu_i \hat{D}_a \quad (3.75)$$

where μ_i is the time-average of the input drive, and \hat{D}_a is a variable with unitary mean, which describes the instantaneous value of the input as compared to its average value. By doing this, every linear coordinate transformation will be reflected in μ_i and not to \hat{D}_a .

We can now expand k in eq. 3.74, and get

$$\partial_t \nu_t = \nu_t^2 \frac{9}{8} \frac{1}{\hat{D}_a} \frac{\sigma_i}{\mu_i} \sqrt{\frac{r_i}{\mu_o^3}} (\mu_o - \nu_t) \quad (3.76)$$

It is possible to note that our final equation depends now only on the ratio between the standard deviation of the input drive and its mean ($\frac{\sigma_i}{\mu_i}$). This quantity is invariant under linear transformation in the space of the input drive and, consequently, our model is invariant under these transformations, as expected. Other quantities appearing in the given equation are \hat{D}_a (instantaneous input drive - relative to the mean value μ_i), r_i (rate on which the input drive change), μ_o (desired firing rate of the neuron).

3.3.3 Linear transformation in the output space

We can investigate how the model is affected by a linear transformation in the output space (neural firing rate). We express the instantaneous firing rate ν_t as follows

$$\nu_t = \mu_o \hat{\nu}_t \quad (3.77)$$

where μ_o is the desired firing rate. By doing this, we can write

$$\partial_t \hat{\nu} = \hat{\nu}_t^2 \frac{9}{8} \frac{1}{\hat{D}_a} \frac{\sigma_i}{\mu_i} \sqrt{r_i \mu_o} (1 - \hat{\nu}_t) \quad (3.78)$$

It is possible to note that μ_o doesn't completely disappear from the equation. This means that the model is not invariant to linear transformations in the output (neural firing rate) space. The reason is that adaptation is driven by inference, and inter-spike intervals carry the information used by the inference process. For this reason, the higher neural firing rate is, the faster inference and adaptation will be implemented.

3.4 Results

To experimentally test our model predictions (equation (3.78)), we collaborated with Matteo Carandini (University College London, UK) and Andrea Benucci (Riken Brain Science Institute, Japan), who generously shared with us their data from Benucci et al. (2013). The database consists of in-vivo recordings of the response of a population of cat V1 neurons to a stimulus with random orientations. We analyzed the given database to extract the time-course of neural adaptation following a jump (Heaviside theta) in the stimulus, and we compared our findings with the predictions of our model. In order to fit the parameters of the model, we exploited the known dynamical properties of the stimulus used during the experiment. Therefore, our predictions are achieved without fitting any free parameter of the

model to the data.

3.4.1 Experimental setting

A 10x10 electrode array was used to record the spiking activity from V1 of anesthetized cats. Electrode sites were spaced 400 μ m apart, and activity of neurons was recorded for 6 consecutive seconds in each experiment, for a total of 7 experiments. The temporal resolution of the recording was 8.3 ms. Neural responses were characterized using sequences of static gratings with random orientation and phase, each presented for 32 ms. Axis of preferred orientation were then divided evenly into 12 bins, each pooling the activity of neurons with similar orientation preferences. Responses of each bin (neural pool) were then normalized (see figure 3.4).

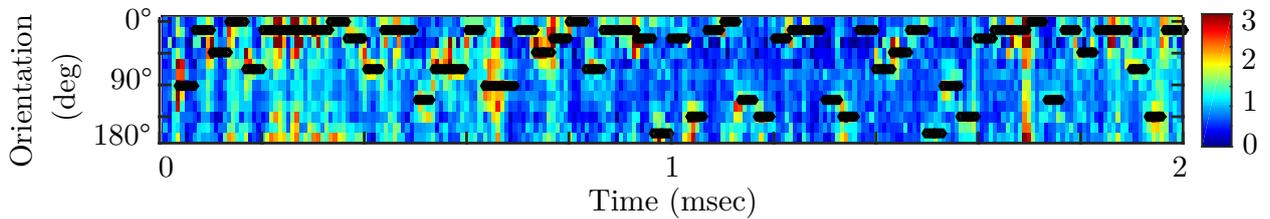


Figure 3.4: Visualization of the data we used to test our model, which was generously shared with us by the authors of (Benucci et al., 2013). Color code represents activation level of different neural pools from cat V1 in response to the stimulus provided. Responses are scaled to a value of 1 at the preferred orientation and 0 at the orthogonal orientation. Neural pools are ordered on the y-axis according to the center of their receptive fields (see fig.3.5). The value of the stimulus provided (grid orientation) is represented by a black line.

3.4.2 Data analysis

By applying regression to these responses, the tuning curves for each of the bins of preferred orientation were extracted (see fig. 3.5).

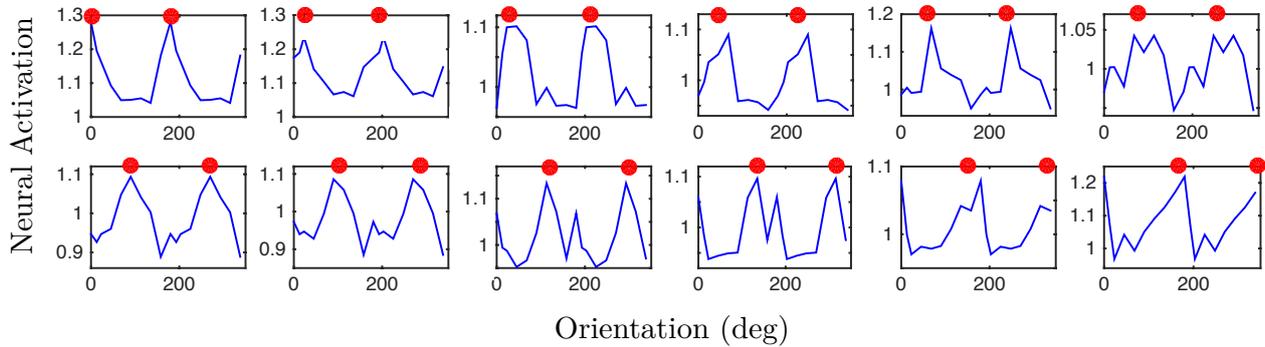


Figure 3.5: The figure represent the receptive field of cat V1 neural pools extracted from data in fig.3.4. Red dots represent the center of the respective tuning curves. The grids used as stimulus in the experiment are non-directional and, as a consequence, neural tuning curves are invariant for 180 rotations. Neural pools are ordered according to the position of their receptive field centers.

To extract the time-course of adaptation from the given database, we considered (for each neural pool) a window of orientations centered at the respective tuning curve center, and with variable width. When the stimulus is outside the relevant region, the pool activity is low, and adaptation acts to increase spiking activity. If, in this regime, the stimulus changes to a new value within the relevant region, activity will suddenly rise, induced both the highly driving stimulus and by the low inherited level of adaptation. At the same time, neural adaptation gets into play and pool activation starts decreasing towards a sustainable regime for the neurons, following a certain time-course. Our aim is to compare this time-course with the time-course predicted by our model of neural adaptation. To do this, each time the stimulus entered inside the relevant region we started recording the trajectory of pool activation, and we recorded as long as the stimulus was inside the given region. We repeated the same process for all the neurons (fig. 3.6). We finally took all the trajectories from all different neurons and we calculated their mean and standard error as a function of time. Different trajectories can have different lengths (according to how long the stimulus stays inside the relevant region), and at each moment we averaged only the trajectories longer than the considered time.

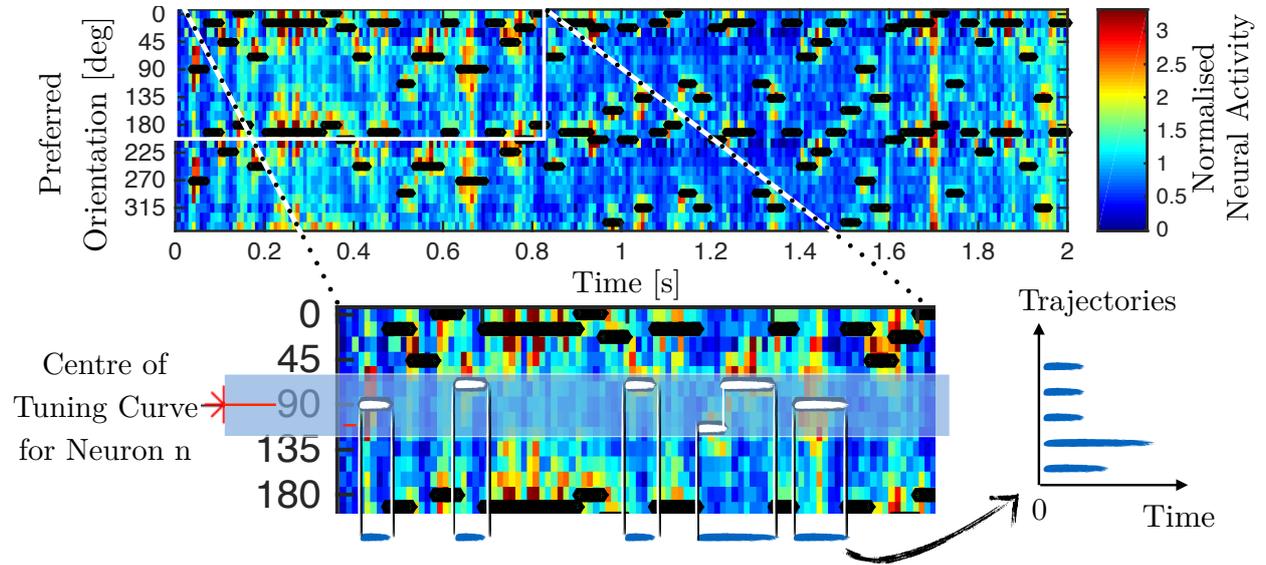


Figure 3.6: Visual representation of data processing used to extract relevant information from the given dataset. For each neuron, we considered a relevant window (centered either at the tuning curve center or at the orthogonal orientation). We then recorded all the neural activation trajectories, starting when the stimulus entered in the relevant window and ending when it first exited the given region. We repeated the same process to collect trajectories for all the neural pools. We superimposed all the trajectories by synchronizing the initial times (when the stimulus entered the relevant window). We then calculated mean and standard error of all the trajectories in each moment in time, to extract the average time-course of adaptation.

By doing this, we extracted the average neural activation following a jump of the stimulus, from a region far from the receptive field center to a region close to it. By combining different trajectories, we discarded information about the specific amplitude of each stimulus jump preceding adaptation. Nevertheless, we decided to average all the trajectories together to increase statistical significance, by relying on the fact that the superposition of solutions of our model equation (3.74) will be well approximated by a solution of the same equation. Moreover, the stimulus could assume different values while remaining inside the relevant region. Yet, we discarded these changes, relying on the fact that positive and negative jumps within the relevant region will statistically sum to zero.

3.4.3 Validating model predictions on experimental results

We repeated the same procedure utilizing different relevant windows. We used different windows widths (specifically 24, 48 and 72) and we considered either a window centered at the center of the neural pool's tuning curve, or the complement region. By doing this, we investigated the average time-course of adaptation following either positive or negative jumps of the neural input drive (according to where we centered the relevant window), as well as jumps of different amplitude (according to the varying window width).

For each tested condition (width and window center) we compared the extracted time-course to the analytical predictions from our normative model (fig. 3.7). Since our model relates properties of neural adaptation to the time-structure of the encoded stimulus, we used the dynamical properties of the stimulus used in the experiment to fix all the parameters of our adaptation equation (3.74). Specifically, we used the expanded form given in 3.78:

$$\partial_t \nu_t = \nu_t^2 \frac{9}{8} \frac{1}{\hat{D}_a} \frac{\sigma_i}{\mu_i} \sqrt{\frac{r_i}{\mu_o^3}} (\mu_o - \nu_t) \quad (3.79)$$

and we set r_i to 31.25 s^{-1} , which is the inverse of the interval used in the experiment to update the grid orientation (32 ms). We extracted the value of the ratio between standard deviation and mean of the input drive ($\frac{\sigma_i}{\mu_i}$) from the same ratio calculated over the neural response (0.1). We did this by exploiting the fact that the given ratio is preserved for linear transformations over the input space, and that the effect of adaptation over this ratio should be negligible given the random nature of the stimulus jumps. We made various in-silico simulations to verify the validity of the given approximation. We set the value of desired firing rate of the neuron μ_o to the average firing rate of cat V1 neurons (10 s^{-1}) (Legendy and Salcman, 1985). Finally, We approximated the value of the normalized instantaneous input drive \hat{D}_a to be constant (1), being impossible to extract the value from the experimental setting. Furthermore, since the ratio between standard deviation and mean of the input drive ($\frac{\sigma_i}{\mu_i}$) was estimated to be around (0.1), the instantaneous value of the normalized input

drive \hat{D}_a would not deviate more than 10% from 1 in most cases.

Using the given parameters, we solved the finite-difference equation 3.79 to derive the time-course of adaptation (following jumps in the input drive with different amplitudes and directions), and we compared our model predictions with the experimental results in fig. 3.7.

As we can see from fig. 3.7, analytical predictions from our model are compatible with experimental results in all the considered scenarios (except from 3.7 D, where the very low input drive received by the neuron generates a pathological behavior related to network sustained dynamics). Specifically, we can see that solutions of equation 3.79 for different initial conditions match both the shape and the timescale of experimental adaptation trajectories.

Moreover, fig. 3.7 shows that model predictions capture experimental time-course of adaptation better than exponential fits (the shared decay constant of the exponential fits was selected as to best predict the set of 6 different conditions. This finding is supported by a consistently higher variability of the data explained by the model predictions (generalized R^2 in fig. 3.7 condition A=0.57, B=0.87, C=0.64, E=0.61, F=0.79) compared to the exponential fits (generalized R^2 in fig. 3.7 condition A=0.24, B=0.77, C=0.64, E=0.54, F=0.71). Specifically, the presence of the term ν^2 in equation 3.78 provides the solution of the equation to decay faster for large starting values, and slower for small ones. This is visible in fig. 3.3, where it is possible to note an asymmetry of the solutions for normalized starting values larger or smaller than one. The given property matches experimental results, where the decay following a positive jump in the input drive has a smaller time-scale than the ones corresponding to a negative jump. This supports the fact that our principled model captures some non-trivial properties of neural adaptation time-course, which are not captured by other popular phenomenological models like the exponential decay, often used in the literature to describe adaptation dynamics (Dean et al., 2008; Westerman and Smith, 1984).

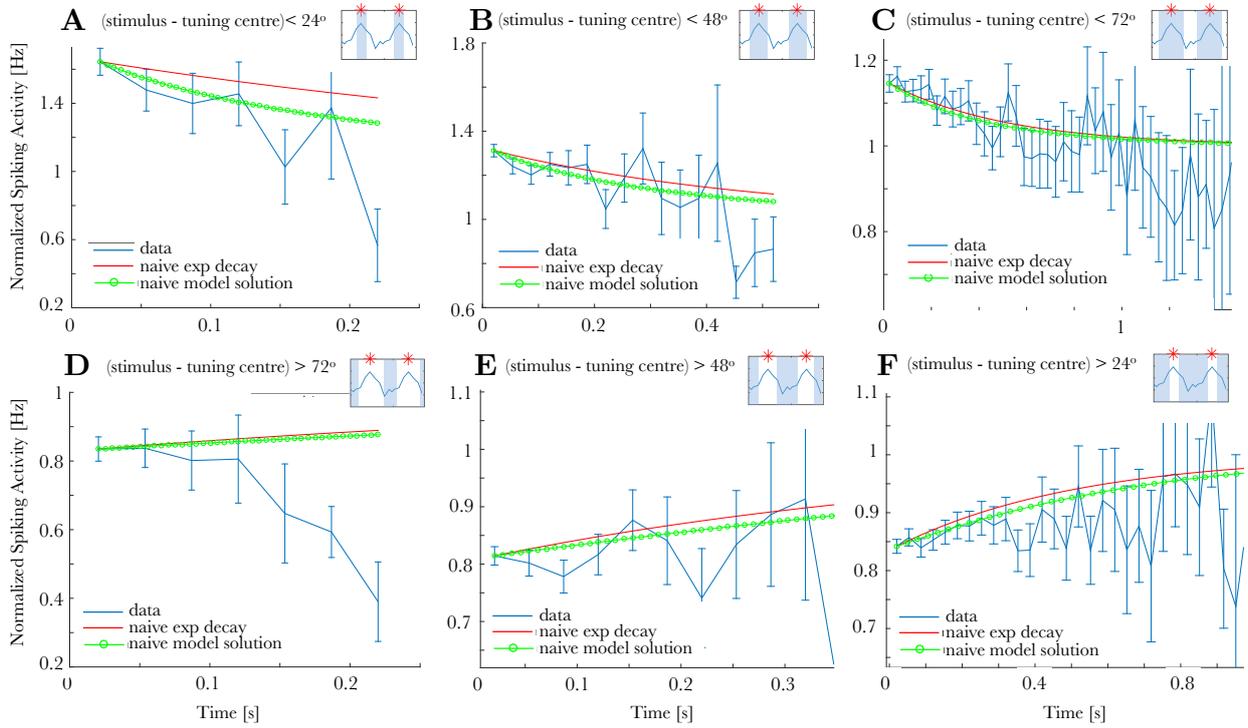


Figure 3.7: Time-course of adaptation (normalized neural activity) from in vivo recordings of cat V1 (blue error bars) (Benucci et al., 2013) is compared with analytical predictions from our model (green line) (equation (3.78), fig. 3.3). Red lines represent exponential fits, in which the shared decay constant was chosen as to best predict the set of 6 different conditions. Different plots refer to different relevant window sizes and locations (as described in the top-right diagrams). All the parameters in the analytical model were fixed using dynamical properties (timescale of change, distribution and dynamical range) of the stimulus used in the experimental setting. Plots A, B and C represent time-course of adaptation relative to a relevant window centered at the tuning curve center of the neural pools. Plots D, E and F refer to a relevant window centered at the orthogonal orientation. We can note that in plots A, B, C, E and F analytical predictions match both timescale and shape of adaptation (generalized R^2 A=0.57, B=0.87, C=0.64, E=0.61, F=0.79). In addition model predictions explain experimental data consistently better than the exponential fit (generalized R^2 A=0.24, B=0.77, C=0.64, E=0.54, F=0.71), which is often used in the literature to fit adaptation (Dean et al., 2008; Westerman and Smith, 1984). Experimental results in plot D show a non-adaptive behavior, which differs from the predictions of our model. In the discussed plot, the relevant window is very small and centered at the orthogonal orientation. As a consequence, the neuron receives a very low input drive and resilient network effects become dominant, introducing an inertial component to neural activity decay.

3.5 Biophysical implementation

The proposed model of neural adaptation (3.2.6) consists of a linear system of 4 coupled first-order non-homogeneous differential equations with multiplicative delta function terms (Nedeljkov and Oberguggenberger, 2012)

$$\frac{dc_t}{dt} = \hat{k}_t (\delta(t - t_{\text{spike}}) - \mu_o) + \hat{r}_t (\hat{\mu}_t - c_t) \quad (3.80)$$

$$\frac{\hat{\mu}_t}{dt} = \alpha_\mu \delta(t - t_{\text{spike}}) (c_t - \hat{\mu}_t) \quad (3.81)$$

$$\frac{\hat{r}_t}{dt} = \alpha_r (c_t - \hat{\mu}_t) \quad (3.82)$$

$$\frac{\hat{k}_t}{dt} = \alpha_k ((t_{\text{spike } i} - t_{\text{spike } i-1}) - \delta(t - t_{\text{spike}}) \mu_o^2) \quad (3.83)$$

where “spike i” is the last observed spike at time t , μ_o is the desired firing rate for the neuron, α_μ , α_r and α_k are parameters controlling parameters adaptation rates and c_t is the variable controlling neural adaptation according to equation (3.1)

$$\nu_t = I_t \frac{\mu_o}{c_t} \quad (3.84)$$

where I_t is the instantaneous input drive received by the neuron and ν_t is the resulting spiking rate after adaptation (or the instantaneous probability density of generating a spike event). The proposed system of differential equations can be potentially implemented by the internal dynamics of a biophysical neuron exploiting the flow of ions through the membrane potential (controlled by ion channels). To identify molecular candidates for the equation variables, we can follow two different approaches, which reflect two different interpretations of neural dynamics.

A) We consider the spiking activity of the neuron to be the message communicated from one neuron to another. This is the main interpretation used in the literature (Maass, 2002; Averbeck et al., 2006), according to which spike events are considered to be the paramount carrier of information inside the brain. In this case, the message transmitted by a pre-synaptic neuron needs to be decoded by a post-synaptic neuron (Warland and Reinagel, 1997; Arabzadeh and Panzeri, 2006). However, in order to understand the received message, post-synaptic neurons need to be aware of the neural code used and, therefore, of the adaptation parameters controlling pre-synaptic neuron dynamics ($\hat{\mu}$, \hat{r} and \hat{k}). Adaptation parameters are continuously changing, and they can't be inferred only by observing the spiking activity of the neuron as a consequence of the nature of the code implemented (which maximize entropy of the information channel and removes correlation). Therefore, this would lead to an inconsistent scenario, where neurons are communicating an encoded message without providing to other neurons the code to understand it and use it for computations. We could assume that these parameters are transmitted through the synapses using chemical messengers, but this appears to us as an unlikely and unreliable solution considering what we know about the biophysics of synapses (Zador et al., 1990; Cohen-Cory, 2002) and the strong dependence of the message encoded to the adaptation parameters.

B) We interpret neuron action potential as an internal communication channel, which allows propagating information through the neuron. This information is then decoded locally inside the neuron, and only the decoded variable is communicated to other neurons. According to this interpretation, spiking activity of the neuron encodes information using a complex coding scheme that maximizes information per spike event (increasing metabolic efficiency of the communication channel). The message encoded in spikes can not be decoded by an external observer if the neural coding parameters (which depends on the dynamical properties of the stimulus) are not known. Yet, this information is represented inside the neuron in an explicit form (c_t), and this explicit representation is communicated to other neurons through synapses.

If we assume the validity of interpretation (B), we can look for a molecular candidate for the adaptation variable c_t by investigating which quantity in the neuron controls synaptic communication. It turns out that in various Central Nervous System (CNS) neurons, the release of vesicles of neurotransmitters into the synaptic cleft has been shown to be controlled by the concentration of calcium at the axonal termination (Schneppenburger and Neher, 2000). This makes calcium a plausible candidate for implementing the role of c_t inside the neuron. However, in order to be seriously considered, a molecular implementation of c_t should satisfy five additional requirements:

1. It should **influence the response function of the neuron** by modulating the probability of generating new spikes (introducing neural adaptation). Coincidentally, calcium has been observed to control neural response (gain normalization) by modulating calcium-dependent potassium currents at the synaptic level in different regions of the brain (e.g. cerebellar nucleus (Feng and Jaeger, 2008; Feng et al., 2013) and vestibular nucleus (Smith et al., 2002)).
2. Its dynamics should **be consistent in different regions of the neuron**, and its concentration should be proportional (yet possibly scaled) across all the neuron and in each moment in time. Again, intra-neuron calcium concentration has been observed to satisfy this requirement (Schiller et al., 1997; Augustine et al., 2003).
3. The candidate quantity should **be compatible with the derived dynamics** (equation (3.50)), which controls the evolution of the adaptation parameter according to the recent spiking history of the neuron. The given equation corresponds to the dynamics of a leaky integrator and, for an appropriate choice of the adaptation parameters, it is compatible with the known time-course of calcium concentration inside the neuron (Augustine et al., 2003; Loewenstein and Sompolinsky, 2003; Schiller et al., 1997).
4. The candidate quantity should **correlate with the stimulus encoded** by the neuron and, as a consequence, with animal behavior. It has been shown in the literature that

calcium concentration in neurons can track the stimulus encoded by the neuron (Kato et al., 2014a) and drive behavior (e.g. eliciting food acceptance or avoidance response (Marella et al., 2006)).

5. Altering the dynamics of the candidate quantity encoding external world inside the neurons should **affect animal behavior**. In this respect, it has been observed experimentally that calcium channel mutations can alter synaptic release of neurotransmitters (Ayata et al., 1999) and influence animal behavior (Kato et al., 2014a).

Therefore, calcium appears as a natural candidate for representing the adaptation variable c_t inside the neuron, and it is compatible with all the requirements predicted by our model for this quantity. In our model of neural adaptation, c_t is the variable encoding the state of the world inside the neuron. Identifying calcium with c_t can provide a new perspective on the computational role of calcium and spikes in the working brain. According to this new perspective, calcium is the quantity controlling local neural processing, and encoding relevant features of the world. On the contrary, spikes only serve to encode new available information (received by other neurons and not yet encoded by c_t) and to propagate this information to different regions of the neuron in order to improve the representation of the world by c_t . This new picture of the working neuron has consequences also on the interpretation of other neural processes, such as on the computational role of synapses, membrane potential fluctuations, spike generation mechanisms and ions channels dynamics, and on how they are controlled by calcium, spikes and pre-synaptic inputs. In Chapter 4 we explore in more details this new interpretation of neural dynamics, which is inspired and motivated by our results on neural adaptation.

3.6 Discussion

Starting from a generative model of the stimulus stochastic dynamics, we used Bayesian Statistics and Information Theory to derive from first principles a normative model of how a neuron should optimally adapt in order to encode information efficiently. Under some biophysically motivated assumptions, we show that the resulting model can be expressed in form of a first-order linear non-homogeneous ordinary differential equation, and that it is consistent for various generative models of the stimulus dynamics. The proposed model links properties of adaptation (e.g. time-scale, amplitude and equilibrium point) to the time structure of the encoded stimulus (e.g. dynamical range and time-scale of changes). In this way, we derive from first principles quantitative predictions that link neural dynamics and properties of the stimulus encoded.

The stimulus encoded can evolve its structure with time (for example luminance statistics varies in different season), and for this reason we expanded our model to enable the neuron to learn the optimal adaptation parameters, and to update their value on-line to the experienced stimulus. The resulting model can be expressed in form of four coupled first order non-linear differential equations: one equation controlling the neural gain and three equations controlling the evolution of the adaptation parameters, which provide these parameters to converge to the optimal values for a given time structure of the stimulus. The proposed model provides the neuron both to adapt to the encoded hidden variable on a short time-scale (ms/s) and to tune its adaptation behavior on-line to the dynamical properties of the experienced stimulus on a longer time-scale (min/day).

We also show that the resulting coding scheme (predictive coding) can be decoded by a post-synaptic neuron. Coincidentally, the neural dynamics required to decode the adapting neuron is compatible with known effective models of neural dynamics (leaky integrate and fire)(Burkitt, 2006).

As discussed in 3.3.1, our model predicts the average (over various stochastic repetitions) time course of neural spike-rate adaptation for a particular trajectory of the encoded stimulus. Model predictions revealed to be compatible with neural recordings from cat V1 (fig. 3.7), and they captured experimental deviations of firing rate time-course from an exponential fit (often used in the literature to model spiking rate adaptation (Drew and Abbott, 2006)). Furthermore, for a given spiking history and stimulus time structure, our model quantitatively predicts the instantaneous level of neural adaptation and neural firing probability. Additionally, the proposed model allows relating properties of adaptation (i.e. time-scale of adaptation, adaptation amplitude and adaptation equilibrium value) to properties of the encoded stimulus (i.e. timescale on which stimulus change and mean-standard deviation of stimulus statistics). Finally, we predict how adaptation parameters respond to a change in the time structure of the encoded stimulus. Specifically, the predicted time course for the adaptation parameters is often non-trivial, as a consequence of couplings in the flow equations controlling the evolution of different parameters (fig. 3.2). All these predictions can be tested experimentally, and specific experiments can be designed to further validate the given model, and the underlying assumptions about the computational role of neural adaptation.

The proposed mechanism of neural adaptation can be implemented by a biophysical neuron exploiting the flow of ions through the membrane potential (controlled by ion channels), and the given model can be used to identify candidate signaling pathways for neural adaptation within a neuron. Identifying biophysical candidates for the mathematical quantities predicted by our system of coupled differential equations would represent a further validation of the proposed adaptation mechanism, and it would provide a deeper understanding of the computational role of specific molecules and processes within neural dynamics.

In Section 3.5 we saw that calcium appears as a natural candidate for encoding the adaptation variable c_t inside the neuron, and that it seems to satisfy all the requirements predicted by our model for such a quantity. This suggests that calcium might play an important role inside the neuron, and it would imply a new interpretation of the role of calcium and Spikes

in brain dynamics. According to this new interpretation, calcium would be the quantity encoding the state of the world inside the brain, while spikes would only be used to internally synchronize the concentration of calcium inside the neuron (using a special code that can only be decoded efficiently by the neuron itself). This would have direct consequences also on the interpretation of other aspects of neural dynamics, and specifically on the computational role of synapses, membrane potential fluctuations, spike generation mechanisms and ions channels dynamics.

In Chapter 4, we will explore a new model of neural dynamics inspired by these findings, and by the consequent new interpretation of the role of calcium and spikes in the brain. We will develop a normative model for the dynamics of synapses, membrane potential, spike initiation and calcium concentration inside the neuron, and we will show how these different processes can work together to encode and communicate information efficiently. In this new picture, spike-rate adaptation (as discussed in this Chapter) will rise as a consequence (or side effect) of the proposed neural dynamics.

Chapter 4

Calcium Coding Theory

‘... living matter, while not eluding the “laws of physics” as established up to date, is likely to involve “other laws of physics” hitherto unknown, which, however, once they have been revealed, will form just as integral a part of this science as the former.’

Erwin Schrödinger, “What is Life?”, pp 68

4.1 Introduction

Spikes are traditionally considered as the key representation of information inside the brain (Kepecs and Lisman, 2003). Spiking activity of neurons is extensively used in the literature to extract receptive fields and tuning curves of neurons in different brain regions, investigating what these areas are encoding about the state of the world (Cox et al., 2013), how they adapt (David et al., 2012) and what computations they are performing (Lindeberg, 2013). Spike events are interpreted as the letters utilized by neurons to communicate through synapses (Shadlen and Newsome, 1994), and the relative inter-spike interval is used in the literature as the key ingredient to investigate rules of synaptic plasticity (Markram et al., 2012). Different

models of neural coding have been proposed based on spikes (e.g. rate coding or spike-time coding)(Prescott and Sejnowski, 2008), and spiking activity of neurons is used to analyze properties of brains' dynamics, as spike-time correlation (De La Rocha et al., 2007) or Fano factors (Berry and Warland, 1997).

Yet, trying to understand brains computations from spikes encounters various conceptual and empirical problems, that are preventing the development of a consistent theory of neural coding. The concept of tuning curves, which proved to be very helpful when investigating representations (Piazza et al., 2004) and their transformations (Wilson et al., 2004) in different areas of the brain, is showing its inadequacy in capturing neural language. In this sense, the use of rich stimuli under close loop conditions is revealing the necessity to revise the standard concepts of representation (Fairhall, 2014). Moreover, the high trial-to-trial variability of neural spiking response to a constant stimulus (Reich et al., 1997) challenges our current understanding of neural computations. If, on one side, the idea that neurons are adopting an efficient coding scheme to represent information is widely accepted (Olshausen and Field, 1996; Fairhall et al., 2001), how this is implemented and how this information can be decoded efficiently remains an open question.

Spiking activity can be measured directly, using neural recordings, or indirectly, measuring quantities in the brain that are affected by spike events, as the blood-oxygen-level-dependent (BOLD) response (measured by fMRI)(Amano et al., 2009), or calcium concentration inside neurons (measured by calcium imaging)(Mittmann et al., 2011). Some of these quantities (e.g. calcium) also act as 2nd messengers (Augustine et al., 2003) mediating various internal molecular processes inside the neuron (e.g. vesicle release (Neher and Sakaba, 2008; Borst and Sakmann, 1996)). These quantities both influence and are influenced by the spiking activity of the neuron, and various computational roles have been proposed for their dynamics (Loewenstein and Sompolinsky, 2003).

In Chapter 3 (and specifically in Section 3.5) we introduce some findings involving the role

of calcium in neural dynamics. Specifically, we identify a quantity predicted by our normative model of neural adaptation with intracellular calcium concentration, and we show that calcium dynamics is compatible with various testable requirements of such a quantity. These results motivate and inspire a new (subtly different) interpretation of spikes and calcium in the working brain, which has fundamental implications linking molecular processes to neural coding.

According to this new interpretation, intracellular calcium concentration is the primary and most explicit form of information encoding inside the neuron. It controls synaptic communications between neurons and mediates short and long term plasticity. On the contrary, spikes role is to synchronize and propagate information (calcium level) to different regions within the neuron. Action potentials are very fast in spreading information to various calcium hotspots, and the binary nature of their message enables the neuron to consistently and reliably synchronize different areas avoiding error propagation. But action potentials are also metabolically expensive (Laughlin et al., 1998; Lennie, 2003; Neishabouri and Faisal, 2011), and high firing rate can hardly be sustained for long periods (Neishabouri and Faisal, 2013). For this reason, neurons use an efficient coding scheme to encode information with spikes, in order to maximize information encoded per spike event. The given coding strategy generates spikes only to encode newly available information (predictive coding (Srinivasan and Laughlin, 1982; Rao and Ballard, 1999)), which is not already encoded in current calcium concentration (i.e. it has not been communicated by previous spikes). By doing this, spikes will update the level of calcium in all the neuron, this improving neural encoding of the external world. The adaptive nature of this encoding makes it impossible for an external observer to decode the spiking activity of the neuron without taking into account the current calcium intracellular concentration, and without knowing the coding parameters, which depend on the dynamical properties of the encoded stimulus.

To support our assumption, in this Chapter we propose a normative, yet bottom-up model of Calcium computation, and we derive in a first-principle manner both molecular and neu-

ral adaptation dynamics. First (section 4.3) we derive how Calcium concentration should drive probability of vesicle release at the axonal terminations to maximize information transfer, thereby predicting relationships compatible with experimental observations (Schneggenburger and Neher, 2000). Second (section 4.4) using a Bayesian approach we derive a normative model of how synaptic information encoded in pre-synaptic vesicles releases should be optimally translated into post-synaptic membrane dynamics, thereby reproducing leaky integrate-and-fire neural dynamics (Shepard, 2007). Third (section 4.5) we derive how Calcium concentration can modulate the translation of subthreshold membrane dynamics into spikes. So as to implement predictive coding, we find that local Calcium level needs to act as divisive gain normalization of spike initiation thresholds (Louie et al., 2011). Finally (Section 4.6) we analytically predict the local dynamics of Calcium from spiking, so as to enable a neuron to optimally decode spiking activity. The resulting dynamics is compatible with the known kinetics of voltage-gated Calcium channels (Loewenstein and Sompolinsky, 2003). Our processing pipeline is derived from first Information Theoretic and Bayesian principles and predicts both molecular and firing rate adaptation data in mammals.

According to our proposed model, the probability of initiating of a new spike in the soma/dendrites depends both on the membrane potential (newly available information to be encoded) and on the current level of Calcium concentration (information already encoded by the neuron). When generated, an action potential propagates through all the neuron and it affects the flow of Calcium through the membrane, improving neural representation of the hidden variable. We analytically predict how parameters controlling the flow of Calcium depends on the dynamical properties of the stimulus encoded (e.g. timescale of changes and dynamic range), and how they can be tuned locally and on-line to the experienced stimulus. On the other side, the local level of Calcium affects spike initiation mechanisms, acting as divisive gain normalization by controlling Calcium-dependent currents (Smith et al., 2002). This allows spiking activity to implement predictive coding and to represent only new non-redundant information, which is not already encoded by previous spikes. The proposed mechanism acts

as to increase the entropy of spikes, explaining high spike-trains variability and Fano factors (Stein et al., 2005). Furthermore, the given behavior produces spike-rate short-term adaptation (as explored in Chapter 3) as a side effect, since a recent intense activity of the neuron will increase the concentration of Calcium and, as a consequence, it will reduce the probability of generating new spikes. Analytical model predictions about spike rate adaptation are in good agreement with results from cat V1 in-vivo recording (Benucci et al., 2013).

4.2 Representing information inside the neuron

We assume that the aim of a neuron is to represent a certain quantity v_t and its variations in time. v_t describes some property of the world, i.e. a cause underlying sensory observations. This quantity follows a stochastic dynamics, and the neuron uses all available information to infer and track its current value. According to the proposed model, action potentials serve to propagate information about v_t through the neuron. The given information is then represented locally in form of intracellular calcium ions concentration, and it is used to control local processes, such as synaptic communication, neural adaptation and synaptic plasticity.

4.2.1 Calcium encodes log-odds probability ratio of hidden state

We propose that the quantity v_t encoded by the neuron relates to the probability of a certain event (also referred as filter) currently taking place in the world. The given event is described by a hidden state x_t , which can assume two possible values ($x_t = 1$ or $x_t = 0$). The probability P_t of the hidden state assuming either value changes with time as we receive additional information about the current state of x_t . Specifically, we propose that the variable v_t represents the log-odds ratio (also called Logit function) L_t (Deneve, 2008; Cramer, 2003) of

the current probability, defined as

$$v_t = L_t(P) = \log\left(\frac{P}{1-P}\right) = \log\left(\frac{P(x_t = 1|s_{0 \rightarrow t})}{P(x_t = 0|s_{0 \rightarrow t})}\right) \quad (4.1)$$

where $s_{0 \rightarrow t}$ identifies the synaptic input received from the initial time 0 to the current time t . According to our model, the log-odds probability ratio is encoded inside the neuron in form of calcium concentration. Specifically, we assume that the log-odds ratio v_t is a linear function of the current intra-cellular calcium concentration $[c]_t$

$$v_t = L_t(P) = h_o + h_1[c]_t \quad (4.2)$$

The variable $[c]_t$ is dimensionless, and corresponds to the current local calcium concentration divided by a certain reference concentration $[Ca^{2+}]_{ref}$, which will vary in different regions of the neuron

$$[c]_t = \frac{[Ca^{2+}]_t}{[Ca^{2+}]_{ref}} \quad (4.3)$$

Nevertheless, in the following we will refer to $[c]_t$ as calcium concentration neglecting the reference concentration, for simplicity of discussion.

It can be noted that $[c]_t$ is necessarily positive. This requirement can be satisfied in equation (4.2) for an appropriate choice of the parameters h_o and h_1 . In fact, the log-odds ratio is inferior bounded by a certain value, which is the Logit function of the minimum probability P that will be inferred by the neuron if no evidence (pre-synaptic spike) about the state $x_t = 1$ is received for a long period of time. This value P_{\min} will never converge to 0, since the dynamics of the state x_t is stochastic, and the probability that x_t recently changed its value from 0 to 1 (and didn't generate any observable consequence yet) will never be negligible. This can be also noted by equation (3.50) in Chapter 3, and will be explored in equation (4.18) and in Section 4.6.

4.2.2 Why the log-odds probability ratio is a good choice

The log-odds probability ratio L (also called Logit function) is the inverse of the Logistic function

$$P(v) = L^{-1}(v) = \text{Logistic}(v) = \frac{e^v}{1 + e^v} \quad (4.4)$$

The Logistic function arises as a natural representation for the posterior probability in a binary classification problem (Jordan, 1995). To show this, we can consider a simple case of classification problem where the hidden state x can take 2 possible values $x \in x_0, x_1$. Available observations consist of a d -dimensional random vector $\omega = [\omega_1, \omega_2, \dots, \omega_d]^T$. $P(x)$ represents the marginal (prior) probability of the hidden state x , $P(\omega|x)$ is the conditional (likelihood) probability of ω given the state x and $P(x|\omega)$ is the inferred (posterior) probability of the state x given the observation ω . If we observe a specific vector ω and we want to infer the probability of the hidden state x , we can use Bayes theorem (2.2) to write

$$P(x_1|\omega) = \frac{P(\omega|x_1)P(x_1)}{P(\omega)} \quad (4.5)$$

$$= \frac{P(\omega|x_1)P(x_1)}{P(\omega|x_1)P(x_1) + P(\omega|x_0)P(x_0)} \quad (4.6)$$

$$= \frac{1}{1 + \exp \left\{ -\log \left(\frac{P(\omega|x_1)}{P(\omega|x_0)} \right) - \log \left(\frac{P(x_1)}{P(x_0)} \right) \right\}} \quad (4.7)$$

$$= \frac{1}{1 + e^{-\xi}} \quad (4.8)$$

Therefore, the posterior probability $P(x_1|\omega)$ can be expressed in form of Logistic function, given an appropriate choice of the function argument ξ . Moreover, it is interesting to note that the argument ξ , which is the log-odds ratio the posterior probability, is expressed in form of a linear sum of the log-odds ratio of the prior probability $P(x_1)$ and of a likelihood term, which encodes information from the observation ω . It should be noted that this likelihood term is not the log-odds ratio of the likelihood probability $P(\omega|x_1)$, since the likelihood distribution

is not necessarily normalized, and $P(\omega|x_0) \neq 1 - P(\omega|x_1)$. If additional information about the state of x becomes available, the new likelihood term can be computed, and added linearly to the previously calculate log-odds ratio in a recursive process. Therefore, if a series of independent observations are available about a certain hidden state x , the posterior distribution can be calculated as a linear sum of the prior log-odds ratio and of all the likelihood log-odds terms, corresponding to different observations. This allows to perform Bayesian inference of independent variables by linear summation.

In order to predict how information generated by correlated variables should be combined optimally to infer the probability of the hidden state x , we can go back to our previous example, and assume that the observation vector ω is sampled from a multivariate Gaussian distribution, with identical covariance matrix Σ and mean dependent on the hidden state, respectively μ_0 or μ_1 .

$$P(\omega|x_i) = \frac{1}{(2\pi)^{\frac{d}{2}} |\Sigma|^{\frac{1}{2}}} e^{-\frac{1}{2}(\omega - \mu_i)^T \Sigma^{-1} (\omega - \mu_i)} \quad (4.9)$$

If we expand the exponent of the Gaussian, and we substitute it into the formula for the log-odds ratio of the posterior ξ (equation (4.8)), we can write

$$\xi = \log \left(\frac{P(x_1)}{P(x_0)} \right) + a + b^T \omega \quad (4.10)$$

where

$$a = -\frac{1}{2}(\mu_0 + \mu_1)^T \Sigma^{-1} (\mu_1 - \mu_0) \quad (4.11)$$

and

$$b = \Sigma^{-1} (\mu_1 - \mu_0) \quad (4.12)$$

From this result it is possible to note that, even when observations are sampled from a vector

of correlated variables, in the Logit space we can combine linearly the log-odds likelihood terms generated by different observations ω_j

$$\xi = \log \left(\frac{P(x_1)}{P(x_0)} \right) + \sum_j a_j + b_j \omega_j \quad (4.13)$$

where parameters a_j and b_j will depend on the mean, variance and correlation structure of the variables. Furthermore, it can be noted that the log-odds likelihood term $a_j + b_j \omega_j$ generated by the observation ω_j is a linear function of the value ω_j itself (see figure 4.1). This is a very interesting result, as it allows implementing inference by linearly integrating the value of the received observations, given an appropriate choice of the parameters.

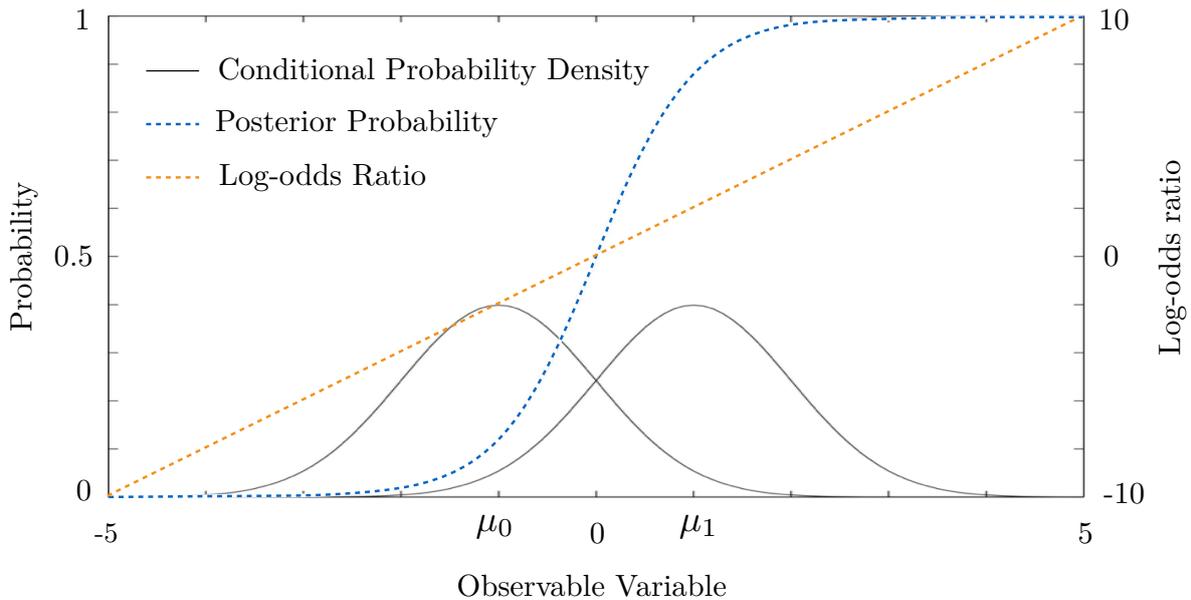


Figure 4.1: The figure represents the result of Bayesian inference in a binary classification problem, where the class-conditional distributions (black continuous lines) are mono-dimensional Gaussians with unit variance and means, respectively, $\mu_0 = -1$ and $\mu_1 = 1$. The posterior probability (blue dashed line) is given by the logistic function $1/(1 + \exp(-2x))$ and the log-odds ratio (orange dashed line) is given by a linear function of the input ($2x$). Figure adapted from (Jordan, 1995).

We propose that neurons use log-odds representations of probabilities to be able to optimally

combine independent sources of information by linearly integrating them in the dendritic tree. We know that dendrites are capable of non-linear summation of inputs. However, often, a surprising balance of boosting and shunting mechanisms is tuned in such a way to implement linear summation of most input combinations (Cash and Yuste, 1999). We showed that a linear summation is sufficient to generate optimal inference also in the case when the observed variables present a correlation structure, as long as the mean value μ_i depends on the hidden state x_i , but the covariance matrix Σ is independent.

If the covariance matrix Σ_i depends on the hidden state, or if the sources of information to be combined present a dependence structure which goes beyond first-order correlation, then in order to implement optimal inference in the Logit space we will have to introduce some non-linearities, that can explain non-linearities experimentally observed in dendritic computations (London and Häusser, 2005).

4.3 Synaptic vesicles release

According to the standard interpretation of synaptic communication, the role of chemical synapses in the brain is to enable connected neurons to interact by translating pre-synaptic action potentials into excitatory postsynaptic potentials (EPSP). This is implemented in four steps:

1. The pre-synaptic action potential opens voltage-sensitive calcium channels in the presynaptic membrane, allowing Ca^{+2} ions to enter the cell.
2. The increased intercellular concentration of calcium causes vesicles containing neurotransmitters to migrate towards the synaptic cleft, and to fuse with the neural membrane.

3. Neurotransmitter molecules are released into the synaptic cleft, and they diffuse to the post-synaptic neuron.
4. Neurotransmitter molecules bind with some receptors on the post-synaptic neuron, causing ions channels to open and ions to flow through the membrane, generating an excitatory postsynaptic potential.

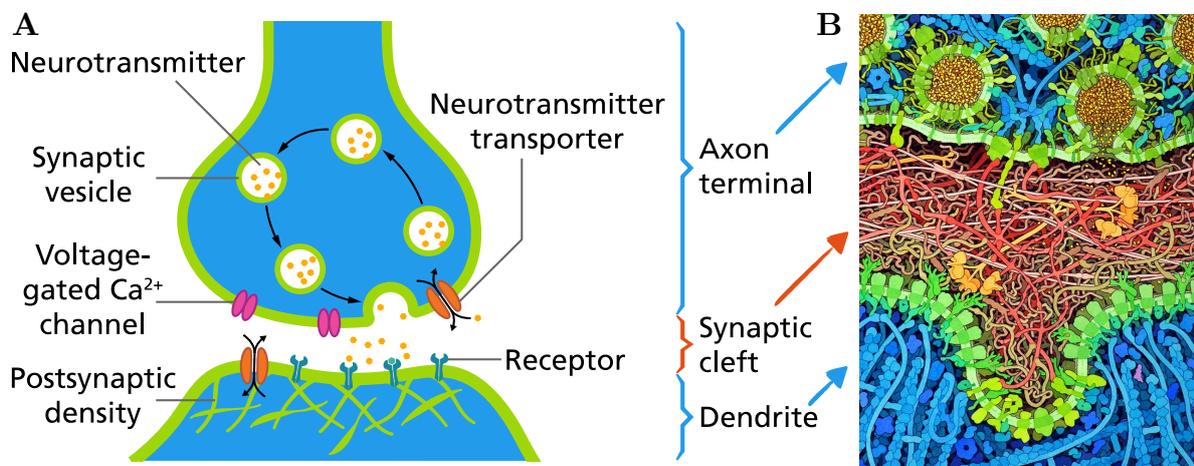


Figure 4.2: Functional (A) and biophysical (B) illustration of a chemical synapse. When an action potential approaches the synapse, the positive charge causes the opening of voltage-gated calcium channels (green). Calcium pours into the synaptic button and binds to several proteins, changing their shape. The activated proteins dynamically rearrange the blue cytoskeleton to transport green vesicles filled with yellow neurotransmitters to the synaptic cleft, which is filled with red adhesion proteins. Calcium-activated (SNARE) proteins bind to both the vesicles and the synaptic membrane, causing the vesicles to fuse with the membrane, turning them inside out and spilling neurotransmitters into the synaptic cleft. Neurotransmitters then bind to proteins on the receiving cell. There are several types of receptor proteins. Sodium (Na^+) channels (excitatory) respond to the neurotransmitter Glutamate. Chloride (Cl^-) channels (inhibitory) respond to the neurotransmitter GABA. Dopamine, Serotonin, and Opioids bind to G-Protein Coupled Receptors (GPCRs) which cause complicated phosphorylation cascades that change the metabolism of the cell.

This process is highly variable. The number of vesicles released per action potential is stochastic, and described by a binomial distribution (the release of each vesicle is an independent event, which can happen with probability p (Branco and Staras, 2009)). Also, the

average number of vesicles released per spike event in the central nervous system is often very small (usually lower than 1) (Harris and Sultan, 1995), meaning that some of the pre-synaptic spikes will fail to be transmitted. Moreover, the probability of vesicle release has been shown to depend not only on the number of vesicles available (Murthy et al., 1997), but also on the spiking history of the neuron (which influence the flow of calcium through the membrane)(Dobrunz and Stevens, 1997). This makes synaptic transmission of spike events unreliable, and it introduces variability that propagates in successive layers of neural processing.

As already mentioned, in Chapter 3.5 our findings inspired an alternative interpretation of the role of calcium and spikes in neural dynamics. This new picture of the working neuron can put a new light also on our understanding of the computational role of chemical synapses and their functioning. According to this new perspective, synapses are not converting pre-synaptic spikes into post-synaptic membrane fluctuations, but they are generating samples from a distribution, which is encoded in the neuron by pre-synaptic calcium concentration. The release of each neurotransmitter vesicle can be considered as an independent inhomogeneous Poisson process, whose rate is controlled by calcium. This new interpretation explains synaptic failures and variability, as well as the low average number of vesicles released per spike and the dependence of the release probability on the spiking history.

4.3.1 Sampling

According to the proposed interpretation of synaptic vesicle release, the log-odds ratio $L(P)$, encoded in the form of calcium concentration in the pre-synaptic neuron, controls the release of neurotransmitter vesicles into the synaptic cleft. We propose that this conversion takes the form of a sampling process: the probability of releasing a vesicle in a reference time interval τ_s given the current calcium concentration is proportional to the probability of the hidden state being positive $P(x_t = 1)$. On extended periods of time, the release of vesicles in the synapse

will be governed by a Poisson process. The average number of vesicles released during a certain time interval τ will be linearly proportional to the number of available vesicles, to the encoded probability P and to the ratio between the considered time interval τ and the reference time τ_s (compatibly with a Poisson distribution). In this sense, the release of vesicles can be interpreted as sampling: each time a vesicle is released, it represents in the brain the realization of a world where the encoded event took place, and vice versa. By doing this, the probability of the given event happening is encoded in form of a rate (frequency of vesicles being released in time). Each event (vesicle released) will be interpreted as an observation by successive neurons in the processing pipeline, which encode correlated features of the world (events). The given neurons will then use available observations to infer the probability of the encoded events, to represent it in form of calcium concentration and, iteratively, to generate new samples.

Our model enables the brain to produce internal simulations of the world, where events more likely (conditionally to the available sensory information) happens more often. Specifically, this would predict that the probability of representing a certain feature in the brain is proportional to the inferred probability of such a feature being present in the world given all the available sensory information, which is compatible with experimental results on humans (Wozny et al., 2010). This behavior is usually referred in the literature as “probability matching in perception”, and it’s predictions are in contrast with traditional “optimal Bayesian decision strategy” (Duda et al., 2012), which predicts that the most likely feature should be the only one to be represented (this maximizing the likelihood of a correct perception). Furthermore, our model predicts that the sampling process taking place in the brain also controls decision making and action selection. This explains experimental observations, which show probability matching behavior of people in choices under uncertainty (Koehler and James, 2009) (Murray et al., 2015). In addition, the given model predicts that spontaneous activity of the brain in the absence of sensory stimuli should be consistent with the known structure of the world: the activity (vesicle release) of different neurons should be correlated if these

neurons encode features which are correlated in the world. Moreover, the statistics of different patterns of activation in the spontaneous activity regime should match the statistics of these patterns in the world, which also correspond to the average persistence of these patterns in the brain when neural activity is driven by natural stimuli (this also matching experimental observations (Berkes et al., 2011)). Finally, the proposed model gives a principled explanation for the quantized nature of neurotransmitter release, which apparently reduces information capacity of the synaptic communication channel compared to using a continuous concentration spectrum.

The hypothesis of Bayesian sampling in the brain (Sampling Brain Hypothesis) has been already proposed in the literature (Fiser et al., 2010; Berkes et al., 2011). Yet, for the first time here we suggest that samples are represented by vesicles released in the synapse, and not by spikes generated inside the neuron. This change of perspective is driven by our new interpretation of the role of calcium and spikes in brain functioning, and it would explain high variability and stochasticity in synaptic functioning. Furthermore, since the number of synapses in the brain is much larger than the number of neurons, using vesicle release instead of spikes for sampling would increase statistical reliability of the implemented simulations. Different synapses of the same pre-synaptic neuron will implement independent sampling processes, this reducing sampling artifacts and improving simulation accuracy (according to the central limit theorem). The proposed sampling process plays a role that can also be interpreted in a dithering perspective. Neurons encode the probability of a hidden state, which is represented by a continuous variable between 0 and 1. However, synaptic communication (vesicle release) is binary by nature. In each moment in time, a vesicle is either released or not. The Bayesian maximum-likelihood approach would prescribe to always release a vesicle for probabilities higher than 0.5. In this sense, adding variability to the process does something similar to what dithering does in video/audio processing. Specifically, stochastic vesicle release reduces quantization error, it prevents large-scale patterns and it improves the overall quality of the internal representation. From the binary activity of all the synapses of a given

neuron it would be possible to accurately infer the encoded probability value. Therefore, the information encoded by the neuron is reliably propagated to the rest of the system, which would not happen in the case of a probability matching behavior of synapses.

The proposed model of synaptic communication requires some conditions to be verified by the biophysical synapses. First, the release of a vesicle should generate an effect on the post-synaptic neuron that is equal for different vesicles of the same synapse, and which is independent of the history of the synapse (therefore constant in time). This is a requirement, since we want each sample (vesicle release) to bring a consistent message to the post-synaptic neuron. This condition matches experimental results (Dobrunz and Stevens, 1997), which show that the "potency" of vesicles in central nervous system synapses is independent of the vesicle release history of the synapse, and that it is independent of the vesicle release probability. Additionally, in the central nervous system (specifically in hippocampal area CA1) it has been observed that synapses contain more than ten times the concentration of neurotransmitter necessary to saturate the post-synaptic receptors (Harris and Sultan, 1995). This result is compatible with the idea that the brain considers each vesicle release as a sample, which brings a quantized piece of information, and whose message is consistent in time and independent of the particular state of the brain.

Moreover, our model predicts that, in any given moment in time, the probability of releasing a vesicle in the reference time interval τ_s is given by the logistic function of v_t (equation (4.4)), which is a linear function of calcium concentration $[ca^{2+}]$ (equation (4.2)). Therefore, under the assumption that P is small, we can predict a specific mathematical relationship between release probability and pre-synaptic calcium concentration

$$P([c]) = \frac{e^{h_o+h_1[c]}}{1 + e^{h_o+h_1[c]}} \quad (4.14)$$

If multiplied for the total number of available synapses and divided by an appropriate time constant (τ_s), the release probability enables us to calculate the average vesicle release rate

in the synapse.

$$\nu_s = \frac{N}{\tau_s} P([c]) \quad (4.15)$$

We tested this model prediction with experimental results about vesicle release rate in synapses (Schneppenburger and Rosenmund, 2015), and for an appropriate choice of the model parameters our predictions qualitatively matched experimental observations (figure 4.3). Assuming a linear dependence of excitatory postsynaptic potential (EPSP) on neurotransmitter concentration, relation (4.14) also implies a specific dependence of the average EPSP amplitude from pre-synaptic calcium concentration. This dependence is compatible with experimental results from the end-plate potential of the neuromuscular junction of vertebrates (Dodge and Rahamimoff, 1967). Finally, this result can be extended also to extra-cellular calcium concentration. In fact, the flow of calcium ions through the membrane is a stochastic process, and the rate of ions entering the membrane after a spike is proportional to the concentration of ions available. Therefore, a higher concentration of calcium outside the cell will linearly reflect into higher concentration inside the cell.

4.3.2 Short term synaptic plasticity

It has been experimentally observed that the release rate of vesicles (and consequent EPSP) following a pre-synaptic spike event can vary over time, and depends on the recent spiking history of the neuron (Zucker and Regehr, 2002). This major property of synaptic dynamics is usually referred to as “short term synaptic plasticity”, and it is divided into “short term potentiation” and “short term depression”.

Short term potentiation happens when a recent activation of the synapse enhances synaptic transmission of pre-synaptic spikes. This mechanism has been known for more than 70 years (Eccles et al., 1941), and the dominant component on very short timescales is referred to as “facilitation”, and it is usually investigated by considering pairs of stimuli. If the inter-

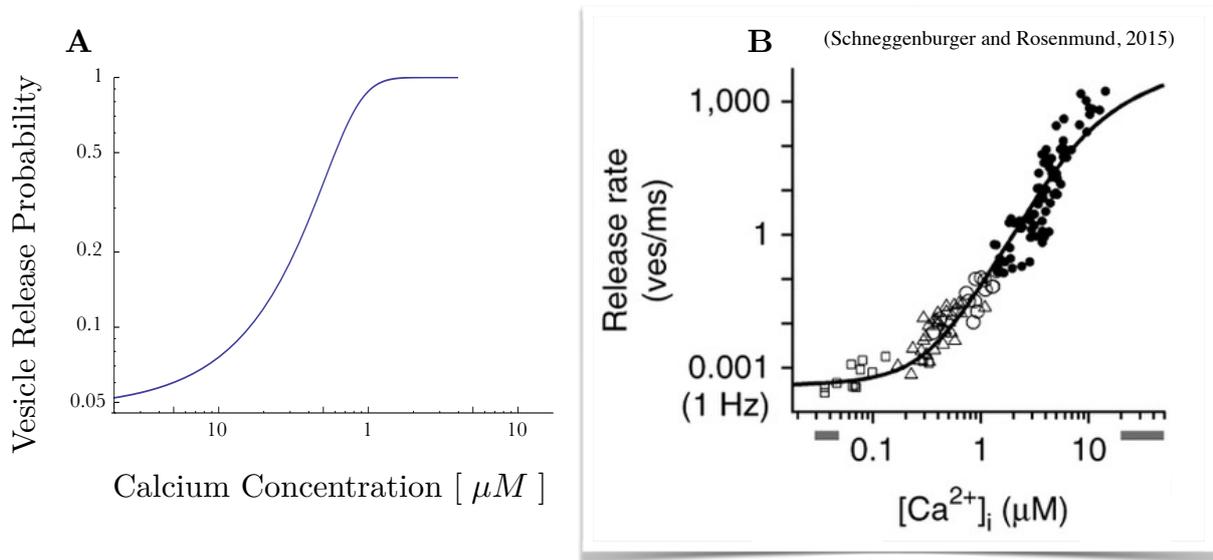


Figure 4.3: Predicted and experimental dependence of release rate of neurotransmitter vesicles on intra-cellular calcium concentration at the axonal termination. (A) Our model predicts that vesicle release rate is proportional to the release probability $P([c])$, which is given by the logistic function of calcium concentration. The graph represents the curve predicted by our model (equations (4.14)-(4.3)), which describe the dependence of release probability on calcium concentration. Parameters $h_o = -3$, $h_1 = 5$ and $[Ca_{ref}^{2+}] = 1\mu M$. (B) Experimental observations about vesicle release rate as a function intra-cellular calcium concentration. Figure adapted from (Schneggenburger and Rosenmund, 2015).

spike interval between two successive spikes/pulses is small (hundreds of ms), the excitatory postsynaptic potential (EPSP) evoked by the second spike/pulse will be larger than the one from the first spike. Facilitation is often measured as a ratio between the EPSP evoked by the two spikes, and it is usually observed to decrease almost-exponentially as a function of the inter-spike interval.

It has been shown that in the neuromuscular junctions, an action potential can cause facilitation even if there is no transmitter release (Dudel and Kuffler, 1961). For this reason, we can assume that facilitation is caused by some process controlled by pre-synaptic action potentials, yet independent from neurotransmitter secretion. This leads naturally to the hypothesis that facilitation is controlled by the influx of Ca^{2+} ions in the pre-synaptic ax-

onal termination (Catterall et al., 2013). Support for this hypothesis comes from (Katz and Miledi, 1968).

Our model of synaptic dynamics predicts short term potentiation as a consequence of calcium accumulating in the pre-synaptic neuron, and of the super-linear dependence of the release probability/rate on calcium concentration for low values of P . Calcium dynamics conditional to pre-synaptic spikes can be derived by equation (4.140) (as will be derived in details in Section 4.6)

$$\begin{aligned} \frac{d}{dt}[c]_t = & 2h_1^{-1} (1 + \cosh(h_o + h_1[c]_t)) \left[\tau_r k_{\sigma_i^2, \mu_o, r_i} (\delta(t - t_{\text{spike}}) - \mu_o) + r_i \mu_i \tau_r \right] \\ & - h_1^{-1} r_i (1 + e^{(h_o + h_1[c]_t)}) \end{aligned} \quad (4.16)$$

where \cosh is the hyperbolic cosine and τ_r is the refractory period of the neuron, as we will see in Section 4.5.3. The given equation can be rewritten as a flow equation for the log-odds ratio v_t using equation (4.2)

$$\frac{d}{dt}v_t = 2 (1 + \cosh(v_t)) \left[\tau_r k_{\sigma_i^2, \mu_o, r_i} (\delta(t - t_{\text{spike}}) - \mu_o) + r_i \mu_i \tau_r \right] - r_i (1 + e^{v_t}) \quad (4.17)$$

We can consider as starting condition v_s the equilibrium value v_{eq} , which is the average value of v_t given the activity of the neuron conditional to the current configuration of the external world. Alternatively, we can consider the situation when the neuron doesn't generate a spike for a long period of time. In this case, v_s will converge to its minimum value v_m , which is given by

$$v_m = \log \left(\frac{\left(\mu_i - \frac{\mu_o k_{\sigma_i^2, \mu_o, r_i}}{r_i} \right)}{\tau_r^{-1} - \left(\mu_i - \frac{\mu_o k_{\sigma_i^2, \mu_o, r_i}}{r_i} \right)} \right) \quad (4.18)$$

where we assume an appropriate choice of the model parameters, which guarantees the quantity (v_m) to be well defined. After the first pre-synaptic spike, the value of v_t will increase

from the starting value v_s of a certain quantity Δ_v^1

$$\Delta_v^1 = 2 (1 + \cosh(v_s)) \tau_r k_{\sigma_i^2, \mu_o, r_i} \quad (4.19)$$

Then, from the new value $v_s + \Delta_v^1$ it will start decreasing back to v_m , according to the differential equation

$$\frac{d}{dt} v_t = 2 (1 + \cosh(v_t)) \left(r_i \mu_i \tau_r - \mu_o \tau_r k_{\sigma_i^2, \mu_o, r_i} \right) - r_i (1 + e^{v_t}) \quad (4.20)$$

whose solution, for the initial condition $v_{t=0} = (v_s + \Delta_v^1)$, is given by

$$v_t = \log \left(\frac{\left(e^{v_s + \Delta_v^1} + 1 \right) \left(r_i \mu_i \tau_r - \mu_o \tau_r k_{\sigma_i^2, \mu_o, r_i} \right) (e^{r_i t} - 1) + e^{v_s + \Delta_v^1} r_i}{r_i (e^{v_s + \Delta_v^1 + r_i t} - e^{v_s + \Delta_v^1} + e^{r_i t}) - (e^{v_s + \Delta_v^1} + 1) \left(r_i \mu_i \tau_r - \mu_o \tau_r k_{\sigma_i^2, \mu_o, r_i} \right) (e^{r_i t} - 1)} \right) \quad (4.21)$$

When the second spike reaches the synapse (after a time τ the log-odds ratio will increase of a quantity Δ_v^2 from the current value v_τ , where

$$\Delta_v^2 = 2 (1 + \cosh(v_\tau)) \tau_r k_{\sigma_i^2, \mu_o, r_i} \quad (4.22)$$

For a given log odds ratio v_t , the probability of vesicle release will be given by the inverse of the log odds probability ratio, which corresponds to the Logistic function (equation (4.4)). In addition, as a first order approximation, we can assume that the probability of a vesicle release depends on the maximum concentration of calcium $[c]$ (and consequent log-odds ratio v), which is reached immediately after the spike. Therefore, the probabilities for the first and second spike will be given by

$$P_1 = \text{Logistic}(v_s + \Delta_v^1) = \frac{e^{v_s + \Delta_v^1}}{1 + e^{v_s + \Delta_v^1}} \quad (4.23)$$

$$P_2 = \text{Logistic}(v_\tau + \Delta_v^2) = \frac{e^{v_\tau + \Delta_v^2}}{1 + e^{v_\tau + \Delta_v^2}} \quad (4.24)$$

where τ is the inter spike interval between first and second spike. The average EPSP will be proportional to probability of vesicle release. For this reason, we can calculate the ratio between the EPSPs evoked by the second and by the first spike, which is

$$\frac{\text{2nd EPSP}}{\text{1st EPSP}}(\tau) = \frac{1 + e^{-(v_s + \Delta_v^1)}}{1 + e^{-(v_\tau + \Delta_v^2)}} \quad (4.25)$$

where

$$v_s = v_m = \log \left(\frac{\left(\mu_i - \frac{\mu_o k_{\sigma_i^2, \mu_o, r_i}}{r_i} \right)}{\tau_r^{-1} - \left(\mu_i - \frac{\mu_o k_{\sigma_i^2, \mu_o, r_i}}{r_i} \right)} \right) \quad (4.26)$$

$$\Delta_v^1 = 2 (1 + \cosh(v_s)) \tau_r k_{\sigma_i^2, \mu_o, r_i} \quad (4.27)$$

$$\Delta_v^2 = 2 (1 + \cosh(v_\tau)) \tau_r k_{\sigma_i^2, \mu_o, r_i} \quad (4.28)$$

and

$$v_\tau = \log \left(\frac{\left(e^{v_s + \Delta_v^1} + 1 \right) \left(r_i \mu_i \tau_r - \mu_o \tau_r k_{\sigma_i^2, \mu_o, r_i} \right) (e^{r_i \tau} - 1) + e^{v_s + \Delta_v^1} r_i}{r_i (e^{v_s + \Delta_v^1 + r_i \tau} - e^{v_s + \Delta_v^1} + e^{r_i \tau}) - (e^{v_s + \Delta_v^1} + 1) \left(r_i \mu_i \tau_r - \mu_o \tau_r k_{\sigma_i^2, \mu_o, r_i} \right) (e^{r_i \tau} - 1)} \right) \quad (4.29)$$

The given result can be generalized from a pair of pre-synaptic spikes (as above) to a pair of pulses (paired pulse ratio) or a pair of stimuli (each presented for a short period of time), given an appropriate choice of the quantity Δ_v , which depends on the intensity and duration of the pulses/stimuli used in the experiment. We compared our model predictions with experimental results about short-term synaptic potentiation in mouse hippocampal CA1 neurons (Moresco, 2002) and mice hippocampal Tc1 (Witton et al., 2015). Given an appropriate (biophysically reasonable) choice of the model parameters, our predictions are compatible with experimental observations about the time-course of synaptic facilitation (and specifically of the paired-pulse ratio for different inter-stimulus intervals)(figure 4.4). Moreover,

since we assume that facilitation is controlled by calcium concentration, our model requires the timescale of facilitation decay to match the timescale of calcium recovery inside the neuron. This prediction is supported by experimental evidence (Delaney and Tank, 1994), as the two given timescales seem to be linearly correlated (see figure 4.5).

If a brief period of synaptic activation can generate short term potentiation, sustained synaptic activity in time can produce an opposite behavior, called short term depression (Zucker and Regehr, 2002). This phenomenon provides the synapse to be less responsive to incoming spikes, and to reduce both vesicle release rate and excitatory postsynaptic potentials. Evidence suggests that this behavior can be caused by depletion of available synapses to be released (Schneggenburger et al., 2002). In fact, in the synapse there is a limited number of available vesicles filled with neurotransmitters (tens-hundreds per synaptic bouton, according to the neuron type (Ikeda and Bekkers, 2009; Schikorski and Stevens, 1997)), and the replenishment of vesicles in the axonal terminal can be relatively slow (e.g. it's 0.82 vesicles per synapse per second for hippocampal CA1 neurons in rats (Awatramani et al., 2007)). For this reason, when high levels of activation are sustained for long periods of time, the release of neurotransmitters in the synaptic cleft is limited by vesicle replenishment rate (Harata et al., 2001; Abbott et al., 1997). To limit the consequences of this situation, which in the prospective of our model can be interpreted as unwanted side-effects, we assume that the brain took some precautions: (A) Synapses with higher release probability (rate) per single vesicle (and that are therefore more likely to get depleted) should be refurnished with a higher number of available vesicles (Dobrunz and Stevens, 1997). By doing this, the reserve of vesicles will be proportional to the expected maximum need, as a result of a trade-off between reliability and biophysical/metabolic costs. Experimental evidence shows that there is considerable variability in the number of docked vesicles at individual synapses, both across different type of synapses (Rao-Mirotnik et al., 1995; Schikorski and Stevens, 1999) and within the same synapse type (Harris and Sultan, 1995; Zucker and Regehr, 2002). (B) In each synaptic bouton there should be various groups of vesicles, with a different predis-

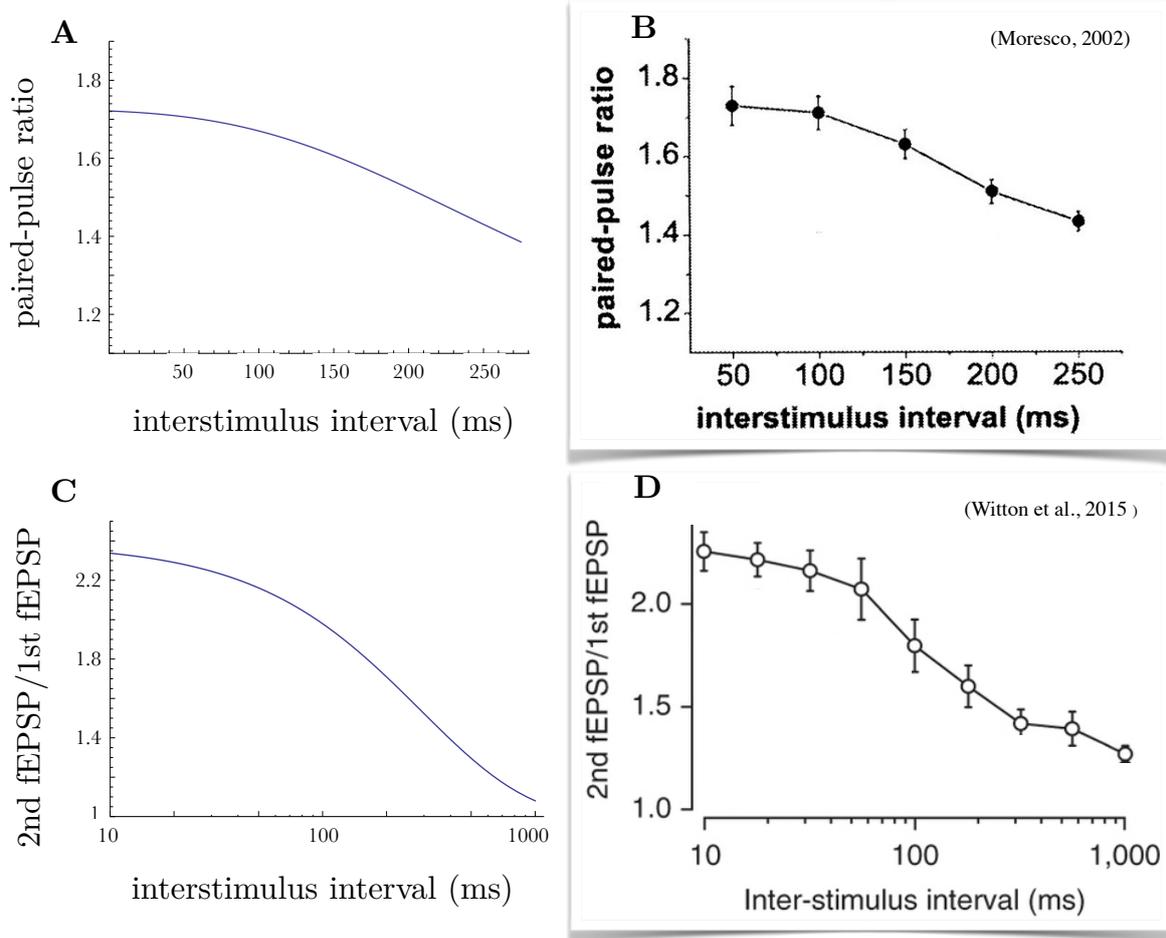


Figure 4.4: Model predictions about short-term potentiation in synaptic transmission of action potentials (A-C) are compared with experimental observations (B-D). All the graphs represent the ratio between the EPSP evoked by the second and by the first pre-synaptic spike, as a function of the inter-spike interval. (A) Curve predicted by our model of synaptic communication (equation (4.25)). (B) Experimental paired-pulse ratio from mouse hippocampal area CA1 (in-vitro). Schaffer collaterals axons were stimulated using bipolar tungsten electrodes (FHC) with enough current (50-s pulses) to reliably elicit synaptic responses. Figure adapted from (Moresco, 2002). (C) Curve predicted by our model (equation (4.25)) plotted in a semi-log plot. (D) Ratio between first and second field excitatory postsynaptic potential (fEPSP) amplitudes in response to a pair of 175 A stimulations in mice hippocampal Tc1. Figure adapted from (Witton et al., 2015).

position to be released. In this way, reactive vesicles will make the synapse responsive in communicating a weak stimulation. However, if a high level of activation is sustained for a

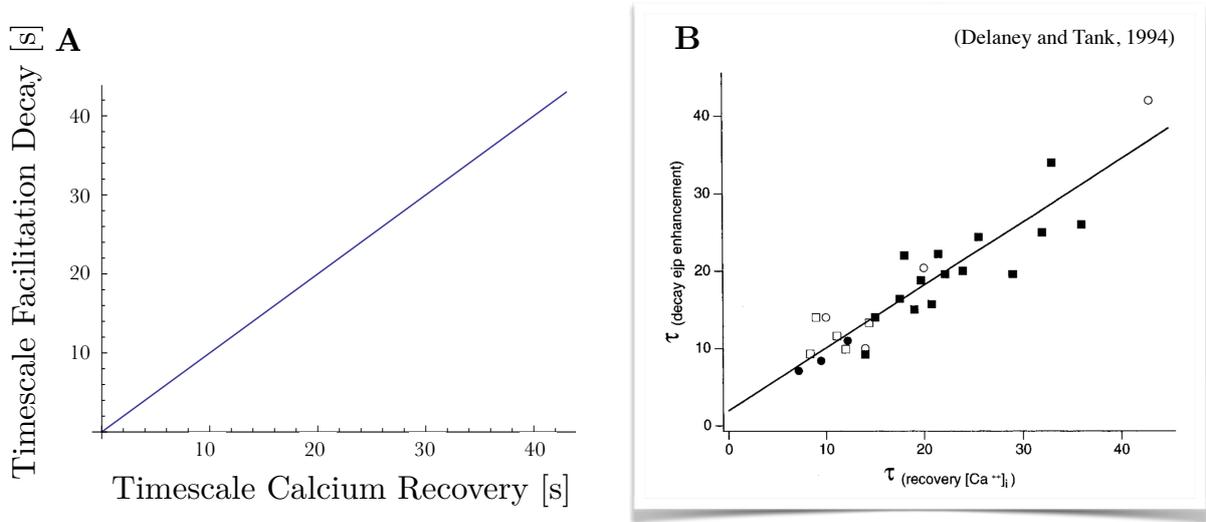


Figure 4.5: Experimental results show correlation between the timescale of calcium dynamics and synaptic potentiation decay, validating our predictions. (A) Our model predicts that facilitation is controlled by calcium concentration and, as a consequence, the timescale of facilitation decay should match calcium recovery. (B) Experimental observations in the crayfish claw opener neuromuscular junction, showing a linear correlation between the timescales of calcium and synaptic potentiation. Figure adapted from (Delaney and Tank, 1994).

long period of time, the residual resilient vesicles will prevent a complete depletion of the synapse, which will still be able to partially communicate outgoing messages. Vesicles that had recently entered the releasable pool are known to have lower release probability than those that had been in the pool for more than 30 s (Wu and Borst, 1999). Furthermore, resilient vesicles are known to be able to transform into responsive vesicles through a process which is controlled by calcium concentration (Wang and Kaczmarek, 1998). (C) To avoid depletion of synapses, calcium concentration (and corresponding log-odds ratio) in a working neuron should usually correspond to relatively small probabilities. This implies that, most of the time, the log-odds probability ratio v should assume negative numbers, corresponding to release probabilities lower than 0.5 (left side of the logistic function). This would explain an approximation often used in the literature (Stockbridge and Moore, 1984; Bertram et al., 1996; Dodge and Rahamimoff, 1967; Katz and Miledi, 1970), according to which the neu-

rotransmitter release rate N_r is assumed to be proportional to the fourth power of calcium $[ca^{2+}]^4$.

$$N_r = n [ca^{2+}]^4 \quad (4.30)$$

In fact, the dependence of release probability on calcium predicted by our model (equation (4.14)) is well approximated by the fourth power of calcium $[c]^4$ for values of calcium concentration that correspond to small (negative) values of the log-odds ratio v (see figure 4.6).

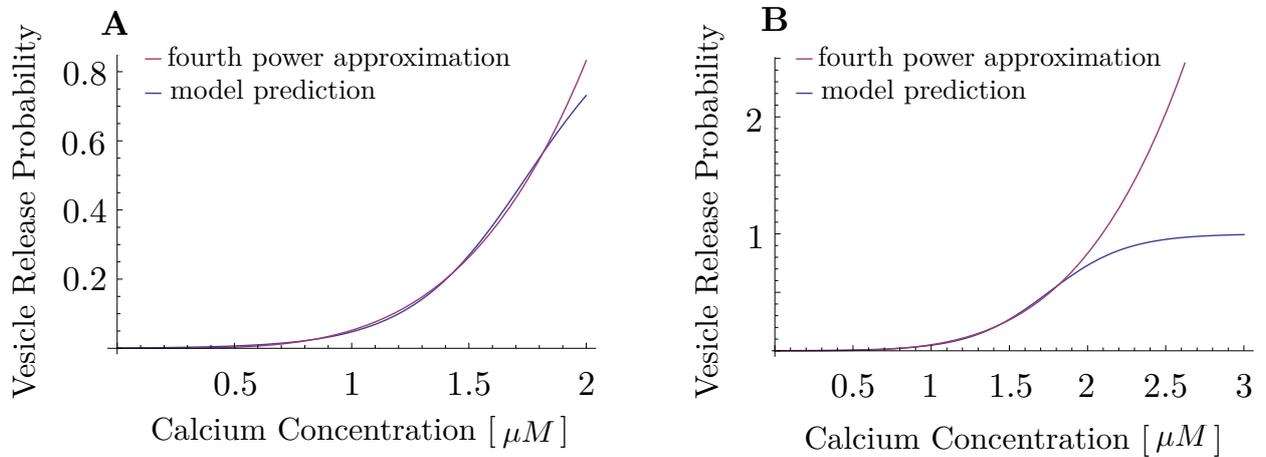


Figure 4.6: The probability of release predicted by our model (equation (4.14), blue line) is well approximated by the fourth power of calcium concentration (equation (4.30), red line) for small values of calcium (A). (B) For large values of calcium, neurotransmitter release diverges in the case of the power approximation, while it converges to a finite value according to our model. Parameters used in the plots are $h_0 = -7$, $h_1 = 4$ and $n = 0.05$.

Here, we introduced a model of how synapses encode and transmit information. Specifically, we propose that the neuron encodes the log-odds ratio of the probability of a certain hidden state in the world, and that this information is encoded in the synapse in the form of calcium concentration. We also propose that the release of neurotransmitter vesicles into the synaptic cleft takes the form of a sampling process, whose probability is controlled by the current log-odds ratio. We provide a theoretical justification for our proposal, and we show that our model can explain biophysical observations about synaptic dynamics (e.g. vesicle release

rate, short term potentiation and timescale of facilitation,) and that it is compatible with effective models used in the literature to describe synaptic behavior. If verified, such a model would provide a new interpretation of the computational role of calcium in synaptic functioning, and it would deliver a principled explanation of synaptic dynamics, allowing to better understand neural mechanisms such as short term potentiation. Additionally, a better understanding of synaptic computations would allow developing more advanced models of neural population dynamics, computations and learning.

In the following, we will expand our model of calcium-based neural computations to describe the dynamics of the post-synaptic neuron. We will see how the message transmitted by the synapse can be interpreted by the post-synaptic dendrite, and we will show how this information can be encoded locally (exploiting calcium dynamics) and it can be spread to different regions of the neuron in an efficient way (exploiting spikes and a predictive coding scheme).

4.4 Dendritic integration

The information encoded in synaptic vesicles release is used by the post-synaptic neuron to infer the configuration of the external world and, specifically, the probability of a certain hidden state x_t . To infer such a probability, the neuron estimates on-line the instantaneous release rate of the synapse, and it recurrently uses this quantity as an observed variable in a binary regression problem, efficiently combining the information received from different synapses. This is implemented in the dendrite, which is the region of the neuron responsible for integrating synaptic inputs (see figure 2.2).

4.4.1 Decoding synaptic vesicles release

According to our model, vesicle release is described by an inhomogeneous Poisson process, whose instantaneous rate ν_s is controlled by the probability encoded in pre-synaptic calcium concentration. Here, we assume that the aim of the post-synaptic dendrite is to infer the posterior distribution of the instantaneous hidden rate ν_s which generated the observed history of release events $\{t_r\}$, and to calculate, represent and communicate the mean value of this distribution. This task is solved on-line: each time a new vesicle release is observed, the information encoded in the inter-release interval is used to update and improve the current posterior distribution. We can recognize in this problem various similarities to the task addressed in 3.2.2, and we will use a similar approach to derive a normative model of synaptic ions flow that allows implementing near-optimal inference. Since we are considering a Poisson process without adaptation, we will achieve slightly different results than what was derived in Chapter 3 (equation (3.49)). Therefore, we will briefly repeat part of the derivation here for clarity; it is possible to make reference to 3.2.2 for more details.

We assume that the release rate ν_s follows a stochastic dynamics. We modeled this dynamics using either a stochastic jump process or a random walk, achieving consistent results. In the following we will consider the case of a stochastic jump process, since it represents the more complicated problem to be solved. Moreover, the latter process better describes the stochastic dynamics of calcium concentration in the pre-synaptic neuron, which presents a sudden jump each time a spike reaches the axonal termination. The stochastic jump process can be modeled mathematically using an underlying Poisson process with rate r_i . Each time an event takes place in the given process, ν_s is set to a new random value, sampled from a baseline distribution $B(\nu_s)$. For mathematical tractability, we will assume $B(\nu_s)$ to be a Gamma distribution, with parameters α_i and β_i .

$$B(\nu_s) = g(\nu_s; \alpha_i, \beta_i) ; \quad g(\nu_s; \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \nu_s^{\alpha-1} e^{-\beta \nu_s} \quad (4.31)$$

In the following, we assume that the post-synaptic dendrite exploits both the information received by vesicle release and the knowledge of the hidden variable ν_s stochastic dynamics to implement optimal inference, and to update the posterior distribution online. We can refer to $D_t(\nu_s)$ as the current posterior belief about the current release rate ν_s . Each time a new observation (inter-release interval) is available, this information can be used to derive a corresponding likelihood distribution, and update $D_t(\nu_s)$ accordingly. The distribution of inter-event intervals in a Poisson process is exponentially distributed and, as a consequence, it can be shown that the likelihood $L_\tau(\nu_s)$ corresponding to a certain interval τ observed between two consecutive events is Gamma distributed

$$L_\tau(\nu_s) = \tau^2 \nu_s e^{-\nu_s \tau} = g(\nu_s; \alpha_L = 2, \beta_L = \tau) \quad (4.32)$$

Furthermore, since the gamma distribution is the conjugate prior of the exponential distribution (Gelman et al., 2003), we can consistently express the posterior distribution both before $D_t(\nu_s)$ and after $D_{t+\tau}(\nu_s)$ the observation τ in form of Gamma distribution, and write

$$D_{t+\tau}(\nu_s) = K L_\tau(\nu_s) D_t(\nu_s) = g(\nu_s; \alpha_L + \alpha_t - 1, \beta_L + \beta_t) = g(\nu_s; \alpha_{t+\tau}, \beta_{t+\tau}) \quad (4.33)$$

Since we are interested in the mean value of the posterior distribution $D_t(\nu_s)$, we can use the relation for mean and variance of a Gamma distribution

$$\mu = \frac{\alpha}{\beta} \quad ; \quad \sigma^2 = \frac{\alpha}{\beta^2} \quad (4.34)$$

and rewrite equation 4.33 in form of update equation for mean and variance of the posterior

distribution

$$\mu_{t+\tau} = \frac{\alpha_t + \alpha_L}{\beta_t + \beta_L} = \frac{\frac{\mu_t^2}{\sigma_t^2} + 1}{\frac{\mu_t}{\sigma_t^2} + \tau} \quad ; \quad \sigma_{t+\tau}^2 = \frac{\alpha_t + \alpha_L}{(\beta_t + \beta_L)^2} = \frac{\frac{\mu_t^2}{\sigma_t^2} + 1}{\left(\frac{\mu_t}{\sigma_t^2} + \tau\right)^2} \quad (4.35)$$

We will now consider the assumption that the average interval $\langle \tau \rangle$ between two consecutive vesicle releases is smaller than the inverse of the timescale r_i on which the hidden rate ν_s changes.

$$\langle \tau \rangle \ll \frac{1}{r_i} \quad (4.36)$$

This can be seen as a necessary condition, since the synapse would not be able to efficiently transmit information about a variable ν_s whose value changes faster than the average rate of vesicle release. Under the given assumption, we can Taylor expand equation (4.35)

$$\mu_{t+\tau} = \mu_\tau + \frac{\sigma_t^2}{\mu_t} (1 - \mu_t \tau) - \frac{\sigma_t^4 \tau}{\mu_t^2} (1 - \mu_t \tau) \left(1 + O\left(\frac{\sigma_t^2 \tau}{\mu_t}\right) \right) \quad (4.37)$$

$$\sigma_{t+\tau}^2 = \sigma_t^2 + \frac{\sigma_t^4}{\mu_t^2} (1 - 2\mu_t \tau) \left(1 + O\left(\frac{\sigma_t^2 \tau}{\mu_t}\right) \right) \quad (4.38)$$

In addition, as a certain time interval τ elapses, there is a certain probability $p \approx r_i \tau$ that the variable ν_s has undertaken a jump, which set it to a new value. We assume that the post-synaptic dendrite is aware of the stochastic dynamics controlling the hidden variable ν_s , and that it exploits this knowledge to improve its representation. Therefore, the posterior distribution $D_{t+\tau}(\nu_s)$ after a time τ will be given by a weighted mixture of the previous posterior distribution $D_t(\nu_s)$ and of the base distribution $B(\nu_s)$

$$D_{t+\tau}(\nu_s) = (1 - r_i \tau) D_t(\nu_s) + r_i \tau P(\nu_s) = (1 - r_i \tau) g(\nu_s; \alpha_t, \beta_t) + r_i \tau g(\nu_s; \alpha_i, \beta_i) \quad (4.39)$$

Under the assumption that $\langle \tau \rangle \ll r_i^{-1}$, it is possible to show (see 3.2.3) that the posterior distribution $D_{t+\tau}(\nu_s)$ can still be accurately approximated using a Gamma distribution, and exploiting the relation for mean and variance of weighted mixture of distributions

$$E(X) = E[E(X|\theta)] \quad ; \quad \text{Var}(X) = E[\text{Var}(X|\theta)] + \text{Var}(E[X|\theta]) \quad (4.40)$$

we can write a flow equation for mean and variance of the posterior distribution $D_t(\nu_s)$, which accounts for the information loss derived by the stochastic dynamics of the hidden variable ν_s

$$\mu_{t+\tau} = (1 - r_i\tau)\mu_t + r_i\tau\mu_i = \mu_t + r_i\tau(\mu_i - \mu_t) \quad (4.41)$$

$$\sigma_{t+\tau}^2 = (1 - r_i\tau)(\mu_t^2 + \sigma_t^2) + r_i\tau(\mu_i^2 + \sigma_i^2) - ((1 - r_i\tau)\mu_t + r_i\tau\mu_i)^2 \quad (4.42)$$

Where μ_i and σ_i^2 are, respectively, mean and variance of the base distribution $B(\nu_s)$. Exploiting approximation $\langle \tau \rangle \ll r_i^{-1}$ one more time, we can expand equation (4.42) and write

$$\sigma_{t+\tau}^2 = \sigma_t^2 + (\sigma_i^2 - \sigma_t^2 + (\mu_t - \mu_i)^2)r_i\tau + O(r_i^2\tau^2) \quad (4.43)$$

Finally, by combining the update equations (4.37)-(4.38) that account for the information gain derived by the observed inter-release interval τ , and equations (4.41)-(4.43) that account for the information loss derived by the stochastic dynamics of the hidden variable ν_s , we can write

$$\mu_{t+\tau} = \mu_t + \frac{\sigma_t^2}{\mu_t} (1 - \mu_t\tau) + r_i\tau(\mu_i - \mu_t) \quad (4.44)$$

$$\sigma_{t+\tau}^2 = \sigma_t^2 + \frac{\sigma_t^4}{\mu_t^2} (1 - 2\mu_t\tau) + r_i\tau(\sigma_i^2 - \sigma_t^2 + (\mu_t - \mu_i)^2) \quad (4.45)$$

Each time a new vesicle is released and received by the post-synaptic neuron, the observed interval τ between the last and the previous vesicle release is used to update the parameters

of the Gamma distribution describing the inferred posterior $D_t(\nu_s)$. Equations (4.44)-(4.45) describe how mean and variance of the inferred distribution are updated on-line accordingly.

To reduce our system of coupled equations to an independent finite difference equation, we can rely on condition $\langle \tau \rangle \ll r_i^{-1}$, which implies that, in average, multiple release events are observed before there is a change in the hidden value ν_s . Therefore, we can assume that the variance has time to converge to its equilibrium value, and we can approximate σ_t^2 in equation (4.44) with the steady state solution σ_e^2 of equation (4.45), which satisfies

$$\langle \frac{\sigma_t^4}{\mu_t^2} (1 - 2\mu_t\tau) + r_i\tau(\sigma_i^2 - \sigma_t^2 + (\mu_t - \mu_i)^2) \rangle = 0 \quad (4.46)$$

It is possible to note that we are following a similar procedure to what we did in 3.2.5. However, we are here considering the case of a Poisson process without adaptation, which is leading us to different results. We can solve (4.46) by applying Ito's Lemma, and exploiting the fact that τ appears only as a linear term in the equation (4.46). Approximating $\langle \tau \rangle = \mu_i^{-1}$ we can write

$$\frac{\sigma_t^4}{\mu_t^2} \left(1 - 2\mu_t \frac{1}{\mu_t} \right) + r_i \frac{1}{\mu_t} (\sigma_i^2 - \sigma_t^2 + (\mu_t - \mu_i)^2) = 0 \quad (4.47)$$

The given equation has only one positive solution, which is

$$\sigma_e^2 = \frac{\mu_t r_i + \sqrt{\frac{r_i^2}{\mu_t^2} + 4 \frac{r_i}{\mu_t^3} \tau ((\mu_t - \mu_i)^2 + \sigma_i^2)}}{-\frac{2}{\mu_t}} \quad (4.48)$$

We can now invoke assumption $\langle \tau \rangle \ll r_i^{-1}$ one more time, which translates here into $r_i \ll \mu_t$, and expanding to the first order in $\sqrt{\frac{r_i}{\mu_t}}$ we can write

$$\sigma_e^2 \approx \sqrt{\mu_t r_i (\sigma_i^2 + (\mu_t - \mu_i)^2)} \quad (4.49)$$

Substituting equation (4.49) in (4.44), we can rewrite the coupled system of finite-difference equations (4.44)-(4.45) into a first-order linear non-homogeneous ordinary finite-difference equation

$$\mu_{t+\tau} = \mu_t + H(\mu_t) (1 - \mu_t \tau) + r_i \tau (\mu_i - \mu_t) \quad ; \quad H(\mu_t) = \sqrt{\frac{r_i}{\mu_t} (\sigma_i^2 + (\mu_t - \mu_i)^2)} \quad (4.50)$$

which describes the evolution of the posterior mean μ_t as new vesicles are released. The variable μ_t is the quantity of interests for the post-synaptic dendrite, since it is the best estimation of the hidden variable ν_s given the available information and the knowledge of the stochastic dynamics controlling the variable. The given equation can be rewritten in form of differential equation

$$\frac{d\mu_t}{dt} = H(\mu_t) \delta(t - t_{\text{release}}) - H(\mu_t) \mu_t + r_i (\mu_i - \mu_t) \quad (4.51)$$

The biophysics of receptors in the dendrite detect the release of a vesicle by sensing the presence of neurotransmitters in the synaptic cleft, and the received signal is converted into an excitatory postsynaptic current (EPSC). We propose that vesicles releases are translated into EPSCs according to equation (4.50). Specifically, we assume that the post-synaptic current I_t is a linear function of the decoded quantity $I_t = m\mu_t + n$. It is possible to note that, if we approximate H to be independent of μ_t , then we can write

$$\frac{d\mu_t}{dt} = H \delta(t - t_{\text{release}}) + (H + r_i) \left(\mu_i \frac{r_i}{r_i + H} - \mu_t \right) \quad (4.52)$$

The given equation can be interpreted as follows: Each time a vesicle is released, the concentration of neurotransmitters (and, consequently, synaptic current) increases of a fixed

quantity, which reflects into a positive jump of μ_i with amplitude H . Moreover, in absence of input drive (vesicle release), a continuous-in-time diffusion of neurotransmitters provides synaptic current to decay exponentially and, as a consequence, μ_t to decay to the equilibrium value $\mu_i \frac{r_i}{r_i+H}$ with timescale $H + r_i$. In this sense, we can interpret the original dependence of $H(\mu_t)$ on μ_t as a consequence of non-linear couplings between neurotransmitter concentration, ions concentration, instantaneous synaptic current and ions channels open/close configuration.

Another possible approximation we can make to investigate properties of the given equation is to consider the regime $\sigma_i^2 \gg |\mu_t - \mu_i|$ and to discard the term $\mu_t - \mu_i$ in equation (4.51), as follows

$$H(\mu_t) \approx \sqrt{\frac{r_i}{\mu_t} \sigma_i^2} \quad (4.53)$$

Such an approximation in a similar scenario is discussed and motivated in 3.2.5 and in fig. 3.1B, where we show that the error introduced by the approximation is relatively small in the region of interest. If an elevated number s of vesicles is released in a short time (and we neglect receptors saturation effects) the continuous-in-time decay of neurotransmitter will not significantly affect the postsynaptic potential immediately after the stimulus is received. Therefore, according to approximation (4.53), the resulting max value of the EPSC after the stimulation would satisfy the differential equation

$$\frac{d\mu}{ds} = \sqrt{\frac{r_i \sigma_i^2}{\mu_t}} \quad (4.54)$$

whose solution predicts a power law dependence of the resulting postsynaptic current on the number of vesicles released

$$\mu = \left(\frac{3}{2}\right)^{\frac{2}{3}} \left(c_1 + s\sqrt{r_i \sigma_i^2}\right)^{\frac{2}{3}} \quad (4.55)$$

where c_1 is an arbitrary constant, which depends on the initial value of the potential before the beginning of the stimulation.

4.4.2 Combining different synapses

According to our model, the postsynaptic membrane current I_t evoked in the dendrite represents a linear function of the mean μ_t of the posterior distribution $D(\nu_s)$, which is inferred by the neuron about the rate ν_s controlling synaptic release. The given rate is proportional to the probability of a certain event, which we assume to be containing some information about the state x_t encoded by the post-synaptic neuron.

According to 4.2.1, we propose that the post-synaptic neuron is going to use this information to implement a binary classification problem, and infer the probability P_t that the hidden event x_t is currently in the active state 1. To infer this probability, the neuron exploits information encoded in an observation vector ω , which is randomly sampled from a d -dimensional distribution $P(\omega|x)$. We propose that this observation vector ω is given by the postsynaptic current of different branches in the dendritic tree, corresponding to different synapses ($\omega_i = V_i$). Considering how the quantity I_t has been generated in our model, it is reasonable to assume that the marginal distribution ω_i can be well described by an offset Gamma distribution. The reason is that $\omega_i = I_i$ is a linear function of the mean μ_t , which is Gamma distributed, and the space of Gamma distributions is close under scaling. For analytical tractability, we will approximate here the Gamma distribution with a Gaussian. Different variables ω_i refer to distinct events in the world, which might be correlated in their nature. Therefore, the d -dimensional distribution $D(\omega)$ will be approximated by a multivariate Gaussian, with covariance matrix Σ and x -dependent mean vector $\mu(x)$. Again, for mathematical simplicity, we will assume for now that Σ is independent on the state x . In the following this approximation can be removed, and the consequences can be explored. We can recognize in this situation the same task that we considered in 4.2.2, and we can proceed in a similar fashion to derive the optimal solution (i.e. logistic regression). According to equation (4.10),

the log-odds probability ratio $\xi = L(P)$ is given by a linear sum of the observations ω_i

$$\xi = L(\text{Prior}) + a + b^T \omega \quad (4.56)$$

$$a = -\frac{1}{2}(\mu_0 + \mu_1)^T \Sigma^{-1}(\mu_1 - \mu_0) \quad (4.57)$$

$$b = \Sigma^{-1}(\mu_1 - \mu_0) \quad (4.58)$$

This result shows that, by doing a linear sum of the received dendritic currents with some appropriate weights, the neuron can implement optimal inference of the log-odds probability for the hidden state x . As we saw in 4.2.2, this is a promising result, since we know that neurons are often finely tuned to implement linear summation of most input combinations (Cash and Yuste, 1999). Additionally, it has been observed that neurons compensate for the morphological disposition of synapses in such a way that EPSCs evoked in distal dendrites in average affect the soma as much as EPSCs generated in dendrites proximal to the soma (Magee and Cook, 2000). This is compatible with the idea that weights used in the linear summation of EPSCs depend on the correlation structure Σ between the inputs, and not on their locations. Finally, as we saw in 4.2.2, non-linear summations can be explained by considering in our model a covariance matrix Σ which is conditional on x .

In equation (4.58) we consider the case of a single logistic regression iteration. Yet, in the neuron this inference is implemented as a continuous process: information about the hidden state is received in each moment in time, and the value of ξ is updated on-line accordingly. Thanks to the Bayesian framework used in (4.56), the equation can be implemented recursively. Let's assume for now that τ_c is the average timescale on which the observations ω are correlated. For now we can consider each vector $\omega(t)$ at intervals of τ_c as an independent observation, and write

$$\xi_{t+\tau_c} = \xi_t + a + b^T \omega_{t+\tau_c} \quad (4.59)$$

which can be rewritten in differential form as

$$\frac{d\xi_t}{dt} = \frac{1}{\tau_c}(a + b^T \omega_t) \quad (4.60)$$

Moreover to the likelihood provided by available observations, the stochastic dynamics of the hidden variable x_t introduces a time-dependent information loss. We model this stochastic dynamics as a Markov Chain process, following an approach first proposed in (Deneve, 2008). According to the proposed dynamics, a positive state x_1 will turn into a negative state x_0 with rate r_\downarrow , while a negative state will turn positive with rate r_\uparrow . According to the given dynamics, we can write how the probability P of x being in the positive state x_1 evolves in time

$$\frac{dP_t}{dt} = -P r_\downarrow + (1 - P) r_\uparrow \quad (4.61)$$

and, consequently

$$\frac{dL(P_t)}{dt} = \frac{d}{dt} \log \left(\frac{P}{1 - P} \right) = r_\uparrow \left(1 + \frac{1 - P}{P} \right) - r_\downarrow \left(1 + \frac{P}{1 - P} \right) \quad (4.62)$$

or, expressing everything as a function of $\xi_t = L(P_t)$

$$\frac{d\xi_t}{dt} = r_\uparrow (1 + e^{-\xi}) - r_\downarrow (1 + e^\xi) \quad (4.63)$$

By combining equations (4.60) and (4.63), we can write

$$\frac{d\xi_t}{dt} = \frac{1}{\tau_c}(a + b^T \omega_t) + r_\uparrow (1 + e^{-\xi}) - r_\downarrow (1 + e^\xi) \quad (4.64)$$

It can be noted that, if we neglect observations ω ($a, b = 0$), the value of ξ will converge to the solution of

$$r_\uparrow (1 + e^{-\xi}) = r_\downarrow (1 + e^\xi) \quad (4.65)$$

which is given by

$$P_{eq} = \frac{1}{1 + \frac{r_{\downarrow}}{r_{\uparrow}}} \quad (4.66)$$

According to equation (4.64), the membrane potential of a neuron will linearly integrate a constant flow of ions a/τ_c and all the received input currents ω_t with some appropriate weights b/τ_c . Furthermore, a potential dependent flow of ions will drive the membrane potential to converge to the equilibrium value P_{eq} . In (Deneve, 2008) it is shown that, under some assumptions, this potential-dependent flow of ions approximates the leakage term in the leaky integrate-and-fire model of neural dynamics, with the only difference that the timescale of exponential decay depends on the membrane potential, and specifically it is given by

$$\tau_{\xi} = \frac{1}{r_{\uparrow}e^{-\langle \xi \rangle} + r_{\downarrow}e^{\langle \xi \rangle}} \quad (4.67)$$

where $\langle \xi \rangle$ is the average value of the log-odds likelihood. According to the proposed model, the weights used by the neuron to integrate dendritic inputs ω depend on the mean values $\mu_1 - \mu_0$, on the covariance structure Σ and on the average timescale of autocorrelation τ_c . By implementing this dynamics, the membrane potential fluctuations of the neuron will infer and track the log-probability ratio of the probability P_t of the hidden state x_t .

The proposed model of membrane fluctuations (which has been derived from first Bayesian principles) is compatible with existing models of neural dynamics and, if validated, it would provide a computational explanation of what we know about neural integration of dendritic currents. In the following, we will see how the information encoded in the membrane fluctuations can be efficiently translated into spikes, in order to propagate it to distant regions of the neuron, and to coordinate different local processes.

4.5 Action potentials

There is a large variety of models describing how membrane fluctuations control action potentials generation. Some models are relatively simple, as it is the case of the integrate and fire model (Burkitt, 2006), which prescribes that every time the membrane potential V_t reaches a certain threshold V_{th} a spike is generated, and the membrane potential is reset to value V_o . Other more complicated models contemplate a dynamically changing threshold (Mihalaş and Niebur, 2009) or they introduce some random component to the membrane fluctuations dynamics to introduce stochasticity (Sacerdote and Giraudo, 2011). Some models substitute the concept of threshold with a probability of spiking $p(V)$ which depends on the membrane potential V_t (Jolivet et al., 2006). Some of these models are very detailed and biophysically realistic (Hodgkin and Huxley, 1952) while others are purely effective (Brette, 2005; Mihalaş and Niebur, 2009) or they provide a principled explanation of the computational role of the proposed dynamics in information encoding and processing (Deneve, 2008).

In the following we will not select any particular spike-initiation model. On the contrary, we will develop an abstract framework that can be applied to any chosen dynamics, as long as the used model produces a firing rate that reflects the membrane fluctuations of the neuron (as is the case in (Deneve, 2008)). We will exploit some common properties shared by various spike-initiation models, and we will apply the computational mechanism we developed in Chapter 3 to introduce spike-rate adaptation, which allows making the spiking activity of the neuron more efficient by implementing a predictive code.

In the literature, action potentials are usually interpreted as bits of a quantized code used by neurons to communicate information with each other. Spikes are commonly considered as the letters underlying neural language, and the information encoded in the spiking activity of neurons is investigated using either information theory analysis (Bhumbra and Dyball, 2005; Yarrow and Challis, 2012) or correlating neural dynamics with perception (Lee et al., 2013) and behavior (Paninski and Fellows, 2004). Here, we propose a slightly differently perspec-

tive: action potentials are generated for one precise reason, which is to propagate information from one region to the other within the neuron, and to coordinate and synchronize different local processes. In fact, the advantage of action potentials is that they are quantized and, as a consequence, their message is remarkably robust and consistent when propagated over long distances and in the presence of noise. However, action potentials have the disadvantage of being metabolically expensive (Laughlin et al., 1998). In addition, the dynamic range of spiking rate is limited on one side (low rates) by low consequent information capacity and on the other (high rates) by biophysical constraints, as neural refractory periods or the capacity of mitochondria to produce ATP. For this reason, neurons use an efficient coding scheme to propagate information with spikes and, specifically, we propose that they use a coding strategy that is predictive (Hosoya et al., 2005). Neurons internally decode information already encoded in their spiking history, and use this information to control spike initiation mechanisms. By doing this, only non-redundant new information is encoded by action potentials, which was not already encoded in previous action potentials.

4.5.1 Gain rate normalization

Some models of spike-generation translate membrane fluctuations into spikes and firing rates using a stochastic dynamics, while others are fully deterministic. Yet, in each model the neuron receives inputs through a stochastic process, which is synaptic vesicle release. Moreover, the input received is the last product of a (often long) processing pipeline. Each neuron receives a stochastic input, integrates it and passes it in form of a stochastic output to the next neuron in the chain. Variability introduced by each iteration accumulates and propagates. Therefore, it is reasonable to assume that the final relationship between neuron spiking activity (output of the processing pipeline) and the encoded feature of the world x (input) can be well described by an inhomogeneous Poisson process. This idea is compatible with the rate-coding approach to information encoding in the brain (Fuglevand and Winter, 1993), yet

other approaches are possible (Gautrais and Thorpe, 1998). In the following, we will accept and follow the rate-coding hypothesis.

We propose that the message (hidden rate) encoded in the spiking activity of the neuron is decoded and represented locally. This decoded message, which in Chapter 3 was referred to as c_t , will be represented here by the intracellular concentration of calcium ions $[c]$. It should be mentioned that the variable c_t (from Chapter 3) and $[c]$ (used in this Chapter) represent different quantities. In fact, they refer to two different levels of description and, as we will see in Section 4.6, they are controlled by different dynamics. For this reason, they should not be confuse, despite the fact that they play an analogous role in the two models. Recursively, the instantaneous calcium concentration will interact with spike generation mechanisms, and it will act as divisive gain normalization. Therefore, if a neuron is encoding a stimulus that in absence of adaptation would produce a firing rate ν , adaptation will act in such a way that the effective firing rate of the neuron will be ν^* , given by

$$\nu_t^* = \frac{\nu_t}{y([c])} \quad (4.68)$$

where $y([c])$ is a function of calcium concentration. The choice of $y([c])$ depends on the spike generation model we use, and on the assumptions we make about what variable (e.g. probability P or Log-odds ratio L) controls the average firing rate of the neuron. The case $y([c]) = [c]$ matches equation 3.15 we used in Chapter 3, and it corresponds with the circumstance of the log-odds ratio linearly controlling the firing rate. However, here we want to provide a general description compatible with different spike-initiation models and, for this reason, we don't select a particular function $y([c])$. The proposed divisive gain normalization can be implemented through different biophysical mechanisms, and the choice of the implementation will also depend on the specific model of spike generation we are working with. In the following we will propose a candidate implementation of this mechanism, and we will use a normative approach to predict the potassium ions channels dynamics that

allows to execute it.

4.5.2 Calcium dependent potassium current

The discussed gain rate normalization can be implemented through different neural mechanisms. For example, various models of neural dynamics propose a dynamic spike-initiation threshold. If working with such a model, an appropriate choice of a calcium-dependent threshold $V_{th}([c])$ would enable the neuron to implement the desired firing-rate regulation. Here, we propose an alternative implementation that exploits supplementary (adaptation) ions currents to control neural excitability and normalize spiking activity. In the proposed model we assume that the mentioned currents I_c are controlled by calcium concentration, and they act as to balance the predictable component of dendritic input I_ω as follows

$$I_{\text{total}} = I_\omega - I_c = b^T(\omega_t - E(\omega|P_t)) \quad (4.69)$$

where P represents the current knowledge of the neuron about the state of x that is already encoded in the spiking history. We can express the expected dendritic currents ω given the current knowledge P as

$$E(\omega|P) = P \int \omega P(\omega|x=1) d\omega + (1-P) \int \omega P(\omega|x=0) d\omega \quad (4.70)$$

$$= P\mu_1 + (1-P)\mu_0 \quad (4.71)$$

$$= \mu_0 + P(\mu_1 - \mu_0) \quad (4.72)$$

Moreover, it comes by definition that, solving equation 4.64 for $\omega = E(\omega|P)$, the resulting value will track the quantity L_c , which is the log-odds ratio of the probability P . We assume that the goal of such a mechanism is to infer the quantity $\Phi = \xi - L_c$. It is possible to show

that the given quantity can be inferred using the following differential equation

$$\frac{d\Phi_t}{dt} = \frac{1}{\tau_c}(b^T\omega_t - b^T(\mu_0 + P(\mu_1 - \mu_0))) - \Phi(r_\uparrow e^{-L_c} + r_\downarrow e^{L_c}) \quad (4.73)$$

The given equation was developed under the approximation that, in average, $|\xi - L_c| \ll |L_c|$. Such a condition is a consequence of the fact that both ξ and L_c track the same variable, which is the log-odds ratio of the hidden state, and, therefore, their relative difference should be smaller than their absolute value.

By substituting the definitions (4.57) and (4.58) in equation (4.73), the adaptation current I_c can be expanded as

$$I_c = -b^T(\mu_0 + P(\mu_1 - \mu_0)) = -(\mu_1 - \mu_0)^T \Sigma^{-1} \mu_0 - P(\mu_1 - \mu_0)^T \Sigma^{-1} (\mu_1 - \mu_0) \quad (4.74)$$

and equation 4.73 can be rewritten as

$$\frac{d\Phi_t}{dt} = \frac{1}{\tau_c}(I_\omega - I_c) - \Phi(r_\uparrow e^{-L_c} + r_\downarrow e^{L_c}) \quad (4.75)$$

The given equation represents a leaky integration of the received input currents, where the leakage term depends on calcium concentration (L_c). According to our model, this dynamics is implemented by the membrane potential of the neuron and, as we will see later, the resulting value is consequently used to control the probability of generating action potentials. Finally, information transmitted by spikes is decoded and represented locally in form of calcium concentration $[c]$. Therefore, calcium level $[c]$ represents a linear function of the best approximation of the log-odds ratio $L_c = L(P)$ given the information encoded in the spiking history of the neuron (equation (4.2)). This also implies that the current knowledge P of the neuron about the hidden state x can be expressed as

$$P = \frac{1}{1 + e^{-(h_o + h_1[c]_t)}} \quad (4.76)$$

and, as a consequence, we can express the calcium-dependent ions current in equation (4.74) as

$$I_c = -\frac{1}{1 + e^{-(h_o+h_1[c]_t)}}(\mu_1 - \mu_0)^T \Sigma^{-1}(\mu_1 - \mu_0) \quad (4.77)$$

Equation (4.77) describes the flow of ions predicted by our model, which compensates for the predictable (given spiking history) component of dendritic currents. Coincidentally, there is a class of potassium (K^+) channels that are regulated by calcium concentration. These calcium-dependent potassium channels were first discovered in (Gardos, 1958) and their existence was confirmed in (Gardos, 1958). These channels are usually referred to as "big conductance" (BK), "intermediate conductance" (IK) channels or "small conductance" (Sk) potassium channels, according to their ions permeability. calcium-dependent potassium channels (and especially BK channels) are known to regulate neuronal excitability, controlling interspike intervals and spike-frequency adaptation (Vergara et al., 1998; Gu et al., 2007). Additionally, SK channels contribute to control the medium afterhyperpolarization potential (mAHP) after spikes (Deister et al., 2009) and, therefore, to lower neuron's intrinsic excitability following sustained activity.

Calcium-dependent potassium channels C can take two possible states: "open" (C_1) or "close" (C_0). In the open state (C_1) channels become permeable to potassium, allowing K^+ ions to flow passively outside the neuron returning the electrochemical gradient, generated by a higher concentration of potassium ions inside the cell. The produced current drives the membrane potential V towards the potassium reversal potential, which is around -90 mV (Trussell and Jackson, 1985). This has the effect to hyperpolarize the cell, reducing excitability and probability of firing.

The dynamics of a channel C can be modeled using a statistical mechanics approach (Sullivan and Holmqvist, 1997). Driven by thermal fluctuations, the channel is free to oscillate between open and close states. To each state it is associated an energy level ϵ_i , which depends on the molecular structure of the open/close channel. Therefore, for a given temperature T , the

partition function Z can be written as

$$Z = \sum_{i=0,1} e^{-\frac{\epsilon_i}{KT}} \quad (4.78)$$

$$= e^{-\frac{\epsilon_0}{KT}} + e^{-\frac{\epsilon_1}{KT}} \quad (4.79)$$

and the probability of the channel being open is given by

$$P_{\text{open}} = \frac{e^{-\frac{\epsilon_1}{KT}}}{Z} \quad (4.80)$$

$$= \frac{e^{-\frac{\epsilon_1}{KT}}}{e^{-\frac{\epsilon_0}{KT}} + e^{-\frac{\epsilon_1}{KT}}} \quad (4.81)$$

which can be rewritten as

$$P_{\text{open}} = \frac{1}{1 + e^{-\frac{\epsilon_1 - \epsilon_0}{KT}}} \quad (4.82)$$

By multiplying the open probability P_{open} with the number of channels N and the average current generated by each open channel i , we can write an expression for the resulting current I_K

$$I_K = -Ni \frac{1}{1 + e^{-\frac{\epsilon_1 - \epsilon_0}{KT}}} \quad (4.83)$$

Where the minus sign is generated by the outward direction of positive charges $K+$ or, equivalently, by the highly negative reversal potential of potassium (-90 mV). In equation (4.82) we can recognize the same structure we derived in equation (4.77) for the current necessary to compensate for the expected dendritic drive. Specifically, if we assume that the difference of the energy levels $\epsilon_1 - \epsilon_0$ is linearly proportional to calcium concentration $[c]$

$$\frac{\epsilon_1 - \epsilon_0}{KT} = h_o + h_1[c]_t \quad (4.84)$$

the resulting current I_K (4.83) matches the necessary current I_c (4.77) for an appropriate choice of the parameters N and i . Our model predictions are compatible with experimental

observations (fig 4.7) from single channel recordings of calcium dependent potassium channels (specifically BK channels (Díez-Sampedro, 2006) and mSlo channels (Sullivan and Holmqvist, 1997)). It should be highlighted that the given ions channel dynamics was predicted by our model from first principles, and the existence of ions channels matching the necessary requirement is a promising result for the proposed model of neural functioning.

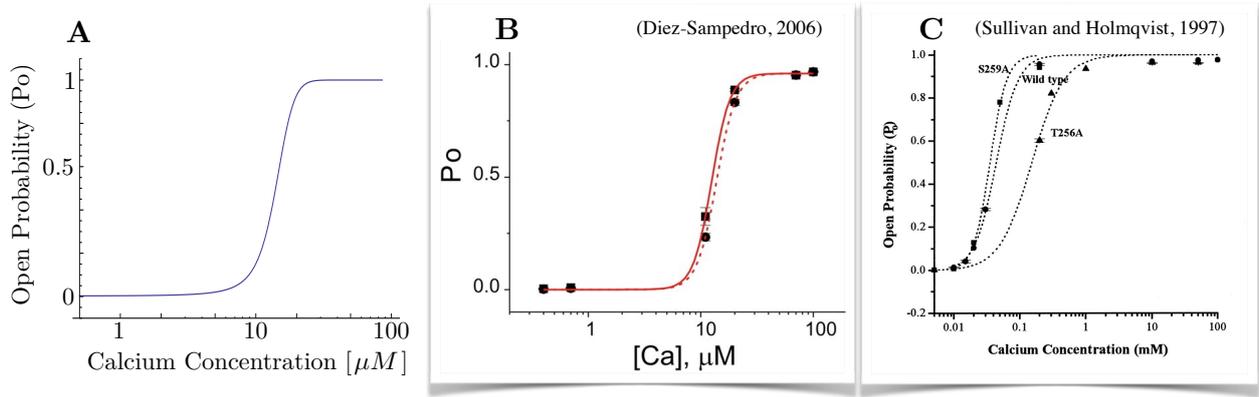


Figure 4.7: Open probability of a calcium dependent potassium channel as a function of local calcium concentration. Model predictions are compared with experimental observations. (A) Potassium channel open probability P , as predicted both by our normative model of potassium ions flow (equation (4.77)) and by the proposed statistical mechanics description of ions channels dynamics (equation (4.82)-(4.84)). Parameters used are $h_0 = 0.2$, $h_1 = 0.4$ and $[Ca_{ref}^{2+}] = 1\mu M$ (referring to equation (4.3)). (B) Single channel recordings of BK channels in Chinese hamster neurons. Different lines represent open probability for different mutations of BK channels at different Ca^{2+} concentrations (0.4, 0.7, 10, 20, 70, and 100 μM) measured at +40 mV . Figure adapted from (Díez-Sampedro, 2006). (C) Open probability for different mutations of mSlo potassium channels as a function of calcium concentrations. Values were measured at +20 mV . Figure adapted from (Sullivan and Holmqvist, 1997).

4.5.3 Action potentials initiation

We mentioned that the divisive spike-rate gain normalization mechanism in (4.68) is compatible with different neural spiking models, and according to the chosen mechanism the function $y([c])$ will assume a different form. Here, we want to discuss an example of such a mechanism. The model of spike initiation we introduce here is not adapted from the literature. Yet, its

dynamics is inspired by the interpretation of membrane fluctuations that our general model of neural functioning introduced. The proposed model of neural spiking has been developed as an alternative to existing models, but we want to highlight the fact that its validity is not required for the consistency of the general framework discussed in the previous sections.

In Section 4.4.2 we have seen that a linear function of the membrane fluctuations (Φ) can be interpreted as the difference between the Log-probability ratio $L(P)$ of a certain binary state x and its predictable component L_c , which is removed by calcium-dependent potassium currents. For clarity, in the following we will refer to this new probability as P^* and to the corresponding Log-odds ratio as $L^* = \Phi$. These quantities will satisfy

$$L^* = L - L_c \quad (4.85)$$

where $L = \xi$ represents the log-odds probability ratio inferred from the synaptic observations ω , and L_c is the log-odds ratio encoded by calcium concentration $[c]$. We propose that, at each moment in time, the spiking rate will depend on the probability P^* encoded by the membrane fluctuations. Specifically, we propose that there is a reference time interval τ_r , and that P^* represents the probability of generating an action potential in such a time window. τ_r can also be identified with the time between the beginning of the action potential and the end of the “absolute” refractory period of the neuron, since its inverse represents the maximum spiking rate of the neuron (corresponding to $P^* = 1$). Therefore, we can express the spiking activity of the neuron as an inhomogeneous Poisson process, whose instantaneous rate is given by

$$\nu^* = \frac{P^*}{\tau_r} \quad (4.86)$$

The given expression is an approximation. In fact, in a Poisson process the probability of an event happening in a certain finite time interval is not linearly proportional to the time duration. Moreover, the refractory and rising period will influence slow and high rate regimes differently. However, we will accept the given approximation assuming that the average firing

rate of the neuron is much smaller than the maximum firing rate $\nu_{max} = \tau_r^{-1}$

$$\langle \nu^* \rangle \ll \nu_{max} = \tau_r^{-1} \quad (4.87)$$

This assumption is biophysically reasonable, given that the average inter-spike interval of cortical neurons is usually within the range $50\text{ms} < \tau < 1\text{s}$ (Roxin et al., 2011) while the “relative” refractory period are usually within $7\text{ms} < \tau < 13\text{ms}$ (Chen et al., 2006) and the “absolute” refractory period is even smaller. According to the proposed model, the rate at which spikes are initiated is proportional to the logistic function of $\Phi = L^*$, which is a linear function of the membrane potential V .

$$\nu^* = \frac{1}{\tau_r} \frac{1}{1 + e^{-\Phi}} \quad (4.88)$$

The given model describes the probability (rate) of initiating an action potential. The response function described in equation (4.88) is compatible with existing effective models used in the literature (Keat et al., 2001) (Luo et al., 2012) (Geffen et al., 2009) (Kaardal et al., 2013). Furthermore, a logistic response function is often used when building Boltzmann machines (Neelakanta et al., 1991). Also, it has been shown that the logistic function corresponds to the maximum noise entropy (MNE) response of a neuron encoding a sensory stimulus (Fitzgerald et al., 2011), which maximizes entropy and minimizes model bias (equivalent to random field models in machine learning (Bouman and Shapiro, 1994)).

Implementation

We will interpret the logistic response function in 4.88 as the result of a fixed threshold model in the presence of Gaussian membrane-potential noise. We can assume that the coupled non-linear dynamics of ions channels and membrane potential in the neuron generates an unstable

fixed point V_{th} in the membrane potential V dynamics.

$$\left. \frac{dV}{dt} \right|_{V=V_{th}} = 0 \quad (4.89)$$

If V is lower than the given fixed point V_{th} , the dynamics of the system in absence of external inputs will drive V towards an equilibrium value $V_0 < V_{th}$. On the contrary, if $V > V_{th}$ a single cyclic oscillation (action potential) will be generated, that will eventually bring the value of V back below V_{th} . We assume that the value of V is a linear function of the log-odds ratio $\Phi = L^* = \text{Logit}(P^*)$ plus an additional random term

$$V_t = g_0 + g_1 \Phi_t + W_t \quad (4.90)$$

where W_t is a stationary Gaussian process. We will assume that, in each time \hat{t} , the marginal distribution $D(W_{\hat{t}})$ is normally distributed with zero mean $\mu_{\hat{t}} = 0$ and constant standard deviation $\sigma_{\hat{t}} = \sigma_w$.

$$D(W_{\hat{t}}) = N(0, \sigma_w^2, W_{\hat{t}}) \quad (4.91)$$

In addition, given a pair of times \hat{t}_1 and \hat{t}_2 , the multivariate normal distribution $D(W_{\hat{t}_1}, W_{\hat{t}_2})$ is completely determined by the covariance function

$$K(\hat{t}_1, \hat{t}_2) = \sigma_w^2 \rho_w(|\hat{t}_1 - \hat{t}_2|) \quad (4.92)$$

where ρ_w is a correlation function. For stationarity, we assume that the function ρ_w depends only the absolute difference between the considered times. Different covariance functions $K(t_1, t_2)$ will define different processes. For example, a delta distributed function

$$K(t_1, t_2) = \sigma_w^2 \delta_{t_1, t_2} \quad (4.93)$$

corresponds to Gaussian noise, while an exponential decaying covariance function

$$K(t_1, t_2) = \sigma_w^2 e^{-\frac{|t_1 - t_2|}{d}} \quad (4.94)$$

corresponds with an Ornstein-Uhlenbeck process, which is a stationary version of a Brownian motion. For our purposes, we will consider a correlation function ρ_w that decays on a timescale of τ_r . By doing this, we can approximate our process with a series of almost independent samples w_{t_i} distant τ_r one from the other. Therefore, if in a given time \hat{t} the Log-probability ratio is given by $\Phi_{\hat{t}}$, according to equation (4.90) the membrane potential can be expressed as

$$V_{\hat{t}} = g_0 + g_1 \Phi_{\hat{t}} + W_{\hat{t}} \quad (4.95)$$

and the probability of generating an action potential in the window $(\hat{t}, \hat{t} + \tau_r]$ will be given by

$$P^*(\Phi) = \int \theta(V_{\hat{t}} - V_{th}) P(W) dW \quad (4.96)$$

$$= \int \theta(W - (V_{th} - g_0 - g_1 \Phi)) P(W) dW \quad (4.97)$$

$$= \int_{V_{th} - g_0 - g_1 \Phi}^{\infty} P(W) dW \quad (4.98)$$

$$= \int_{V_{th} - g_0 - g_1 \Phi}^{\infty} \sigma_w N(0, 1, W) dW \quad (4.99)$$

$$= \int_{\sigma_w^{-1}(V_{th} - g_0 - g_1 \Phi)}^{\infty} N(0, 1, W) dW \quad (4.100)$$

$$= NCDF \left(\frac{V_{th} - g_0 - g_1 \Phi}{\sigma_w} \right) \quad (4.101)$$

where $NCDF$ is the cumulative distribution function of the standard normal distribution.

It can be shown that

$$NCDF(x) = \frac{1}{2} \left(1 + erf \left(\frac{x}{\sqrt{2}} \right) \right) \approx \frac{1}{1 + e^{-qx}} \quad (4.102)$$

for an appropriate choice of the parameter q , where erf represents the error function. Specifically, we selected the value of q that minimizes

$$q = \arg \min \int_{-\infty}^{\infty} \left(\frac{1}{2} \left(1 + erf \left(\frac{x}{\sqrt{2}} \right) \right) - \frac{1}{1 + e^{-qx}} \right)^2 \quad (4.103)$$

We used a gradient descent method to find the optimal value $q = 1.70099 \approx 1.7$. In figure 4.8 the two expressions in equation (4.102) are compared and their difference is plotted. From the figure it is possible to see that the logistic function is a very good approximation of the normal cumulative distribution function.

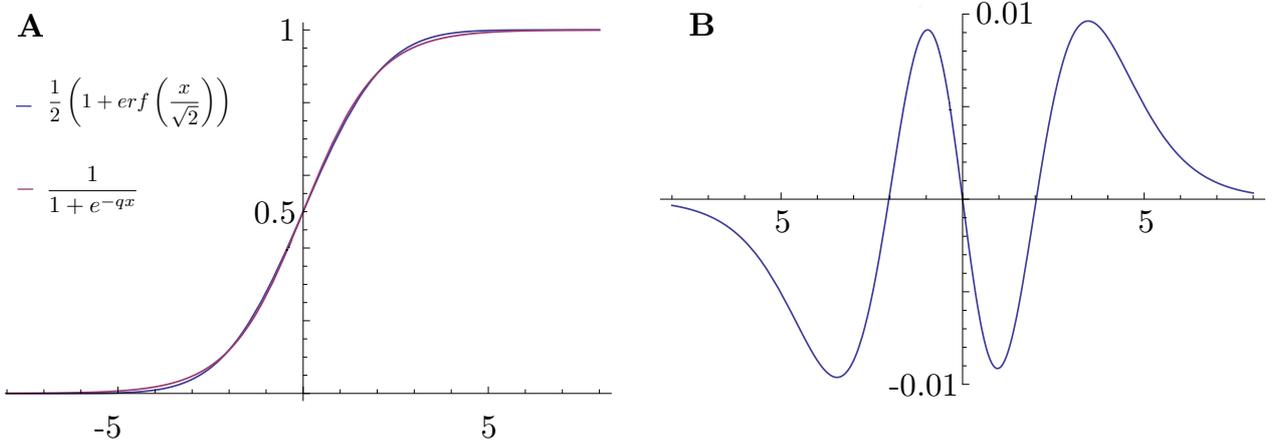


Figure 4.8: Graphical comparison of the expressions in approximation (4.102). (A) The logistic function (blue line) is a good approximation of the normal cumulative distribution function (NCDF - red line) for an appropriate choice of the parameter $q = 1.7$. (B) The difference between the two equations is plotted for $q = 1.7$.

Therefore, we can substitute the logistic equation to the NCDF in (4.101)

$$P^*(\Phi) \approx \frac{1}{1 + e^{\frac{1.7}{\sigma_w}(V_{th} - g_0 - g_1 \Phi)}} \quad (4.104)$$

and, for an appropriate choice of the parameters

$$g_0 = V_{th} \quad \text{and} \quad g_1 = \frac{\sigma_w}{1.7} \quad (4.105)$$

we can write

$$P^*(\Phi) \approx \frac{1}{1 + e^{-\Phi}} \quad ; \quad \nu^*(\Phi) \approx \frac{1}{\tau_r} \frac{1}{1 + e^{-\Phi}} \quad (4.106)$$

which matches the desired firing rate probability (equations (4.86)-(4.88)).

4.5.4 Integrating computational and algorithmic description

In Section 4.5 (and in Chapter 3) we developed an abstract (computational) model of spike rate adaptation. We propose that calcium concentration controls divisive gain normalization to implement predictive coding. The resulting instantaneous firing rate is given by

$$\nu_t^* = \frac{\nu_t}{y([c])} \quad (4.107)$$

where ν_t would be neural firing rate in absence of adaptation, and $y([c])$ is a function of current calcium concentration. The given function $y([c])$ depends on the quantity encoded in the firing rate of the neuron, and on the model of neural spiking we consider. For example, if the firing rate of the neuron is proportional to the log-odds ratio encoded in membrane fluctuations, $y([c])$ will be proportional to calcium concentration $[c]$, which encodes the Log-odds probability ratio inferred from the spiking history. Alternatively, if neural spiking reflects the probability P (as in the spike generation mechanism described in Section 4.5.3), $y([c])$ will correspond to the Logistic function of $[c]$ (i.e. the inferred probability).

In Sections 4.5.2 and 4.5.3 we proposed a more specific (algorithmic) model of how calcium concentration can interfere with spike-generation, and how information encoded in membrane fluctuations controls the initiation of action potentials. Specifically, we proposed that a calcium-dependent potassium current removes the predictable component of dendritic drive, and that the instantaneous probability (rate) of neural spikes is proportional to the

probability P encoded in neural fluctuations.

$$\nu^* = \frac{P^*}{\tau_r} = \frac{1}{\tau_r} \frac{1}{1 + e^{-\Phi}} \quad (4.108)$$

where Φ is a linear function of membrane potential V . We also described a simple biophysical (implementation) model of how this spiking behaviour can be achieved exploiting neural instability and intrinsic noise.

Here, we try to create a bridge between the different levels of analysis previously discussed, showing that the two approaches are compatible and consistent. We introduce a new quantity k such as

$$k = \tau_r \mu_o \quad (4.109)$$

where τ_r is the refractory period and μ_o is the desired average firing rate of the neuron, which depends on biophysical and metabolic constraints. As we saw previously, absolute refractory period is usually on the scale of 1 to few ms , while the average inter-spike interval of cortical neurons is usually within $[50ms < \tau < 1s]$. This means that $k \ll 1$, which is a property that we will exploit in the future. k is not a probability. Yet, since $0 < k < 1$, we are allowed to consider its logistic function L_k

$$L_k = \text{Log} \left(\frac{k}{1 - k} \right) \quad (4.110)$$

We will redefine the parameter g_0 in equation (4.105) as

$$g_0 = V_{th} + L_k \quad (4.111)$$

Exploiting equation (4.104), it is possible to show that this is equivalent to leave g_0 unchanged, and to redefine Φ as $\Phi + L_k$. Therefore, according to equation (4.85) we can write

$$\Phi = L^* = L - L_c + L_k \quad (4.112)$$

According to eq. (4.108), we can express the instantaneous firing rate of the neuron as

$$P^* = \frac{1}{1 + e^{-\Phi}} = \frac{1}{1 + e^{-L+L_c-L_k}} \quad (4.113)$$

We can expand

$$P^* = \frac{1}{1 + \left(\frac{1-P}{P}\right) \left(\frac{P_c}{1-P_c}\right) \left(\frac{1-k}{k}\right)} \quad (4.114)$$

$$= \frac{P(1-P_c)k}{P(1-P_c)k + (1-P)P_c(1-k)} \quad (4.115)$$

and, exploiting the regime $k \ll 1$, we can approximate

$$P^* \approx \frac{Pk(1-P_c)}{P_c(1-P)} \quad (4.116)$$

where P is the probability of the hidden state x inferred from the dendritic input, and P_c represents the same probability, as inferred from the spiking activity of the neuron. Therefore, we can assume that $P \approx P_c$. We also assume that both P and P_c are in average small. This last assumption is not completely justified from a mathematical prospective, but it is compatible with the idea that the world is composed of a very high number of events/features, and only a limited subset of them is occurring in every moment in time. Given these 2 assumptions we can approximate

$$P^* \approx \frac{Pk}{P_c} = \tau_r \mu_o \frac{P}{P_c} \quad (4.117)$$

which corresponds to an instantaneous firing rate

$$\nu^* = \frac{P^*}{\tau_r} \approx \mu_o \frac{P}{P_c} \quad (4.118)$$

We can note that the addition of the term $k - L_c$ to the log-odds ratio controlling spike

generation (which corresponds to a shift of a quantity $g_1(L_c - k)$ for the membrane potential) is compatible with the proposed divisive gain normalization rule (4.107), for

$$y([c]) = \frac{1}{k} \frac{1}{1 + e^{-L_c}} \quad \text{where} \quad k = \tau_r \mu_o \quad \text{and} \quad L_c = h_0 + h_1[c] \quad (4.119)$$

The term $k = \tau_r \mu_o$ provides the average firing rate of the neuron to match the desired value μ_o

$$E(\nu^* | \text{Spike History}) \approx E \left(\mu_o \frac{P}{E(P | \text{Spike History})}, \text{Spike History} \right) \quad (4.120)$$

$$\approx \mu_o \frac{E(P | \text{Spike History})}{E(P | \text{Spiking History})} \quad (4.121)$$

$$\approx \mu_o \quad (4.122)$$

In order to further validate the considered approximations, we numerically compared the desired abstract behaviour (equation 4.117) with the results of our spike generation mechanism (equations (4.108) and (4.112)). The detailed model accurately matched the abstract rule for small values of the probability P inferred from the dendritic drive, while it represented a reasonable approximation of the desired rule for higher values of P , especially in the region of interest $P_C \approx P$ (figure 4.9).

In this section, we developed a model of neural spiking activity on three different levels of description. On the computational level, we defined an abstract divisive spike-rate gain normalization rule, which improves the efficiency of neural spiking in transmitting information. On the algorithmic level, we proposed a mechanism that allows combining information received from the dendritic input (P) with information already encoded in the recent spiking history (P_c) by linearly summing their Log-odds ratios into the membrane potential. Moreover, we proposed that the probability of generating an action potential is controlled by the current membrane potential V , accordingly to a logistic response function. On the implementation level, we showed how ions channels can be used to translate information encoded in

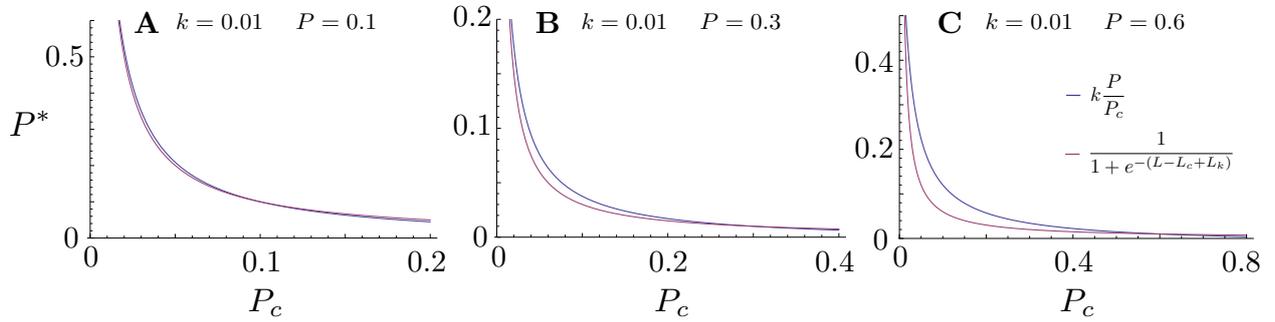


Figure 4.9: The desired gain normalization rule (equation 4.117) is compared to the behaviour of our model of spike initiation (equations (4.108) and (4.112)). As predicted, the detailed model matches the abstract rule for small values of the probability P (A, $P = 0.1$), while for larger values of the probability (B, $P = 0.3$)-(C, $P = 0.6$) the model is a reasonably good approximation of the desired behaviour. $k = \tau_r \mu_o = 3 \text{ ms} \times 0.03 \text{ ms}^{-1} = 0.01$. It can be noted that in the point of interest $P_c = P$ the 2 equations match in all the scenarios, while in the region $P_c \approx P$ the model is a good approximation of the abstract rule also for high values of P .

the pre-synaptic input and in calcium concentration into membrane fluctuations that match our algorithmic description. Furthermore, we provided a statistical mechanics description of ions channels dynamics that enables generating the desired flow of ions through the neural membrane. Finally, we showed how we can exploit the unstable dynamics of membrane potential as well as random membrane fluctuations (modeled using a Gaussian process) to generate the desired response function for spike initiation.

All these different levels of descriptions have been shown to be compatible and consistent (at least at a first order approximation), and the derived predictions about ions channel dynamics and neural spiking have been shown to match experimental observations about neurons.

The proposed model of neural spiking predicts that the initiation of an action potential is a highly variable process, controlled by membrane fluctuations. However, the divisive gain mechanism allows to increase information capacity of neural spiking, and to produce an accurate representation of the encoded variable P despite the intrinsic stochasticity of the process. If, for example, unlikely random fluctuations produce some unwanted spike events,

this will cause to overestimate the variable P_c ($P_c > P$). Yet, as a consequence, the high value of P_c will produce adaptation, and the activity of the neuron will be inhibited as long as the variable P_c doesn't converge back to the desired value P . This encoding scheme explains very high variability observed in neural spiking response to identical stimuli, as well as high Fano factors (Berry and Warland, 1997). At the same time, the proposed feedback mechanism enables the neuron to accurately and reliably represent the encoded variable. The given model proposes a new interpretation of the computational role of calcium, membrane potential, ions channel dynamics and random membrane fluctuations, which can potentially lead to a deeper understanding of neural computations and brain dynamics.

Up to now we assumed that the value of calcium concentration $[c]$ is a linear function of the Log odds ratio (L_c) of the probability P_c , which is the expected value of the probability P conditional to the recent spiking history of the neuron. In Section 4.6 we will see how P_c and L_c can be inferred optimally from the spiking train of the neuron using a Bayesian approach, and how this inference process can be implemented by the intercellular concentration of calcium, thanks to the non-linear dynamics of calcium channels.

4.6 Calcium dynamics

In Section 4.5 we introduced a model that describes how spikes are generated. We proposed that the instantaneous probability of firing is linearly proportional to encoded probability P , and that the spiking activity of the neuron can be modeled using an inhomogeneous Poisson process, with time-dependent rate ν^* given by equation (4.118)

$$\nu^* = \mu_o \frac{P}{P_c} \quad (4.123)$$

where μ_o is the desired firing rate of the neuron and P_c is the expected value of P conditional to the recent spike history $P_c = E(P|\text{Spike History})$. We assume that the Log-odds ratio of

the probability P_c is encoded in the neuron by calcium concentration, and its value is inferred on-line exploiting the non-linear dynamics of calcium channels.

Here, we derive an algorithmic model of how the value P_c can be inferred from the spiking history of the neuron. Using a Bayesian approach we derive from first principles a normative model that describes how P_c should evolve in time, according to observed spike events. We then derive how this dynamics can be converted into a flow equation for the Log-odds ratio $L_c = \text{Logit}(P_c)$ and, finally, we show how the given inference process can be implemented by calcium concentration and ions channels.

4.6.1 Inferring the instantaneous neural spiking rate

Previously in this Chapter (and more in detail in Chapter (3)) we derived a coupled system of finite difference equations ((4.44) and (4.45)) that allow to infer on-line the mean and variance of the posterior distribution $P(\nu)$ over the instantaneous firing rate ν of an inhomogeneous Poisson process

$$\mu_{t+\tau} = \mu_t + \frac{\sigma_t^2}{\mu_t} (1 - \mu_t \tau) + r_i \tau (\mu_i - \mu_t) \quad (4.124)$$

$$\sigma_{t+\tau}^2 = \sigma_t^2 + \frac{\sigma_t^4}{\mu_t^2} (1 - 2\mu_t \tau) + r_i \tau (\sigma_i^2 - \sigma_t^2 + (\mu_t - \mu_i)^2) \quad (4.125)$$

The given system of equations was derived as an near-to-optimal solution of the original inference problem, given some biophysically reasonable assumptions about the generative process underlying the dynamics of ν . Such a model was derived under the assumption that the instantaneous firing rate of the neuron is given by ν . However, in the considered case the instantaneous rate is given by

$$\nu^* = k \frac{\nu}{P_c} \quad \text{where} \quad \nu = \frac{P}{\tau_r} \quad \text{and} \quad k = \tau_r \mu_o \quad (4.126)$$

which can be rewritten as

$$\nu^* = \mu_o \frac{\nu}{\nu_c} \quad \text{where} \quad \nu_c = E(\nu | \text{Spike History}) = \mu_t \quad (4.127)$$

In this scenario, we can proceed in a similar fashion to what we did in Chapter 3 Section 3.2.4. In fact, an inter-spike interval τ observed in a Poisson process with rate ν corresponds with an observation $\tau_* = \tau \frac{\mu_t}{\mu_o}$ in a process with rate ν^* , and both observations bring a corresponding amount of information about the current value of the original variable ν . Therefore, we can replace τ with $\tau_* \frac{\mu_o}{\mu_t}$ and write

$$\mu_{t+\tau_*} = \mu_t + \frac{\sigma_t^2}{\mu_t} (1 - \mu_o \tau_*) + r_i \tau_* (\mu_i - \mu_t) \quad (4.128)$$

$$\sigma_{t+\tau_*}^2 = \sigma_t^2 + \frac{\sigma_t^4}{\mu_t^2} (1 - 2\mu_o \tau_*) + r_i \tau_* (\sigma_i^2 - \sigma_t^2 + (\mu_t - \mu_i)^2) \quad (4.129)$$

We can note that the substitution $\tau \rightarrow \tau_* \frac{\mu_o}{\mu_t}$ was made only in the first term of both equations (which corresponds to the gain of information conditional to the observed spike event), while τ was simply replaced by τ_* in the latter part of the expression (corresponding to the time-dependent information loss, caused by the stochastic dynamics of ν).

Proceeding as in Section 3.2.5 it is possible to show that the variance term σ_t^2 can be approximated as a function of the mean μ_t (equation (3.48))

$$\sigma_e^2 \approx \frac{9}{8} c_t \sqrt{\frac{r_i \sigma_i^2}{\mu_o}} \quad (4.130)$$

According to the given approximation, we can write

$$\mu_{t+\tau_*} = \mu_t + k_{m_i, \sigma_i^2, \mu_o, r_i} (1 - \mu_o \tau_*) + r_i \tau_* (\mu_i - \mu_t) \quad ; \quad k_{\sigma_i^2, \mu_o, r_i} = \frac{9}{8} \sqrt{\frac{r_i \sigma_i^2}{\mu_o}} \quad (4.131)$$

which can be rewritten in form of differential equation

$$\frac{d}{dt}\mu_t = k_{m_i, \sigma_i^2, \mu_o, r_i} (\delta(t - t_{\text{spike}}) - \mu_o) + r_i(\mu_i - \mu_t) \quad (4.132)$$

4.6.2 Calcium ions flow

Equation (4.132) can be used to derive the evolution equation for the quantity $P_c = \nu_c \tau_r = \mu_t \tau_r$

$$\frac{d}{dt}P_{c, t} = \tau_r k_{\sigma_i^2, \mu_o, r_i} (\delta(t - t_{\text{spike}}) - \mu_o) + r_i(\mu_i \tau_r - P_{c, t}) \quad (4.133)$$

and the flow equation for the Log-odds ratio $L_c = \text{Logit}(P_c)$. In fact

$$L_c = \text{Log} \left(\frac{P_c}{1 - P_c} \right) \quad (4.134)$$

$$\frac{d}{dP_c} L_c = \frac{1}{P_c - P_c^2} \quad (4.135)$$

and, therefore

$$\frac{d}{dt} L_{c, t} = \frac{1}{P_c - P_c^2} \left[\tau_r k_{\sigma_i^2, \mu_o, r_i} (\delta(t - t_{\text{spike}}) - \mu_o) + r_i(\mu_i \tau_r - P_{c, t}) \right] \quad (4.136)$$

$$= \frac{1}{\frac{1}{1+e^{-L_{c, t}}} - \frac{1}{(1+e^{-L_{c, t}})^2}} \left[\tau_r k_{\sigma_i^2, \mu_o, r_i} (\delta(t - t_{\text{spike}}) - \mu_o) + r_i(\mu_i \tau_r - \frac{1}{1+e^{-L_{c, t}}}) \right] \quad (4.137)$$

$$= (1 + e^{-L_{c, t}}) (1 + e^{L_{c, t}}) \left[\tau_r k_{\sigma_i^2, \mu_o, r_i} (\delta(t - t_{\text{spike}}) - \mu_o) + r_i(\mu_i \tau_r - \frac{1}{1+e^{-L_{c, t}}}) \right] \quad (4.138)$$

$$= 2 (1 + \cosh(L_{c, t})) \left[\tau_r k_{\sigma_i^2, \mu_o, r_i} (\delta(t - t_{\text{spike}}) - \mu_o) + r_i \mu_i \tau_r \right] - r_i (1 + e^{L_{c, t}}) \quad (4.139)$$

Similarly, using equation (4.2) we can derive the flow equation for the normalized calcium concentration $[c] = h_1^{-1}(L_{c,t} - h_o)$

$$\begin{aligned} \frac{d}{dt}[c]_t = & 2h_1^{-1} (1 + \cosh(h_o + h_1[c]_t)) \left[\tau_r k_{\sigma_i^2, \mu_o, r_i} (\delta(t - t_{\text{spike}}) - \mu_o) + r_i \mu_i \tau_r \right] \\ & - h_1^{-1} r_i (1 + e^{(h_o + h_1[c]_t)}) \end{aligned} \quad (4.140)$$

The given calcium dynamics can be potentially implemented by calcium channels dynamics. In fact, the intracellular calcium concentration $[c]$ predicted in (4.140) can be interpreted as the result of:

- a continuous in time ions flow of ions I_b

$$I_b \propto 2h_1^{-1} \left[-\tau_r k_{\sigma_i^2, \mu_o, r_i} \mu_o + r_i \mu_i \tau_r \right] - h_1^{-1} r_i \quad (4.141)$$

- an action potential dependent flow I_{ap}

$$I_{ap} \propto 2h_1^{-1} \tau_r k_{\sigma_i^2, \mu_o, r_i} \delta(t - t_{\text{spike}}) \quad (4.142)$$

- a calcium dependent flow I_c

$$I_c \propto 2h_1^{-1} \cosh(h_o + h_1[c]_t) \left[-\tau_r k_{\sigma_i^2, \mu_o, r_i} \mu_o + r_i \mu_i \tau_r \right] - h_1^{-1} r_i e^{(h_o + h_1[c]_t)} \quad (4.143)$$

- a flow I_{c-ap} , which depends both on action potentials history $\{t_{\text{spike}}\}$ and on current calcium concentration $[c]$

$$I_{c-ap} \propto 2h_1^{-1} \cosh(h_o + h_1[c]_t) \tau_r k_{\sigma_i^2, \mu_o, r_i} \delta(t - t_{\text{spike}}) \quad (4.144)$$

Equation (4.140) describes how calcium concentration should evolve in order to decode near-

to-optimally the information encoded in the spiking history. This allows to use action potentials to reliably propagate information through the neuron, to decode and represent it locally in form of calcium concentration and to use it to control local processes, as for example spike rate adaptation, synaptic release and synaptic plasticity.

The proposed model requires calcium dynamics to be consistent (i.e. linearly proportional) in different regions of the neuron. The neuron should coordinate the parameters controlling calcium flow in different areas, in order to synchronize calcium concentration (and the represented message) through the neuron. Local calcium concentration will be given by

$$[Ca^{2+}]_t = [c]_t [Ca^{2+}]_{ref} \quad ; \quad [c]_t = h_1^{-1}(L_c - h_0) \quad (4.145)$$

where $[Ca^{2+}]_{ref}$, h_1 and h_0 can vary from one region of the neuron to the other. Experimental results support this hypothesis (see fig. 4.10) showing that action potentials propagating through the axon evoke a congruous calcium response in different locations (Schiller et al., 1997), and that the same happens for action potentials back-propagating in various branches of the dendritic tree (Golding et al., 2001). This provides the time-course of calcium concentration in different regions to be consistent, and linearly proportional to a baseline value (Augustine et al., 2003).

Additionally, according to the proposed model, calcium concentration plays a paramount role in neural dynamics, as it encodes inside the neuron the state of the external world, and it controls neural communication. For this reason, it is reasonable to assume that modifying calcium dynamics (e.g. by genetically mutating proteins controlling ions channels functioning) will have significant consequences on animal behaviour. This has been confirmed by various experimental studies, which showed for example that altering calcium channels mechanics (through G-alpha protein mutations) affects the steering behavior in *C. Elegans* (Kato et al., 2014b).

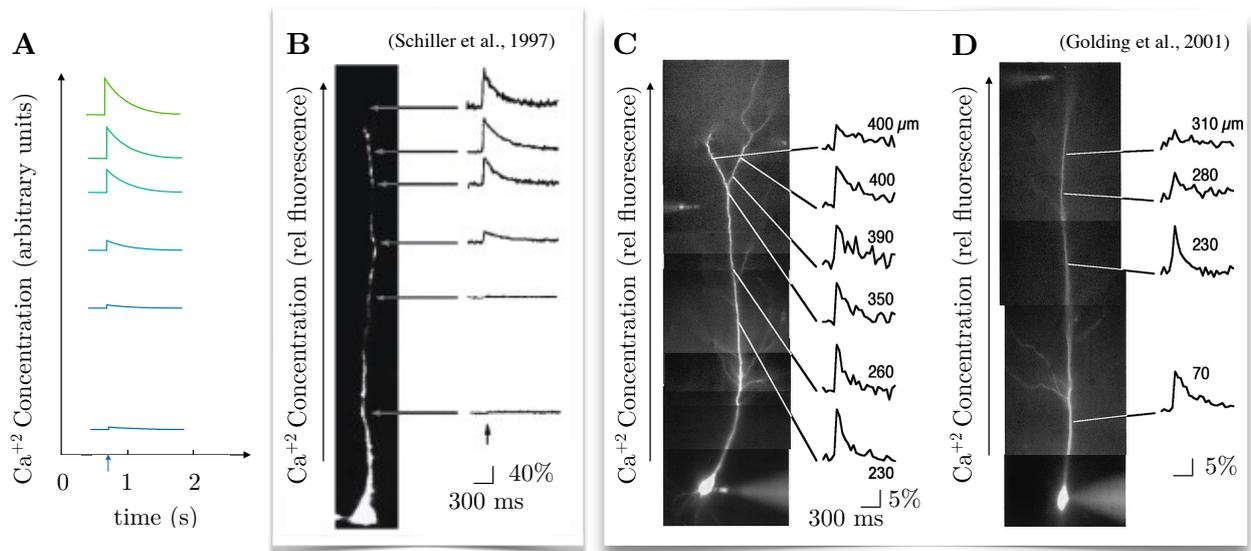


Figure 4.10: Spike evoked calcium concentration is consistent in proximal and distal regions of the neuron. (A) Qualitative model predictions about the time-course of calcium concentration. We predict that spike evoked calcium concentration in different parts of the neuron should be linearly proportional to a baseline value (equation (4.145)). (B) Calcium action potential, propagating along the optical axon of a cortical neuron after synaptic stimulation. Results are expressed as relative change in fluorescence ($\Delta F/F$). Figure adapted from (Schiller et al., 1997). (C-D) Calcium influx caused by strong (C) and weak (D) back-propagation of an action potential in the dendrites of different pyramidal neurons. Figure adapted from (Golding et al., 2001).

4.6.3 Local learning of calcium (hyper) parameters and model predictions

The proposed model for calcium evolution (equation (4.140)) is different (and more complex) than equation (3.50) derived in Chapter 3. Yet, the two models can be considered compatible and equivalent. In fact, equation (3.50) describes the evolution of an abstract quantity (algorithmic level), which is the expected average firing rate of the neuron in the absence of adaptation. On the contrary, (4.140) describes the flow of calcium (implementation level), which physically represents the desired abstract value. By following the given equation, calcium allows to implicitly encode the expected firing rate, and to normalize the instantaneous

firing rate consistently with the abstract model. For this reason, all the considerations we developed in Chapter 3 also apply to the model of calcium dynamics proposed here.

Specifically, it is possible to adapt the optimal parameters of calcium dynamics k , r_i and μ_i to the experienced stimulus and its dynamical properties (see Section 3.2.6). We can do this by implementing the system of coupled differential equations (3.61)-(3.62)-(3.63), which enables the neuron to learn the parameters controlling calcium evolution locally and online.

$$\begin{aligned} \frac{d}{dt}[c]_t = & 2h_1^{-1} (1 + \cosh(h_o + h_1[c]_t)) \left[\tau_r \hat{k}_t (\delta(t - t_{\text{spike}}) - \mu_o) + \hat{r}_t \hat{\mu}_t \tau_r \right] \\ & - h_1^{-1} \hat{r}_t (1 + e^{(h_o + h_1[c]_t)}) \end{aligned} \quad (4.146)$$

$$\frac{\hat{\mu}_t}{dt} = \alpha_\mu \delta(t - t_{\text{spike}}) \left(\frac{1}{\tau_r} \frac{1}{1 + e^{-(h_1[c]_t + h_o)}} - \hat{\mu}_t \right) \quad (4.147)$$

$$\frac{\hat{r}_t}{dt} = \alpha_r \left(\frac{1}{\tau_r} \frac{1}{1 + e^{-(h_1[c]_t + h_o)}} - \hat{\mu}_t \right) \quad (4.148)$$

$$\frac{\hat{k}_t}{dt} = \alpha_k ((t_{\text{spike } i} - t_{\text{spike } i-1}) - \delta(t - t_{\text{spike}}) \mu_o^2) \quad (4.149)$$

The robustness of the proposed system of equations has been tested in-silico (figure 3.2). The remaining parameters μ_o , τ_r , h_o and h_1 represents, respectively, the desired firing rate of the neuron (μ_o), the refractory period (τ_r) and the scaling and offset parameters of calcium concentration (h_o - h_1). All these quantities are fixed by the neuron, and some of these can change in different regions of the neuron (i.e. h_o and h_1) without affecting the consistency of the encoding scheme. Therefore, by implementing the given system of coupled equations the neuron can generate a local representation of the hidden variable in different regions of the neuron, and at the same time it can adapt the parameters controlling calcium flow to the dynamical properties of the encoded stimulus. This allows to adapt to the experienced stimulus in order to improve the internal representation of the world, and to have different consistent representations of the hidden variable in different regions of the neuron. We propose that these different representations (relative to different linear transformation parameters h_o and h_1) reflect the accuracy-energy-physiology requirements of different neural regions, and

depend on the functional role of calcium in controlling the respective local processes.

Moreover, since the two levels of description (algorithmic and implementation) are equivalent, the observable model predictions derived in Chapter 3 (Section 3.3) also apply to the model of neural dynamics proposed here. For example, the given model predicts the time-course of short-term neural spike-rate adaptation (Section 3.3.1), and it links adaptation dynamics with the time structure of the encoded input (equation (3.78)). The given predictions were tested on experimental results from cat V1 in-vivo recordings (Benucci et al., 2013)(figure 3.6), and the model was able to capture timescale and shape of spiking-rate adaptation using only the known dynamical properties of the stochastic stimulus used in the experiment (figure 3.7).

4.7 Discussion

In this Chapter, we proposed a model of neural dynamics. We discussed our model on the three Marr's levels of analysis (computational, algorithmic and implementation) (Marr and Poggio, 1976), and we showed how different levels of descriptions are consistent and equivalent. For this reason, the given model suggests (new) computational interpretations of known bio-chemical neural quantities and mechanisms, such as synaptic vesicle release, intra-cellular calcium concentration and action potentials. Specifically, we propose that synaptic vesicle release takes the form of a sampling process, which is controlled by the probability encoded in calcium concentration. Consequently, we interpret dendritic integration as logistic regression, which allows inferring the probability of a hidden state x exploiting the information encoded in the synaptic inputs. The inferred probability (and specifically its log-odds ratio) is encoded in form of intracellular calcium concentration. Action potentials are used to control local calcium concentration in different neural regions, and to propagate information through the neuron. This is done exploiting an efficient coding scheme (predictive coding)

which communicates only new (non-redundant) information. The proposed dynamics explains large variability observed in neural spiking activity and, at the same time, it provides to reliably and efficiently communicate information. We also show how calcium channels allow to locally decode neural spiking. According to the proposed model parameters controlling calcium evolution (e.g. decay timescale) depend on the dynamical properties of the encoded stimulus (e.g. timescale of changes), and we show how these parameters can be learned online and locally by the neuron.

We show how these computations can be implemented by a biophysical neuron exploiting neurotransmitter-activated ions channels, linear dendritic integration, calcium-activated potassium channels, unstable membrane dynamics, random membrane fluctuations and potential-dependent calcium channels. All the biophysical requirements of our model are compatible with experimental observations. Furthermore, the given model generates various observable predictions, e.g. synaptic short-term potentiation and spike-rate adaptation, which matched experimental results.

The new interpretation of known biophysical quantities proposed in this Chapter can potentially lead to a deeper understanding of neural computations, and it can inspire models of brain dynamics on a population level. For example, the proposed model suggests that dendrites implement binary regression to infer the state of a hidden variable x , and that calcium concentration encodes the log-odds ratio of the probability of such a variable. This can lead to the derivation of a synaptic plasticity rule that enables the system to update the integration weights on-line, and to implement dynamic logistic regression (McCormick et al., 2012) (Penny and Roberts, 1999). Such a synaptic plasticity rule would depend on post-synaptic calcium concentration (as observed in (Zucker, 1999; Castellani et al., 2001; Castellani, 2005)), and it would act as coincidence detection, compatibly with Hebbian learning (Caporale and Dan, 2008).

In next Chapter, we will focus on neural network dynamics and learning. In order to develop

a model of neural behaviour on the population level we will use a simplified model of neural dynamics (compared to the one discussed in this Chapter). This will reduce mathematical complexity and improve simulations tractability. However, we will base our neural population model on some of the concepts introduced in this Chapter, such as the idea of neural sampling, and the instrumental role of noise and random fluctuations in brain functioning.

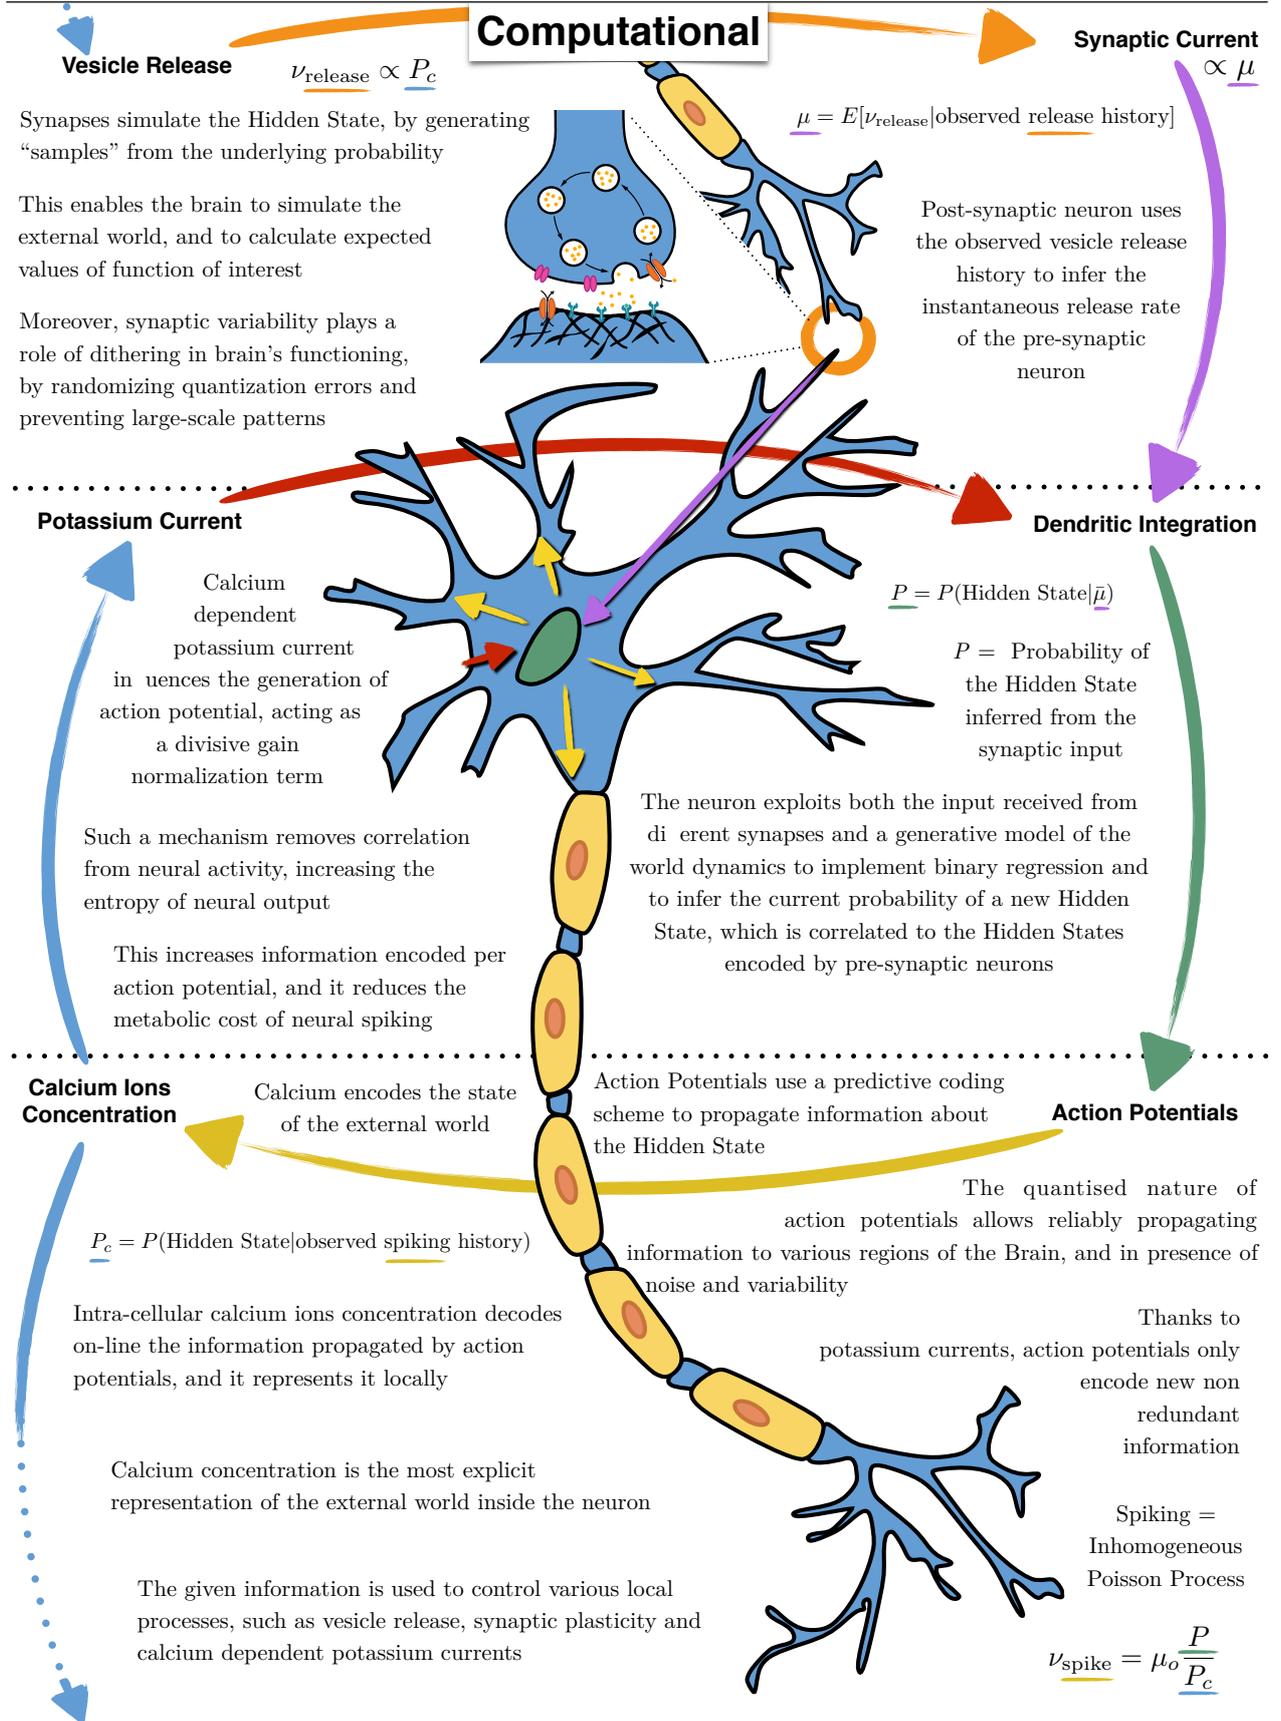


Figure 4.11: Calcium Coding Theory. Computational description of neural functioning.

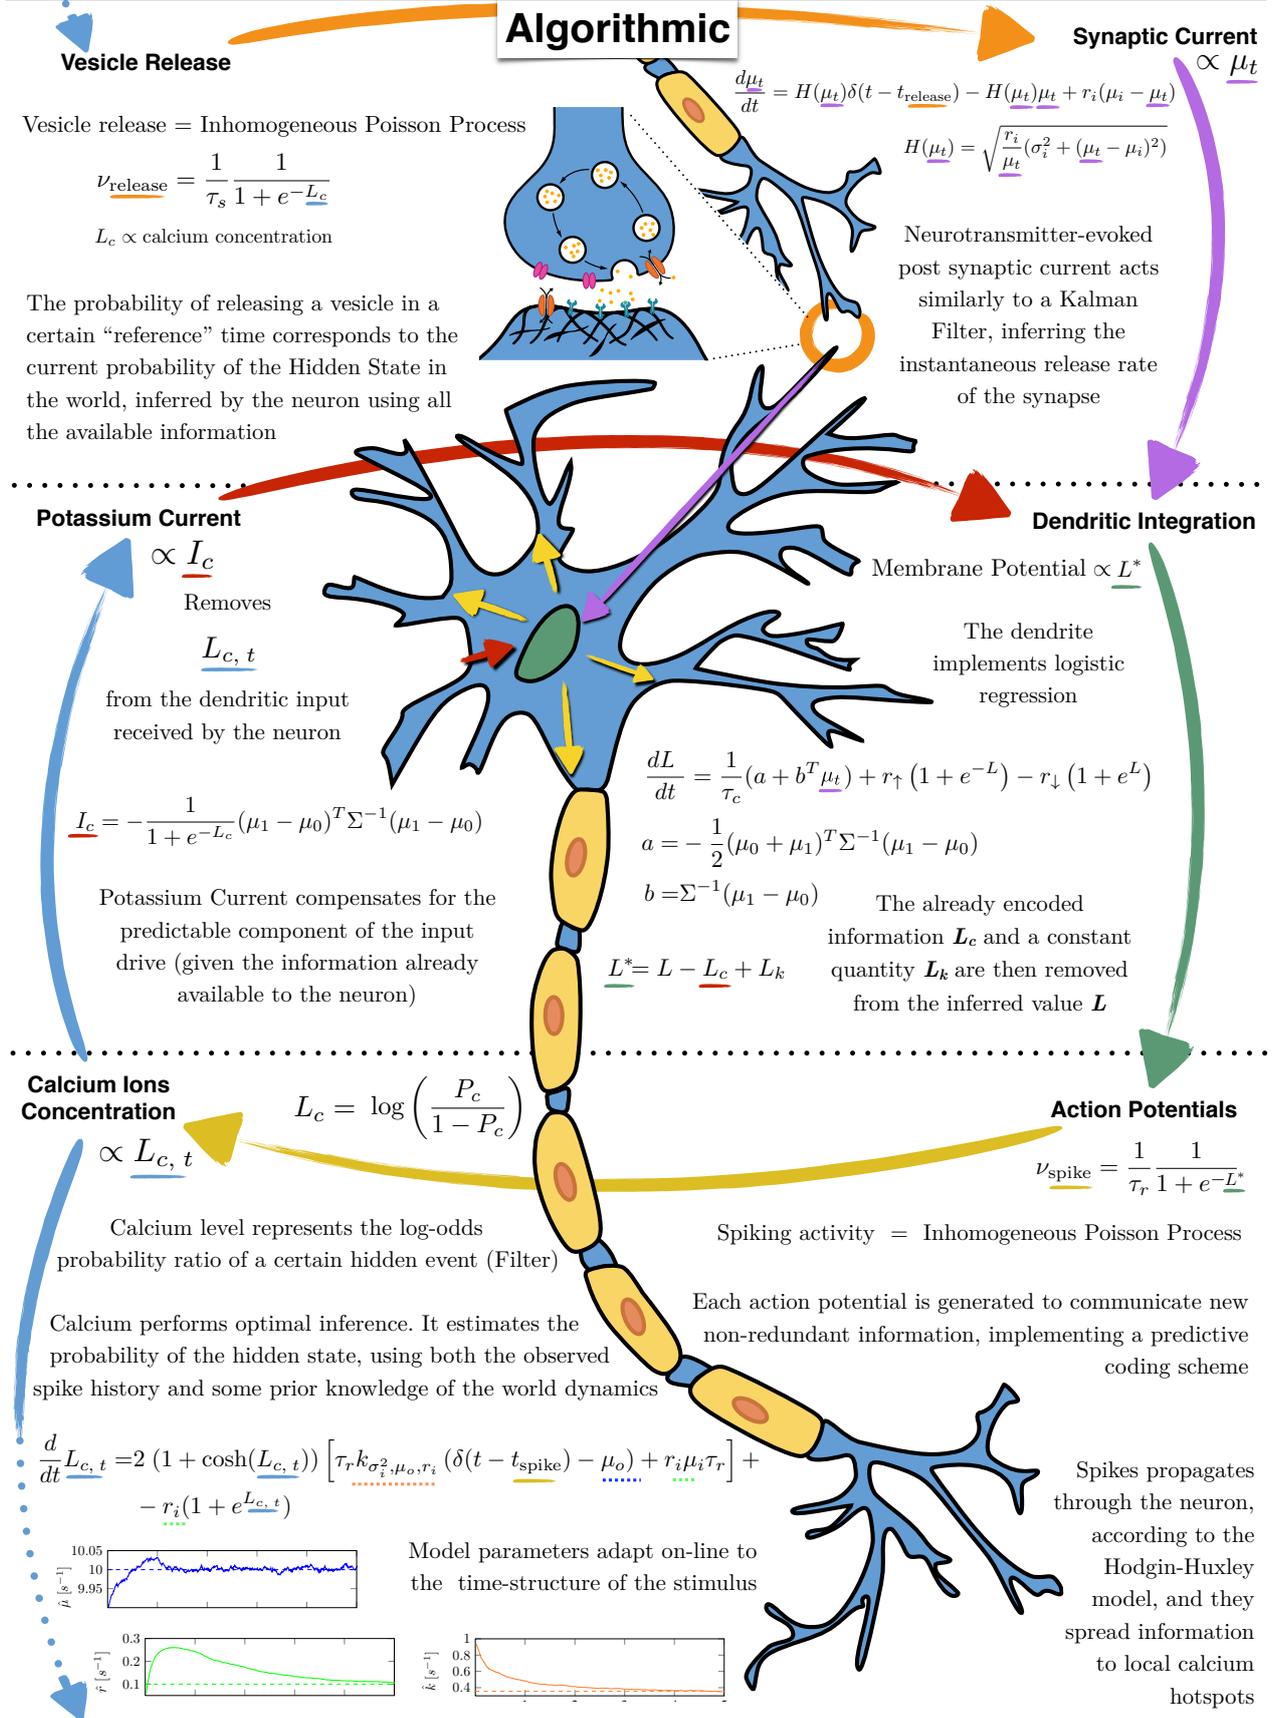


Figure 4.12: Calcium Coding Theory. Algorithmic description of neural functioning.

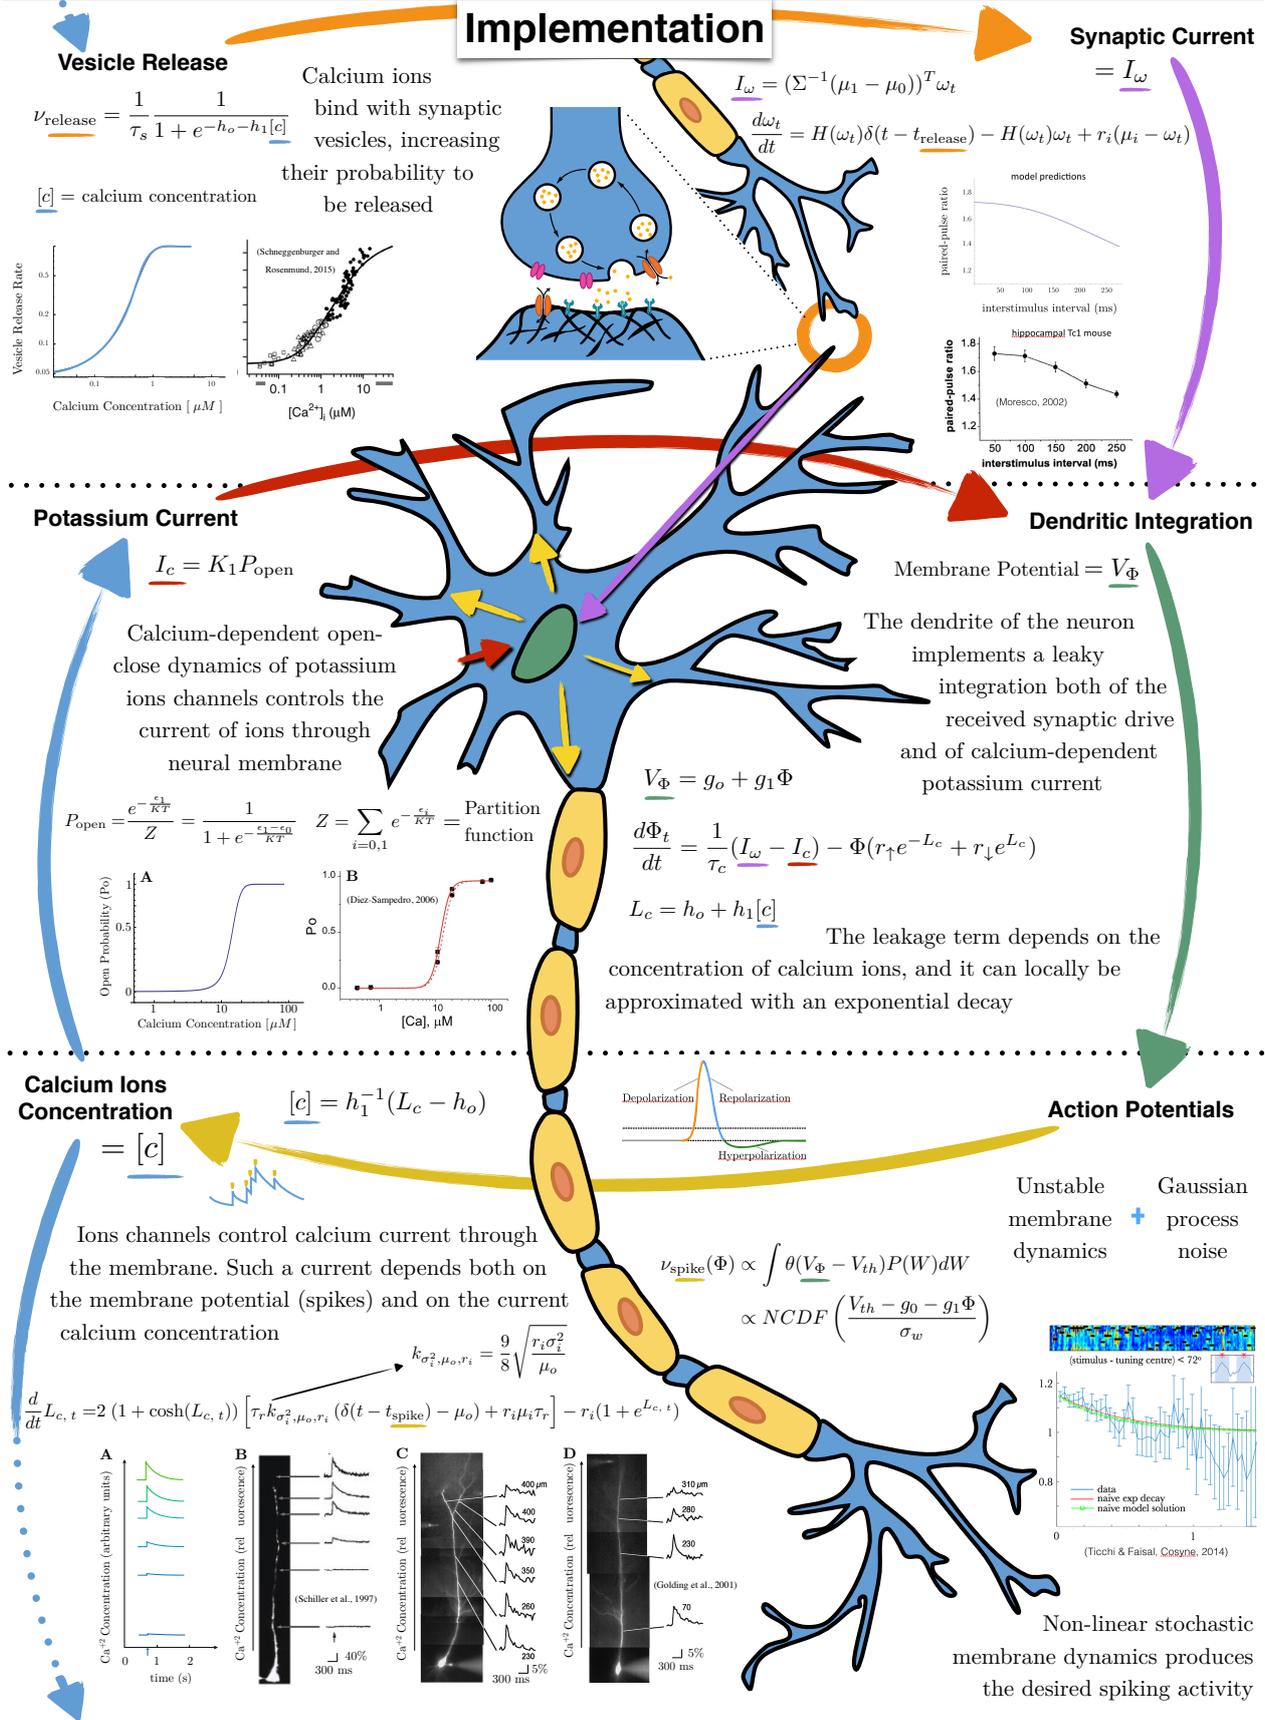


Figure 4.13: Calcium Coding Theory. Implementation description of neural functioning.

Chapter 5

Inference by sampling from noisy neuronal circuits

‘Pictures, propagated by motion along the fibers of the optic nerves in the brain, are the cause of vision.’

Isaac Newton, “Opticks”, 1704

5.1 Introduction

Von Helmholtz first introduced the idea that perception may be interpreted as the continuous unconscious process of inferring causes from their sensory effects and, following this interpretation, Barlow successively suggested that neural activity of sensory neurons may reflect the belief that stimuli are present in the world. Since then, Bayesian Decision Theory (Körding and Wolpert, 2003) has emerged as a principled way to describe our perception and behavior under uncertainty, this being supported by ample psychophysics and behavioral evidence (Wolpert et al., 1995; Ernst and Banks, 2002b; Körding and Wolpert, 2004).

Bayesian Decision Theory normatively predicts that different sensory cues sharing a common cause should be combined optimally by weighting each cue with the corresponding sensory variability. Moreover, it predicts that also prior knowledge and sensory information, presented as probabilities, should be combined following Bayes theorem, which requires the brain to encode prior knowledge of the world in some form. This information can be stored by synaptic connectivity patterns in designated regions of the brain but also, as highlighted by recent electrophysiological evidence (Berkes et al., 2011), could be encoded by the spontaneous activity of the same neurons which encode the state of the variable in the presence of sensory cues.

While the neuronal correlates of cue integration and inference remain elusive (but see (Kepecs et al., 2008)), a family of models under the name of Probabilistic Population Codes have lit the path to how neuronal populations can perform such computations (Ma et al., 2006). These models rely on being driven by feed-forward sensory information processing and perform inference by integrating neural activity under appropriate encoding and decoding assumptions. Here, we propose an alternative model that performs inference by sampling exploiting the inherently noisy activity of neurons to actively generate possible solutions to the inference problem (Ticchi and Faisal, 2012, 2013a,b). Specifically, our model is able to (a) learn the prior statistics of the encoded stimulus. (b) In the absence of sensory drive, it generates spontaneous activity that matches the stimulus statistics previously learned by the population. (c) It evaluates the likelihood of the current internal representation and (d) it combines this additional source of information with the prior knowledge, to generate neural activity that can be interpreted as a sequence of samples from the posterior distribution. Broadly speaking, our model combines feed-forward (sensory) information with feedback information derived from evaluating the likelihood of its inherent noise-driven stochastic activity. For this reason, the proposed model can work in synergy with PPC, and will be extremely valuable whenever a feed-forward processing of sensory information is not sufficient to infer causes from observed consequences. In this situation, our model allows inferring causes by actively

and stochastically proposing guesses, which are then tested on sensory information.

Multiple sources of noise have been shown to contribute to cellular and behavioral trial-to-trial variability in the brain (Faisal et al., 2008), this representing a fundamental problem for signal processing and drastically affecting the reliability of neural computations (Faisal et al., 2005). Yet, evolution may have found a way to use these molecular sources of noise as a resource. For example in the effect of Stochastic Resonance (Longtin, 1993; Faisal et al., 2008; McDonnell and Abbott, 2008) noise can be used to boost detection of weak sensory signals, which is useful when false positive detection has smaller consequences evolutionarily speaking than false negative.

We show here how noise can be used for the purpose of enabling inference and decision-making. A broad class of inference algorithms, usually referred to as Monte Carlo Methods, exploits random number sampling for simulating and investigating stochastic real world problems that can't be solved analytically. For these algorithms, noise and variability represent a precious resource, to the extent that when these algorithms are implemented in-silico most of the computational time is spent in the generation of random variables. We show how, in this perspective, neural variability can be exploited for efficiently generating random variables in the brain, in order to implement this kind of algorithms. In particular, we will focus on a specific class of Monte Carlo algorithms, named Markov Chain Monte Carlo (MCMC), which allows performing optimal inference by generating a sequence of samples from the inferred posterior distribution.

In order to perform Bayes optimal computations, it is first necessary to learn the statistics of the sensory input (prior distribution) and to store this information in the brain for future computations. Experimental evidence shows that brains are able to learn and update their prior knowledge about the world, yet the mechanism controlling how this information is learned, stored and accessed is still unclear. Here, we assume that this information is stored within the same populations that encode the state of the world in the presence of sensory cues.

In fact, behavioral and physiological results prove that neurons encoding information about the state of the world are adapted to the stimulus statistics (Laughlin, 1983, 1981; Berkes et al., 2011). Specifically, it has been shown that neural encoders act as to maximize the mutual information between the state of the world and neural activity, under biological and metabolic constraints (Ganguli and Simoncelli, 2010). We propose a biophysically plausible model of neuronal plasticity and dynamics that, by exploiting subcellular sources of noise, enables a cortical population to encode incoming information efficiently and, concurrently (simultaneously), to learn the statistical proprieties of its sensory input (Prior). Specifically, we provide a model of how this adapted encoder can be learned from the input statistics using a Local Hebbian synaptic plasticity rule, and we show how the stored information can be subsequently extracted for further computations. Furthermore, we propose a biophysically plausible implementation of a MCMC algorithm in spiking neurons, which allows to combine the learned prior knowledge with additional sources of information, and to generate samples from the inferred Bayesian posterior distribution. The model is compatible with experimental results in the literature that match neural physiology, behavior, and stimulus statistics.

The neural population model here presented could potentially be integrated with the single neuron dynamics discussed in Chapter 4, producing consistent results. However, the two models require different levels of detail in the description of neural dynamics. Additionally, the single neuron model presented in Chapter 4 reflects an original interpretation of various neural processes and molecular quantities, which differs from the traditional view in the literature, and that is still waiting to be validated. For this reason, we decided to develop the two models independently, and to use an established single neuron dynamics to build the current population model. This enables the population model to stand autonomously from assumptions and results in single neuron, providing the credibility of the former (population) model not to be affected by the validity of the latter (single neuron) one. Moreover, using a simpler single neuron dynamics to build the population model increases discussion clarity and in-silico tractability of the system.

5.2 Learning the statistics of the world

In the following, we will investigate how a local synaptic plasticity rule can enable neuronal populations to encode information efficiently and, concurrently, to learn prior statistics of the world. In our model, we consider a cortical population (e.g. V1) receiving input from sensory neurons (e.g. via LGN), whose activity is conditional to the state of the external world. Information about the sensory features of the world is encoded in the cortical population according to 1. the response characteristics of the sensory neurons and 2. the dynamically evolving connectivity structure between sensory and cortical neurons (synaptic connectivity pattern and weights, which determine receptive field size and location). While cortical neurons response characteristics are well characterized experimentally, here we focus on the dynamics of the connectivity, which is harder to probe.

5.2.1 The model

We model the spiking activity of sensory neurons as independent Poisson processes, with mean firing rates determined by the state of the external world according to some given neurons tuning curves (specifically, for the tuning curves we used Gaussian functions uniformly distributed in the hidden variable space). Spike events in the sensory population drive the activity of cortical neurons through synapses. We described the dynamics of synapses on a short timescale using the variable P_n^i , which represents the instantaneous activation level of a the synapse connecting neuron i to neuron n . Each time there is a pre-synaptic spike, the activation level P_n^i increases of a fixed quantity ΔP while, in absence of spikes, P_n^i decreases exponentially to 0 with time constant τ_i .

$$\frac{d}{dt}P_n^i = -\frac{1}{\tau_i}P_n^i + \delta(t - t_{\text{pre-synaptic spike}})\Delta P \quad (5.1)$$

We model the activity of cortical neurons (driven by synapses) using a conductance-based Leaky Integrate-and-Fire model, with the addition of white noise to mimic membrane fluctuations. The state of the given neuron is represented by its membrane potential V_n , and it is driven by the product of synaptic activation P_n^i and of the connectivity strength W_n^i , which changes on a slower time-scale according to a plasticity rule (we will discuss the specific plasticity rule later). Furthermore, each time the membrane potential V_n reaches a threshold V_{th} , an action potential is generated, and V_n is reset to the equilibrium value E

$$\tau_m \frac{d}{dt} V_n = (E - V_n) + \sum W_n^i P_n^i (E_n^i - V_n) + \nu dW \quad ; \quad V_n > V_{th} \longrightarrow \begin{pmatrix} \text{Spike} \\ V_n = E \end{pmatrix} \quad (5.2)$$

where τ_m represents the timescale of membrane potential decay, E is the equilibrium potential of the neuron, E_n^i is the reversal potential for different synapses, ν is a parameter controlling the amplitude of random fluctuations and dW is white noise. The proposed system and its in-silico behavior is described in figure 5.1.

5.2.2 Synaptic plasticity

The connectivity W_i^n between sensory and cortical population is dynamic and plastic: we propose a local Hebbian synaptic plasticity rule that enables the cortical neurons to adapt their tuning curves to the sensory input statistics. Our aim is to maximize the mutual information between the cortical population and the state of the world under anatomical and metabolic constraints. We propose a simple synaptic learning rule that exploits the synaptic activity history (represented by parameter P_i^n) and the instantaneous post-synaptic activity to update on-line the connection weights W_i^n . The strength of the synapse connecting the pre-synaptic neuron i to the post-synaptic neuron n (described by the connection strength or weight w_i^n) can be interpreted as reflecting a combination of the number of releasable vesicles, release probability, active zone surface, number of active zones and responsiveness of the active zone to released neurotransmitters, while the activation parameter P_i^n can be

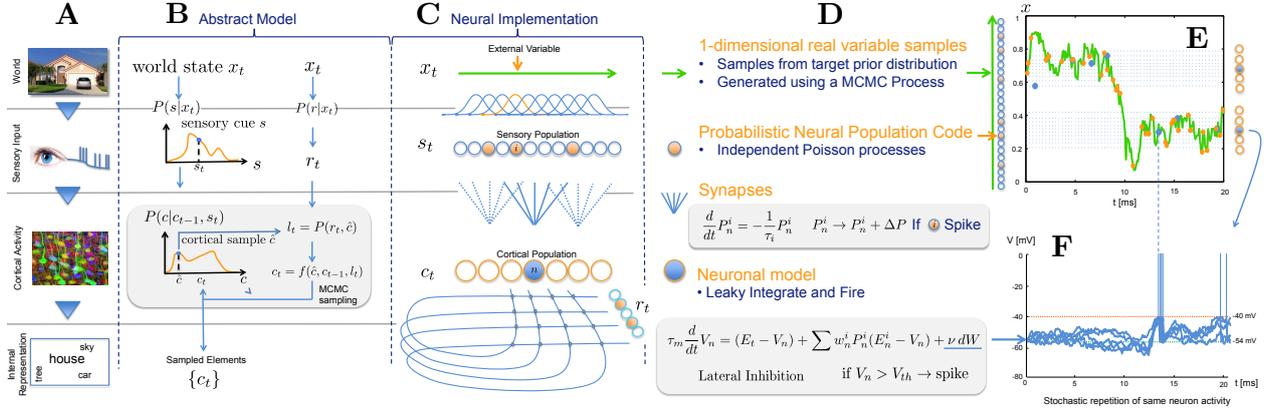


Figure 5.1: Model of neural population dynamics and functioning. A) Abstract representation of the proposed system: The sensory input, conditional to the state of the world, drives the activity of a cortical population, which encodes the state of the world. B) Probabilistic description of the model, which highlights the relations between different system's constituents. C) Neural implementation of the abstract model. D) Dynamics used to simulate neural populations in-silico. Successive values of the external variable are simulated using a Markov Chain Monte Carlo process, which generates correlated samples from a certain underlying statistics. The spiking activity of the sensory population is modeled using independent Poisson processes, whose instantaneous rate is conditional to the state of the world according to a provided receptive field. Synapses dynamics is simulated using an exponential decay of the synaptic activation variable, which increases of a fixed quantity after each pre-synaptic spike. Cortical neurons are modeled using a Leaky integrate-and-fire conductance-based model, with the addition of white noise and lateral inhibition. E) Simulated behavior of the system. The time-dependent value of the encoded world feature (green line) conditions the spiking activity of the sensory population (200 neurons, orange dots), which in turn drives spikes in the cortical population (20 neurons, blue dots). The position of every sensory/cortical neuron in the ordinate axis corresponds to the center of its tuning curve in the feature space. F) Stochastic leaky integrate-and-fire dynamics of cortical neurons: In blue, the membrane potential of a cortical neuron is plotted over repeated stochastic simulations driven by the same pre-synaptic input. The dotted-green line represents the leakage potential (-54 mV) while the dotted-orange line represents the threshold for action potential initiation (-40 mV).

interpreted as a short-time-average of the pre-synaptic firing activity (reflecting local 2^{nd} messengers, synaptic protein expression or upstream regulatory proteins). The proposed

plasticity rule can be mathematically formulated as follows:

$$\frac{d}{dt}W_i^n = \delta(t - t_{\text{post-synaptic spike}})(\alpha P_i^n - rW_i^n) \quad (5.3)$$

$$\frac{d}{dt}r = \rho^- + \delta(t - t_{\text{post-synaptic spike}})\rho^+ \quad (5.4)$$

where r is a regularization variable, and α , ρ^- and ρ^+ are model parameters. According to the proposed model, if a pre-synaptic spike precedes a post-synaptic event, P_i^n will be high when the post-synaptic neuron fires, and the synaptic strength w_i^n will increase. A visual representation of the described plasticity rule can be found in figure 5.2C.

5.2.3 Lateral Inhibition

In addition, to regulate population activity, we require a homeostatic mechanism to control the total average firing activity of the whole cortical population in different sensory input regimes. We implemented this regulatory mechanism by introducing lateral inhibition in the cortical population. To achieve the desired behavior we tested three alternative approaches, which produced consistent results: 1. We introduced a global memory parameter in the system (e.g. local 2nd messenger signals such as Nitric Oxide concentration (Bredt and Snyder, 1992; Artinian et al., 2012)) whose dynamics depends on the activity of the cortical population and which acts on the neurons by modifying the threshold for spike generation in the Integrate-and-Fire model. 2. We reset the membrane potential of all neurons to their resting potential value whenever a spike was generated within the cortical population. This winner-take-all behavior could be biophysically implemented through recurring lateral inhibitory connections. 3. We introduced in the population dynamics a common inhibitory component driven by the average firing rate of the population (e.g. Implemented through an auxiliary inhibitory population). This homeostatic mechanism plays two different roles within the dynamics of the population, by A. Introducing competition in the activity of different

neurons and B. enabling the system to produce spontaneous activity driven by molecular noise in the absence of sensory input.

5.2.4 Results

The synaptic plasticity rule we propose allows a neural population to adapt its encoding scheme to the sensory input statistics, in order to represent information efficiently. Specifically, synaptic adaptation provides the receptive field of the cortical neurons to converge to the optimal configuration (as defined in (Ganguli and Simoncelli, 2010)), which maximize mutual information between neural activity and the state of the world under metabolic constraint. In fact, it is possible to show (Ganguli and Simoncelli, 2010) that, given a population of spiking neurons using rate coding to represent an external variable, the tuning curves that maximize mutual information between neural activity and world state satisfy the following requirements: (a) the density of the neurons receptive fields in the world feature space is proportional to the prior probability of the encoded feature, (b) the width of each tuning curve is inversely proportional to the prior probability and (c) the gain of all the neurons is constant, meaning that the maximum allowed firing rate is the same for each neuron. These mathematical predictions have been confirmed by biophysical and behavioral experiments (Ganguli and Simoncelli, 2010). The local learning rule we propose allows modifying the synaptic connections between sensory and cortical neurons and, as a consequence, to shape the receptive fields of the cortical population. Specifically, the given plasticity rule provides the tuning curves of the cortical population to converge to the optimal configuration (Ganguli and Simoncelli, 2010) without the need of any global supervision, and to reproduce all the previously mentioned features (see figure 5.2E-F, R^2 consistently higher than 0.9). By doing this, the cortical population is also simultaneously learning the prior distribution of the encoded stimulus. In fact, since the population receptive fields are adapted to the statistics of the world features, information about this statistics is implicitly encoded in the synaptic

connections between cortical and sensory populations.

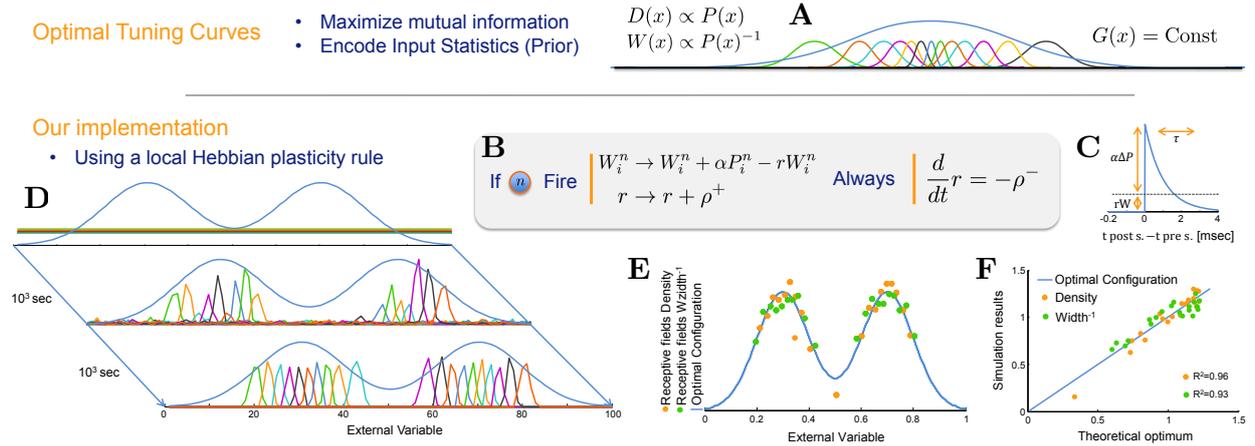


Figure 5.2: Learning the optimal tuning curves: (A) The optimal configuration of tuning curves is plotted for a given stimulus prior (blue line). (B) Mathematical and (C) visual representation of the plasticity rule (TDSP) used in the simulation, which provides the neural population to adapt its tuning curves to the experienced input statistics. W_i^n represents the weight of the synapse connecting pre-synaptic neuron i to post-synaptic neuron n . α is a learning rate parameter and r is a regulatory term. ρ^+ and ρ^- are constants. (D) The cortical population adapts its tuning curves to the input statistics. Different colored lines represent the tuning curves of 20 cortical neurons and their evolution in the feature space, as simulated by our model of neural activity and plasticity. In blue, the prior distribution of the input used in the simulations is plotted, which represents the statistics of the encoded stimulus in the world. The activity of 20 cortical neurons receiving input from 200 sensory neurons was simulated for few hundred minutes (simulation time) starting from a random connectivity matrix. Optimal encoding would predict the gain of the tuning curves to be constant, the tuning curves density to be directly proportional to the prior statistics and the tuning curves widths to be inversely proportional to the prior distribution. (E) The inverse of the distance between the centers of neighbor tuning curves (green dots) and the inverse of tuning curves widths (blue dots) are plotted as a function of the position of the tuning curves centers in the feature space, as produced by our simulations. Optimal encoding would predict that the points lie on the prior distribution (blue line). (F) The results of the simulations for the inverse distance between centers of neighbor tuning curves (green dots) and the inverse tuning curves widths (blue dots) are compared to the values predicted under the optimal encoding assumption. The comparisons provide R^2 values consistently approaching 1.

5.3 Sampling from the prior

The prior information about the input statistics (encoded in the synaptic connections driving the cortical population) needs to be extracted in order to be accessible and used in the inference process. For example, in the absence of relevant sensory input (e.g. eyes are closed) the brain should be able to represent the prior belief about the world. Specifically, for the purpose of this paper, in order to implement a Markov Chain Monte Carlo algorithm the first ingredient we need is a sequence of random samples from the prior distribution. This is a challenge for many other neuronal population models, as these cannot generate endogenous activity in the absence of sensory inputs. In a no-input spontaneous activity regime, our model neurons produce noise-driven spontaneous neural activity, in which neuronal spiking is randomly driven by cellular noise sources (Faisal et al., 2008).

5.3.1 Spontaneous activity

If we assume that the neuronal decoder is adapted to the neural encoder and that each single spike event is decoded independently, we can interpret the spontaneous firing patterns from the neuronal population as samples from the learned prior distribution (see figure 5.3A). In fact, one of the conditions necessary for an optimal encoder is that, when representing the stimulus statistics over a long period of time, each neuron spikes in average the same total number of spikes. Similarly, in the absence of sensory drive, all the cortical neurons are driven by internal noise equally. Thus, the resulting firing rate is constant across the population, and the spiking activity is comparable to what we would have if the population were encoding a sequence of values sampled from the original statistics. Additional structure in the stimulus dynamics (e.g. trends and correlations)(Berkes et al., 2011) could be learned and reproduced by the population exploiting recurrent excitatory and inhibitory connections, but this exceeds the scope of this paper and it will be examined in future works.

5.3.2 Interpreting spontaneous activity

To interpret the role of spontaneous activity in the general picture of the working brain, we need to make some assumptions about how this activity is used and decoded in successive levels of brain processing. Many studies investigate how a neural decoder for Probabilistic Population Codes could be implemented, and test optimality of the proposed decoding methods (Nirenberg and Latham, 2003; Averbeck et al., 2006). From now on, for simplicity, we will assume that the brain decodes every single spike event optimally or near optimally, without making any assumption about how the decoding process is implemented. Moreover, we will assume that the decoding of the spikes history during a certain time interval corresponds to the normalized superposition of the beliefs about the state of the world provided by every single event. Under these assumptions, if we decode the activity of the population in the spontaneous activity regime, the distribution inferred by the decoder about the state of the world will match the previously learned prior statistics. In this perspective, we can interpret each population spike event (and the correspondent distribution inferred by the decoder about the state of the world) as a sample from the prior distribution. We shall note that, following this approach, each sample represents a distribution (with smaller entropy than the original distribution we are sampling from) and not a point, as it is usually the case in the literature. In this sense, in order to go from single samples to the original distribution, we will use the superposition of functions instead of the traditional histogram.

5.4 Inference by sampling (Markov Chain Monte Carlo)

In the absence of sensory drive, there can still be additional sources of information that can be integrated to improve our cortical representation of the state of the world. For example, additional information might come from comparing low-level predictions derived from the current representation of the world state with actual experiences. This is the case

whenever we are provided with sensory input that can't be directly exploited for driving cortical activity, yet it can offer a measure of likeliness for the current cortical representation. Such a process is analogous to what happens when we use a generative approach to infer causes from observations: we propose various plausible causes, we derive their predictions and we test them with the observed facts looking for the best explanation. The model of Bayesian inference we propose allows exploiting these sources of information to generate samples from the posterior distribution through the implementation of a Markov Chain Monte Carlo algorithm (MCMC). As we mentioned before, in our model we will identify spike event with a generated sample. According to the proposed algorithm, after each sample has been generated the brain evaluates its likelihood given the additional sources of evidence, and it uses this likelihood to condition the generation of following samples.

In particular, we explored the implementation of a MCMC algorithm driven by noise, in which the feedback information regarding the likelihood of evaluated samples affects the generation of successive samples by tuning the level of noisiness of the cortical population.

5.4.1 Noise driven inference

As already mentioned, in absence of sensory drive the cortical population dynamics generates spontaneous activity. In this regime, recurrent lateral excitatory connections between neurons make neural activity conditional to the current state of the system. The probability of each neuron to spike at a given moment in time will therefore depend on the recent population spiking history. Specifically, recurrent excitatory connections between neurons will reflect similarity in their receptive fields, and a neuron will be more likely to spike if it is connected recently active neurons. This recurrent connectivity can be learnt exploiting temporal correlation in the driven neural activity, and provides spontaneous activity in the population to evolve similarly to a random walk, where the encoded variable slowly moves away from its current value. Furthermore to recurrent connections, molecular sources of noise stimulate all

the neurons of the cortical population equally, this superimposing a uniform spiking probability distribution over all the possible states to the one previously described. Each sample is therefore generated from the superposition of two different distributions over the available system states: 1) D_1 uniform over all the possible neurons and 2) D_2 , conditional to the last sample and localized around it. The level of noise in the population activity regulates how the two distributions are weighted when superimposed to generate the next sample. We assume that whenever a sample is generated, its likelihood is evaluated through a forward model that compares the sample predictions with additional sources of information (e.g. memory or sensory), and that the evaluated likelihood regulates the level of noise in the neural activity. A biophysical plausible implementation for this mechanism can rely on dopaminergic neurons, which are known to regulate the variability of neural activity and connect different brain regions. This provides that, when a state with high likelihood is sampled, the level of noise in the population is reduced and the next state is sampled from a distribution strongly localized around the current state. On the contrary, when a sample with low likelihood is generated, the level of noisiness is increased, providing next sample to be generated from a distribution with higher entropy. It can be shown that, under certain conditions, the described mechanism is able to generate samples from the posterior distribution. Specifically, under the limit of D_2 having very small entropy and of D_1 weight being very small and proportional to the inverse of the current sample likelihood, it is possible to show that the proposed algorithm converges to a special case of the Metropolis-Hasting Method. Additionally, if we relax these conditions, its behavior can be shown to approach optimality in some realistic working regimes. We tested this algorithm in-silico using spiking neurons and its behavior matched the analytical optimal inference predictions (Kullback-Leibler divergence $< 10^{-4}$). The main problem with the described algorithm is its behaviour in solving high dimensional problems. In fact, its random walk nature strongly affects the performance and time required for generating independent samples in high dimensional spaces, this restricting its effectiveness to simple problems (as for example bistable perception). For this reason, we also investigated other

alternative MCMC algorithms and we explored their possible respective implementations.

5.4.2 Alternative algorithms

We tested an alternative MCMC implementation, which is inspired directly by the Metropolis-Hasting method. In this algorithm, each sample is generated from a distribution conditional on the previous sample (this easily implementable using intra-population connections), and new samples can be rejected or accepting stochastically with probability conditional to their likelihood. We propose an implementation of the acceptance/rejection process through coincidence detection followed by a gating mechanism. The model provided good agreement with the optimal predictions, but its implementation requires ad-hoc hand-designed connection and is not particularly biophysically plausible.

The last algorithm we propose exploits memory and latent variables, identifiable with the hidden biophysical state of the neurons. The likelihood of each spike sample condition the state of the latent variable and, therefore, the probability of generating again the same sample later in the future (on a certain time-scale). This algorithm is much more efficient for dealing with complex variable spaces and complex joint distributions similar to the ones we often face in real world tasks, and its biophysical implementation is more reasonable than for the Metropolis-Hasting method.

5.4.3 Results

We implemented and tested all the proposed MCMC mechanisms using our model of spiking neurons, and the simulation results consistently matched analytical optimal inference predictions for all the different algorithms (see figure 5.3B), although requiring different time scales to converge, and to generate independent samples.

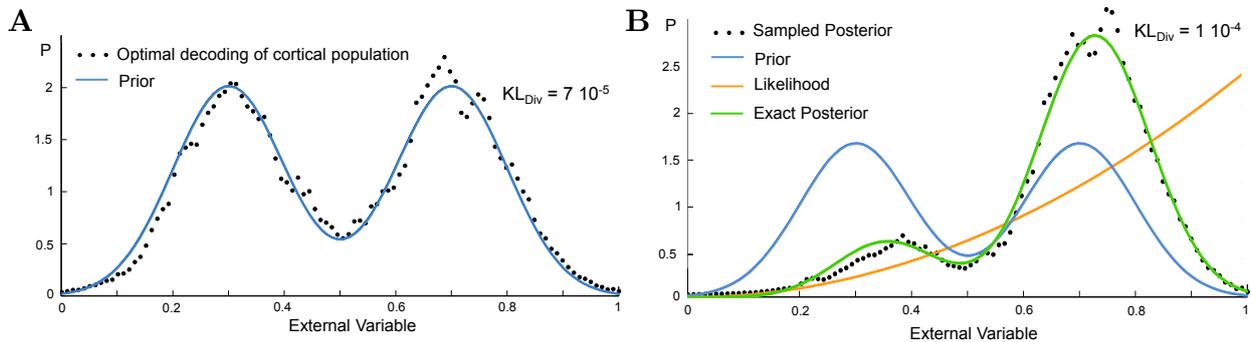


Figure 5.3: Inference by sampling. A) In the absence of sensory drive, the neural population generates spontaneous activity as a consequence of lateral inhibition and, in average, each neuron generates the same number of spike events. If decoded optimally (using a decoder adapted to the encoding scheme of the population in presence of sensory drive) each spike event can be interpreted as a sample from the prior distribution, which corresponds to the statistics of the stimulus in the world. The graph shows simulation results. The blue line represents the mathematical prior, while black dots represent the probability distribution extracted from the frequency of samples observed in-silico. Kullback-Leibler divergence $7 \cdot 10^{-5}$). B) A feedback system allows to evaluate the likelihood of each generated samples, and to use this information to control the generation of next samples, implementing a Markov Chain Monte Carlo (MCMC) algorithm. The graph represents simulation results using the MCMC algorithm based on noise discussed in the text. Blue line represents the prior distribution (encoded by spontaneous activity in absence of feedback) and the orange line represents the likelihood function (encoded by the feedback associated to each sample). The green line represents the mathematical posterior (calculated using the Bayes' equation (2.1)) while black dots represent the distribution extracted from the number of spike events generated in-silico by the neural population. The Kullback-Leibler divergence between the mathematical posterior distribution and the distribution inferred in-silico is $1 \cdot 10^{-4}$. Alternative MCMC mechanisms and their respective neural implementations (as discussed in the text) were also simulated in-silico, and they produced consistent results.

The mechanism for performing optimal inference using MCMC we propose here is expensive in terms of neural wiring, spikes and time required for generating independent samples, and more efficient solutions for implementing similar computations may be considered (Ma et al., 2006), which rely on a feed-forward processing of information in sequential cortical layers. However, as we already stated, in real world tasks brains often face problems that cant be solved with the straightforward processing of known information, requiring instead an

active role of the preceptor in generating possible solution for explaining the received input. This is, for example, the case when we use a complex forward model of the world to test our representation with incoming sensory information, but the map from incoming information to world representation can't be easily inverted. In this context, MCMC offers strong advantages compared to other proposed mechanism. In fact, using MCMC algorithms we only need to be able to evaluate the likelihood of the given samples, without knowing anything about the global distribution and its partition function. Moreover, as already mentioned, the proposed method requires a certain amount of time to generate reliable independent samples, free from random walk behaviors and detached from initial conditions. Interestingly, there is behavioral evidence that our brains face a similar time-accuracy trade off when dealing with certain perceptual or decision-making tasks (MacKay, 1982). Despite this is often attributed to the time needed to accumulate evidence and to integrate sensory information (Usher and McClelland, 2001), often the observed time scale is too big to justify this interpretation. Furthermore, the given trade-off has been shown to depend on variables internal to the subject (e.g. level of focus (Förster et al., 2003), suggesting that it might be caused by some internal process, and not only by the integration of received sensory evidence. In addition, random walk behaviors and dependence of the response on initial conditions has often been reported in behavioral experiments. The given observations transform the original point of skepticism, represented by performance limits and pathologic behaviors, into an additional source of support for the proposed model.

5.5 Discussion

We propose a model of neural population dynamics that, at the same time, allows to efficiently encode incoming sensory information, to generate samples from the prior distribution in the absence of sensory input and to exploit additional (feedback) sources of information to implement inference by sampling.

Specifically, we propose a local Hebbian synaptic plasticity rule that allows learning the connectivity structure and, therefore, the optimal tuning curves for the cortical neural population. This enables the system to efficiently encode sensory information and, at the same time, it provides to learn and internally encode the prior statistics of the stimulus. Our results match biophysical experimental results linking a population receptive field to the input statistics (Ganguli and Simoncelli, 2010), and our model provides quantitative predictions about how population tuning curves adapt to changes in the stimulus statistics.

Moreover, we propose a homeostatic mechanism that produces spontaneous activity in the absence of direct sensory information. Under some assumptions about how population spiking is decoded, spontaneous neural activity can be interpreted as a process generating samples from the learned prior distribution. Our interpretation is supported by physiological results (Berkes et al., 2011).

Finally, we propose three alternative MCMC algorithms for performing Bayesian inference, and we develop a biophysically plausible implementation for them in spiking neurons. The implemented algorithms consistently produced results that matched the analytical optimal inference predictions. The proposed mechanisms, which in the current work has been used to perform optimal inference, can also provide an explanation of how brains solve other real world tasks both in perception and decision making, as for example cue integration (Wozny et al., 2010) and action selection (Gurney et al., 2001).

The model of cortical computation we propose is able to autonomously switch between different operational regimes (sensory encoding or spontaneous activity) according to the received input drive. Furthermore, the proposed model exploits microscopic noise as a resource that enables the system to get in correspondence with the statistical properties of the environment. This places subcellular sources of noise under a new light in the working-brain picture. Noise, which has always been considered a limiting factor for neural processing and encoding, plays a crucial role in our model of neural computations, since it represents the generative

engine supporting active perception.

The model described in this Chapter allows implementing active perception via sampling. It is possible to show that a similar approach can be used to solve decision making and action selection tasks, where an active generation of solutions and a feedback evaluation of proposals is necessary to select the best behavior.

Additionally, various populations can be used to encode different features of the world. In this case, the correlation structure of the encoded features can be learned using recurrent connections between the populations. This allows exploiting all the available information to infer the state of the world and, for example, in the absence of auditory information, visual cues can be used to condition the spontaneous activity of the population encoding sound features (and vice-versa). By doing this, we can produce complex inference behaviors such as cue integration.

Finally, recurrent connections within the same population can be used to learn the time structure of world features, and the transition probabilities between different states. This allows generating a realistic simulation of the external world inside the brain. Each time new sensory information is available, this information is used to update the state of the simulation accordingly, in a Bayesian fashion. On the contrary, if no additional information is available, the internal simulation evolves in a realistic manner, respecting causality and probability rules learned by the system about the external world (also explaining results from (Berkes et al., 2011)). This simulation machine would allow to represent the external world, predict future events, evaluate hypothesis and select actions and behavior according to the expected reward.

The model developed in the current Chapter should not be considered as a description of brain functioning. However, it should serve as a feasibility study of how a population of neurons can exploit intrinsic neural variability to implement inference, and to simulate the stochastic dynamics of the external world. In this sense, the achieved results can be used

as a starting point to derive biophysically detailed model of the dynamics of specific brain areas, and a normative description of their computational role in the working brain.

Chapter 6

Conclusion

‘As long as our brain is a mystery, the universe, the reflection of the structure of the brain will also be a mystery.’

Santiago Ramn y Cajal, “Rules and advices on scientific investigation”, 1987

According to Marr’s tri-level hypothesis (Marr and Poggio, 1976), understanding an information processing system requires three distinct, complementary levels of analysis. First, we need to understand what is the system doing, and why is it doing what it does (computational level). Moreover, we need to comprehend how the task is pursued, i.e. what representations and processes are used to achieve the desired goals (algorithmic/representation level). Finally, we need to discern how these computations are physically realized, starting from the fundamental constituents underlying the used architecture (implementation/physical level).

Nervous systems are able to generate extremely rich and complex behaviors, starting from an overwhelming miscellany of microscopic constituents and interactions. For this reason, extensive Neuroscience research has been focusing either on deriving effective models of large-scale system behaviors, or on meticulously describing local mechanisms and interactions. As

these fields of research are growing, we reached a point that enables us to start combining these different levels of description, and to link them to the computational role underlying the observed dynamics. Recently, various works have been pursuing this goal, creating an elegant framework that provides a principled and simple explanation of rich and complex mechanisms (Deneve, 2008; Ganguli and Simoncelli, 2010; Neishabouri and Faisal, 2011; Berkes et al., 2011; Neishabouri and Faisal, 2014a). Inspired by the approaches used in these papers, we tried in this dissertation to create a top-down explanation of what is observed about the dynamics of neurons and neural population, linking specific mechanisms to their computational role in the brain.

6.1 Summary of thesis achievements

In this dissertation, we propose two different models: a model of single neuron functioning and a model of neural population dynamics. The two models require a different level of detail. For this reason, we decided to develop them independently, and we used a simple established single neuron dynamics to build the population model. This allows the proposed models to stand independently from reciprocal assumptions and results, and it improves discussion clarity and in-silico tractability of the population model.

6.1.1 Single neuron

The single neuron model we propose in this dissertation describes neural functioning on a computational, algorithmic and implementation level, and it shows how different levels of analysis are consistent and equivalent. Using a top-down approach, we derive from first Bayesian principles a normative description of neural behavior, and we show how the proposed computations can be implemented by a biophysical neuron.

On a computational level, the proposed neural dynamics allows to implement binary regression on some hidden states (dendritic integration), to propagate this information efficiently through the neuron (action potentials generation), to encode it locally (calcium intracellular concentration) and to use it to generate samples from the current internal representation of the external world (synaptic vesicle release).

On an algorithmic level, the proposed binary regression takes the form of logistic inference. The information is propagated through the neuron efficiently using a predictive coding scheme, which exploits local calcium concentration to implement divisive gain normalization. The inferred probability is encoded locally by calcium concentration in form of log-odds probability ratio. Finally, local calcium concentration controls the instantaneous probability of synaptic vesicles release.

On the implementation level, we show how the proposed algorithms can be performed by a biophysical neuron exploiting neurotransmitter-activated ions channels, linear dendritic integration, calcium-activated potassium channels, unstable membrane dynamics, random membrane fluctuations and potential-dependent calcium channels. All the biophysical requirements and predictions of our model are compatible with experimental observations.

The proposed model provides an original interpretation of biophysical quantities and processes, as intracellular calcium concentration, action potentials and synaptic vesicle release. Spiking activity, which is usually considered the paramount form of communication between neurons, in our model takes the role of an internal information channel, which serves to coordinate and synchronize local calcium concentration in different regions of the neuron. This information channel exploits an efficient coding scheme that is adapted to the dynamical properties of the encoded stimulus, and the potential dependent dynamics of calcium channels allows to locally decode the message contained in it. Following this new interpretation, intracellular calcium concentration becomes the variable encoding the state of the world inside the neuron, and it controls synaptic communication, neural interactions, spike rate

adaptation and synaptic plasticity. Finally, according to the proposed interpretation, synaptic vesicle release doesn't serve to translate pre-synaptic spikes into excitatory post-synaptic potentials (EPSPs). On the contrary, it actively implements a process that generates samples from the probability encoded in synaptic calcium concentration.

The model we propose link the parameters controlling neural dynamics to the structure (e.g. timescale of changes) of the encoded stimulus. We quantitatively predict the given relations, and we show how the optimal parameters can be learned and tuned on-line to the experienced sensory input. This homeostatic mechanism prevents the need for an externally-supervised fine parameters tuning, and it enables the system to adapt to changes in the dynamical properties of the world.

We use a top-down approach, and starting from some assumptions about the implemented computations and the role of biophysical quantities inside the neuron we derive various observable predictions that match experimental results. For example, the model predicts short-term synaptic potentiation as a consequence of non-linear dependence of vesicle release probability on local calcium concentration, which is compatible with experimental observations about synaptic dynamics. Additionally, the proposed model generates quantitative predictions about the time-course of spike rate adaptation following a jump in the input drive. We experimentally validated the given prediction using in-vivo recordings from cat V1 neurons (Benucci et al., 2013).

6.1.2 Neural population

The model of neural population dynamics we propose should be interpreted as a feasibility study. A simple Hebbian synaptic plasticity rule and a lateral inhibition mechanism suffice to enable a population of spiking neurons to adapt its tuning curves to the statistic of the encoded stimulus, and to represent information optimally given some metabolic and

biophysical constraints. Moreover, in the absence of sensory information, the given population is able to generate an internal representation of the external world (spontaneous activity) that is compatible with the learned statistics, and which can be interpreted as a process of generating samples from the prior distribution (active perception). Finally, the given population can exploit additional sources of information to implement inference by sampling and, specifically, a Markov Chain Monte Carlo (MCMC) algorithm which generates samples from the posterior distribution. This allows to implement optimal Bayesian computations, and to combine different sources of information efficiently in order to infer the state of the world.

The model of cortical computation we propose can autonomously switch between different operational regimes (sensory encoding or spontaneous activity) according to the received input drive. Furthermore, the system functions and evolves on two different timescales: on a short timescale, the neural population simulates the state of the world given the available information while, on a longer timescale, it learns the statistics of the world, and it adapts its internal parameters accordingly.

As already mentioned, the given model should only be considered as a feasibility study, and it doesn't represent a detailed model of neural functioning. For this reason, we propose three different candidate mechanisms for realizing lateral inhibition within the neural population, and we test the implementation of three alternative MCMC methods, achieving consistent results. The given model can be used in the future to derive more detailed and realistic descriptions of specific brain regions and, in this case, the most appropriate MCMC method and lateral inhibition mechanism will be selected.

Our results match biophysical experimental results linking a population receptive field to the input statistics (Ganguli and Simoncelli, 2010), and our model provides quantitative predictions about how population tuning curves adapt to changes in the stimulus statistics. In addition, the interpretation of spontaneous neural activity we propose is compatible with

experimental observations from (Berkes et al., 2011).

6.2 Common themes shared by both models

In this dissertation, we propose two separate models: a model of single neuron functioning and a model of neural population dynamics and computations. The two models are developed independently. However, there are some common themes that link them.

6.2.1 Noise

Both models exploit noise as a resource, and they rely on molecular sources of variability to function properly. In the case of single neuron dynamics, molecular noise manifests itself both at the level of vesicle release and of action potentials generation. The presence of stochasticity in vesicle release allows synaptic communication to take the form of a sampling process. Each vesicle release is considered as an independent sample, conditional to the probability encoded by the neuron. Different synapses belonging to the same pre-synaptic neuron implement independent sampling processes from the same underlying probability. This sampling process plays a role that can also be interpreted as dithering. Neurons represent the probability of a hidden state, which is a continuous variable between 0 and 1. Yet, the nature of synaptic communication (vesicle release) is quantized. In each moment in time, a vesicle is either released or not. Using a maximum-likelihood approach would prescribe to always release a vesicle for probabilities higher than 0.5. Introducing stochasticity in the process does something similar to what dithering does in video/audio processing. Specifically, the presence of variability reduces quantization error, it prevents large-scale patterns and it improves the overall quality of the internal representation. At the level of spike initiation, stochastic membrane fluctuations (modeled as a Gaussian process) together with the unstable

membrane potential dynamics of the neuron allow producing the desired response function. In the case of the neural population model, membrane potential fluctuations drive and control spontaneous population activity in the absence of sensory drive. The given fluctuations enable the population to implement inference by sampling through a Markov Chain Monte Carlo algorithm, and they play a role similar to what “random variables” do in the implementation of Monte Carlo methods in-silico.

6.2.2 Learning

Another common theme to both models is learning. In fact, parameters fitting is a sensitive argument when developing mathematical descriptions of biological systems. Often, the behavior of proposed models heavily relies on a fine-tuning of some model parameters. This fine tuning must be robust to disturbances, and optimal parameters choice might depend on external conditions and on the system environment. This is, for example, the case of the Hodgkin-Huxley model of action potential propagation (Hodgkin and Huxley, 1952), in which spike propagation is highly sensitive to the relative number of ions channels in the neural membrane. Experiments showed that feedback systems enable a neuron to adapt the number of ions channels according to the properties of the solution surrounding the cell, in order to produce the desired behavior (Turrigiano et al., 1994). This homeostatic mechanism allows the neuron to function reliably, and to compensate for changes in the surroundings. Similarly, each biological model that requires fine tuning of environment-dependent parameters to function should always be accompanied by a model that provides these parameters to be tuned and maintained. For this reason, in both our models we include mechanisms that enable the system to update and learn model parameters locally and on-line, and to adapt these parameters to the properties of the experienced stimuli. On a small timescale, the proposed systems receive, encode, process and communicate information, while on a larger timescale they learn properties of the encoded stimulus from the received input, and they adapt their

internal parameters accordingly. In the single neuron case these parameters control internal neural dynamics, while in the neural population model they represent the weights of synaptic connections, which allow to optimally encode the external world and to successfully generate samples from the prior distribution.

6.2.3 Sampling

Both models implement a process that can be interpreted as “sampling from an underlying distribution”. This sampling behavior enables the system to “simulate” configurations of the external world that are compatible with the received sensory information. In the presence of limited knowledge about the current world state, instead of representing the maximum likelihood configuration, the system decides to stochastically select different candidate states following a probability matching behavior. Working in high-dimensional spaces and, often, on skewed non-Gaussian distributions, this behavior allows to avoid systematic errors, and to efficiently calculate expected values of cost functions over these spaces. The resulting variability of behavior also increases the fitness of the system by reducing predictability, supporting exploration and increasing diversification within and across organisms of a certain species, which increases the robustness of evolution to fast changes in the environment.

6.3 Applications

Technological progress is providing more accurate and powerful tools for probing and recording neural dynamics, and this is overwhelming modern Neuroscience with a huge amount of data to be processed and interpreted. For this reason, the research of simple computational principles behind complex observed neural dynamics and interactions can help neuroscientists build models of brain functioning without getting lost in the details of the system. Moreover,

understanding the computational role behind observed neural dynamics is fundamental to successfully interpret experimental results, to design adequate experiments for testing brain theories, and to understand brain's malfunctions and pathologies.

6.3.1 Intra-cellular calcium concentration as a gateway to brain functioning

For example, our model of neural dynamics suggests that intracellular calcium concentration represents the most explicit representation of the external world inside the neuron and, specifically, that calcium is linearly proportional to the log-odds probability ratio of some hidden state, which describes the current world configuration. If this is verified, it will imply that, to understand neural computations and representation, we should look at local calcium concentration in different regions of the brain, and relate these quantities to the stimuli used in the experiments. Furthermore, using spiking activity of neurons instead of calcium concentration would lead to confused and biased results. In fact, the coding scheme used by neural spiking activity varies in different neurons, and it depends on the dynamical properties of the encoded stimulus. For this reason, the spiking trains of neurons often seem to represent unpredicted changes in the encoded stimulus instead of the stimulus itself (Hughes and Maffei, 1966), and spiking response to a certain stimulus strongly depends on the history of the stimulus (spike rate adaptation)(Benda and Herz, 2003). Paradoxically, current experimental procedures often require using complex techniques to infer neural spiking activity from calcium concentration, which can be measured experimentally using calcium-imaging (Mittmann et al., 2011). If the results proposed in this dissertation are validated then, in order to understand what the neuron is saying to the rest of the system, it would be more efficient to just focus on the measured concentration of calcium ions.

6.3.2 Synaptic plasticity and computations

Additionally, understanding what is the message communicated by synapses (and how this message is used by the post-synaptic neuron) is fundamental to understand computations in neural networks. Our research suggests that vesicle release can be modeled as an inhomogeneous Poisson process, whose instantaneous rate is proportional to the probability of a certain event in the world. If validated, this can be used to derive from first principles synaptic plasticity rules that enable the neuron to implement some desired computations, and to associate these rules with the algorithms performed in different brain areas and by specific neuron types.

6.3.3 From neural parameters to the structure of the world

The model we propose creates a quantitative link between parameters controlling neural dynamics and the structure of the stimulus encoded. This bidirectional prediction allows extracting information about the natural statistics and dynamical properties of world features from the properties of neurons encoding them. Moreover, this quantitative link could potentially explain experimental results and artifacts achieved using artificial stimuli whose properties differ from the natural ones. Finally, being able to infer the natural properties of stimuli from the recorded neurons would support the design of more realistic experiments, that allow testing animals in a setting similar to the one they are adapted to.

6.3.4 A detailed, biophysical model of specific brain areas

On a population level, the model of perception and inference proposed in this dissertation can be used as a prototypical example to derive biophysically detailed models of neural dynamics in specific brain areas. As previously mentioned, in the current population model

we propose three candidate mechanisms implementing lateral inhibition as well as three alternative MCMC methods, which produce consistent results. In the future, we can consider specific brain regions to model, and we can select the most appropriate MCMC method and lateral inhibition mechanism accordingly. Furthermore, we can adapt the given model to fit properties of the brain area investigated, as for example average firing rate of neurons, timescale of synaptic filtering, size of the population, synaptic long-term plasticity, dynamics of the encoded stimulus and properties of the population in the spontaneous activity regime.

6.3.5 Spontaneous activity

In addition, understanding the functional role of spontaneous activity in the brain can support the development of new analysis techniques, which allow extracting valuable information from experimental data. Some brain regions have been shown to produce a higher level of activity in the spontaneous activity regime compared to when the animal is subject to sensory stimulation (Berkes et al., 2011). According to our interpretation of spontaneous activity, population activity in the absence of sensory drive carries a large amount of information about the features encoded. This can give us hints about the hierarchic level of representation in different brain regions, and how world causal/generative structure is represented inside the brain.

6.3.6 Brain inspired architectures and prosthesis

As an application to technology, understanding how neurons and neural populations are able to communicate information efficiently, to deal with noisy molecular components and to implement Bayesian computations could help us design brain-inspired architectures. In a similar direction, understanding what is the message encoded by neurons and how this message is communicated can help us building improved brain-machine interfaces and prosthesis.

6.4 Future work

In the future the proposed models can be improved, and the biophysical processes allowing to implement the desired algorithms can be further investigated. Moreover, the given models can be applied to solve novel tasks, and they can be expanded to account for a large number of neural populations. Finally, the models of single neuron and neural population can be combined in a unified model, which explains neural computation on a population level starting from the molecular dynamics of ions channels and neurotransmitters vesicles.

6.4.1 Single neuron

The single neuron model discussed in this thesis links abstract computations with the biophysical mechanisms that allow to implement them. In this direction, further work can be done to develop a more detailed description of molecular interactions that generate the desired behavior. For example, we should look for biophysical correlates of some missing model variables, e.g. the parameters controlling calcium dynamics and the variables controlling post-synaptic neurotransmitter-dependent ions channels. Our model proposes some quantitative predictions about how these parameters control neural functioning, how they depend on the dynamical properties of the encoded stimulus and how they evolve according to the experienced input drive and spiking history. Guided by these quantitative predictions, we should look in the biology of neural interactions looking for candidate quantities that match our requirements, and we should design some experiments to test our hypothesis on biophysical neurons.

Furthermore, the model we propose provides various quantitative predictions about neural dynamics, starting from synaptic vesicle-release rate up to the open-close transition probabilities of specific ions channels. In the future, the given model can be used to derive additional predictions about neural functioning. Also, experiments can be designed to test

these predictions, and to probe the validity of the underlying assumptions.

The model we propose can also be extended to describe other neural processes and phenomena. For example, our model suggests that dendritic integration implements logistic regression over a certain hidden state. Additionally, it suggests that the probability of this hidden state is encoded in the post-synaptic neuron in form of calcium concentration, and that the given inference process is based on the information encoded in synaptic vesicle release. These ingredients are sufficient to derive what the optimal synaptic weights should be, and how these weights can be learned on-line exploiting pre and post-synaptic activity. Therefore, we can derive from first principles a model of synaptic plasticity, which is controlled by vesicle release and by post-synaptic calcium concentration, and which reflects into an effective spike-time dependent plasticity (STDP) rule.

6.4.2 Neural population

The model of neural population dynamics we propose in this dissertation allows generating samples from a target posterior distribution, which is inferred using all the available sensory information about the state of the world. This sampling process can be used to calculate expected values of arbitrary functions over high dimensional spaces. Therefore, the given model can be extended to solve problems that require evaluating alternative options and choices, as for example decision making, motor control or action selections.

Moreover, the proposed model is able to learn the overall statistics of the encoded stimulus. However, stimuli in nature often present particular structures and evolution. For example, the position of a body free moving in space exhibits stereotypical trajectories, and knowing the current position of the body gives us information about what its future location will be. We can assume that these relations are learned by the population encoding the given stimulus, thanks to recurrent excitatory and inhibitory. On a first order, this will translate

into transition probabilities between neighbour states. If the body is in a particular location now, it is reasonable to expect that it will be close to that position in the next future, where the quantitative definition of “close” depends on the average velocity of the given object in nature. On a higher order, we can expect the population to learn non-linear transition probabilities that depend on the current and past state of the system. In the future, we can derive a model of how these transition probabilities can be learned and implemented exploiting synaptic plasticity of recurrent synaptic connections within the given neural population. In the presence of limited sensory information, this mechanism would enable the system to generate a representation of the world that reflects the correct statistics of the stimulus and, at the same time, it evolves compatibly to the stochastic dynamics of the feature in the world.

Finally, in this thesis we describe the model of a single population encoding one feature of the world. Yet, the world is made of a very large collection of features, and many of these features are statistically correlated in nature. We can consider more than one population of neurons encoding an equal amount of world features. The correlation between world features will reflect into correlation in the activity of the given populations, and this correlation can be learned by the system using synaptic connections between neurons of different populations. Specifically, in the future we can design from first principles a Hebbian synaptic plasticity rule that learns correlations between neural activities, and that replicates these correlations in the spontaneous activity of the system. The given mechanism would strongly increase the inference capabilities of the system. For example we can consider two populations, one encoding a visual stimulus (light generated from a certain location) and one encoding an audio stimulus (sound received from a particular position in space). We can assume that the two stimuli are statistically correlated, e.g. they are usually generated by the same physical object. This provides the activity of the respective populations to be correlated and, if we use the proposed mechanism to learn and reproduce the given correlation, we achieve some interesting behaviors. For example, if the subject closes his eyes (absence of visual information) but he carefully listens to the object sound (audio sensory drive), then

the visual population will generate spontaneous activity that is conditional to the activity of the auditory population representing the sound stimulus. As a consequence, the system will produce cue integration (Tassinari et al., 2006; Wozny et al., 2010) and it will exploit all the available information to infer the state of the world. The desired spike-time dependent plasticity rule needs to be able to distinguish between genuine correlation of world features, and correlation between the dynamics of neurons caused by existing recurrent connections (direct or indirect).

6.4.3 Combining single neuron and neural population models

If the model of single neuron functioning receives further validation from experimental testing, it can be considered to integrate it with the model of neural population dynamics. In order to make the two models compatible, it will be necessary to introduce some changes in the current model of neural population dynamics.

Specifically, the neural population model currently identifies spike events with samples. However, in the single neuron model the process of sampling is physically implemented by vesicle release. This change of perspective has one main consequence. Since each neuron communicate to other neurons through independent synapses, each post-synaptic neuron will observe a different realization of the same stochastic process (sampling). This increases stability of the process, since the overall dynamics of the system will not be influenced by one specific realization of the process as much as it would in the previous scenario.

Another significant difference between the two models is learning. In the current population model, learning is controlled by pre and post-synaptic spike timing. Yet, to make synaptic adaptation compatible with the single neuron model, we should create a new plasticity rule controlled by synaptic vesicle release (taking the place of pre-synaptic spiking in the current model) and by post-synaptic calcium concentration (which would substitute post-synaptic

activity). Using the given ingredients, it is possible to build a plasticity rule that generates the desired learning, and that is effectively equivalent to the Hebbian learning rule currently used.

Finally, in the combined model, lateral inhibition will not be controlled by total population spiking activity. Conversely, the inhibitory mechanism will depend on synaptic activity. One or more synapses for each neuron will drive an external inhibitory system, that acts back on the neurons of the population influencing their calcium concentration. This can be computationally interpreted as renormalizing a probability distribution. If we know that each neuron in the population represents a different event, and that all the given events are mutually exclusive in the real world, then we can infer the probability of each event independently from the available information (assuming observations relative to different events are mutually independent), and we can normalize all the derived probabilities to sum up to unity. This can be implemented through a dedicated population of neurons (inhibitory population), which infers the sum of the probabilities of all the neurons from their synaptic activity, and applies the inferred value to the probability represented by different neurons as a divisive normalization factor.

If successfully validated, the resulting (combined) system would integrate the properties of both the single neuron and neural population model we propose in this dissertation. It would adapt to the properties of the encoded stimuli, and it would encode and propagate information efficiently. It would be able to learn the causal structure of the world, to infer and represent its current state and to simulate its evolution. This is implemented exploiting noise as a resource, which allows to get in correspondence with the stochastic nature of the environment. The given model would link molecular quantities and processes inside the neuron to their computational role in brain functioning, integrating different levels of description. Such a model would allow investigating both the questions of “how” do neural populations function and, more importantly, “why” do they function as they do. This can lead in the future towards a universal neural theory (in the physics sense of the world) that,

building on first principles, explains one of the most complex systems in nature, which is the working brain.

‘If the human brain were so simple that we could understand it, we would be so simple that we couldnt.’

George E. Pugh, “The Biological Origin of Human Values”, 1977

List of symbols

$[c]$	Intracellular calcium ions concentration, relative to reference value $[ca^{2+}]_{ref}$
μ_0	Mean synaptic input in the hidden variable is in state 0
μ_1	Mean synaptic input in the hidden variable is in state 1
μ_i	Mean of the neural drive (firing rate in absence of adaptation)
μ_o	Desired average neural firing rate
ν	Instantaneous firing rate of the non-adapted neuron
ν_*	Instantaneous firing rate of the adapted neuron
ω	Observed synaptic current, used to infer the probability P of the hidden state
Φ	Log-odds probability ratio encoded by the membrane potential (equivalent to L^*)
Σ	Covariance matrix of the synaptic input
σ_e	Steady state solution of the posterior variance
σ_i	Variance of the neuronal drive (firing rate in absence of adaptation)
τ	Inter spike interval
τ_c	Correlation timescale of synaptic current

τ_o	Desired average neural inter-spike interval
τ_r	Refractory period of the neuron
ξ	Log-odds probability ratio inferred using synaptic current (equivalent to L_ω)
c	Mean of the inferred posterior distribution of the instantaneous neural input drive
C_m	Membrane capacitance of the neuron
D	Posterior distribution
g	Gamma distribution
I	Neural drive (firing rate in absence of adaptation)
I_c	Calcium-dependent potassium current
I_l	Leakage current
I_ω	Excitatory Post-Synaptic Current (EPSC)
L	Log-odds probability ratio of the hidden event
L^*	Log-odds probability ratio encoded by the membrane potential
L_c	Log-odds probability ratio encoded by calcium concentration [c]
L_ω	Log-odds probability ratio inferred using synaptic current
P	Probability of the hidden event
P^*	Probability of generating a spike event in the time interval τ_r
P_n^i	Activation level of the synapse connecting neuron i to neuron n
P_c	Probability encoded by calcium concentration [c]

P_{open}	Probability of a potassium ions channel being open
r_i	Timescale on which the external stimulus changes
r_{\downarrow}	Hidden state transition rate ($1 \rightarrow 0$)
r_{\uparrow}	Hidden state transition rate ($0 \rightarrow 1$)
v	Latent variable encoded by the neuron. Corresponds to L
V_{Φ}	Membrane potential
V_{th}	Threshold potential for spike initiation
W_n^i	Weight of the synapse connecting neuron i to neuron n
Z	Partition function of a potassium ions channel

Bibliography

- Abbeel, P. and Ng, A. Y. . Exploration and apprenticeship learning in reinforcement learning. In *Proceedings of the 22nd international conference on Machine learning*, 2005. (Cited on pages 4 and 26.)
- Abbott, L. F. and Dayan, P. . The effect of correlated variability on the accuracy of a population code. *Neural computation*, 1999. (Cited on page 4.)
- Abbott, L. F. , Varela, J. A. , Sen, K. , and Nelson, S. B. . Synaptic depression and cortical gain control. *Science*, 1997. (Cited on page 90.)
- Adelman, T. L. , Bialek, W. , and Olberg, R. M. . The information content of receptive fields. *Neuron*, 2003. (Cited on pages 26 and 28.)
- Aho, A. C. , Donner, K. , Hyden, C. , Larsen, L. O. , and Reuter, T. . Low retinal noise in animals with low body temperature allows high visual sensitivity. *Nature*, 1988. (Cited on page 4.)
- Amano, K. , Wandell, B. A. , and Dumoulin, S. O. . Visual Field Maps, Population Receptive Field Sizes, and Visual Field Coverage in the Human MT+ Complex. *Journal of Neurophysiology*, 102 (5):2704–2718, Nov. 2009. (Cited on page 71.)
- Arabzadeh, E. and Panzeri, S. . Deciphering the spike train of a sensory neuron: counts and temporal patterns in the rat whisker pathway. *The Journal of Neuroscience*, 2006. (Cited on page 64.)

- Artinian, L. , Zhong, L. , Yang, H. , and Rehder, V. . Nitric oxide as intracellular modulator: internal production of NO increases neuronal excitability via modulation of several ionic conductances. *European Journal of Neuroscience*, 2012. (Cited on page 145.)
- Augustine, G. J. , Santamaria, F. , and Tanaka, K. . Local Calcium Signaling in Neurons. *Neuron*, 40(2):331–346, Oct. 2003. (Cited on pages 65, 71, and 129.)
- Averbeck, B. B. , Latham, P. E. , and Pouget, A. . Neural correlations, population coding and computation. *Nature Reviews Neuroscience*, 2006. (Cited on pages 64 and 149.)
- Awatramani, G. B. , Boyd, J. D. , and Delaney, K. R. . Effective release rates at single rat Schaffer collateral–CA1 synapses during sustained thetaburst activity revealed by optical imaging. *The Journal of physiology*, 2007. (Cited on page 90.)
- Ayata, C. , Shimizu-Sasamata, M. , Lo, E. H. , and Noebels, J. L. . Impaired neurotransmitter release and elevated threshold for cortical spreading depression in mice with mutations in the $\alpha 1A$ subunit of P/Q type calcium channels. *Neuroscience*, 1999. (Cited on page 66.)
- Azevedo-Filho, A. and Shachter, R. D. . Laplace’s method approximations for probabilistic inference in belief networks with continuous variables. *Proceedings of the Tenth international conference on Uncertainty in artificial intelligence*, 1994. (Cited on page 14.)
- Baccus, S. A. and Meister, M. . Fast and Slow Contrast Adaptation in Retinal Circuitry. *Neuron*, 36(5):909–919, Dec. 2002. (Cited on page 28.)
- Bak, P. , Tang, C. , and Wiesenfeld, K. . Self-organized criticality: An Explanation of 1/f noise. *Physical Review Letters*, 59:381–384, 1987. (Cited on page 16.)
- Barlow, H. B. . Possible principles underlying the transformation of sensory messages. *Sensory communication*, 1961. (Cited on pages 26 and 28.)
- Benda, J. and Herz, A. V. . A universal model for spike-frequency adaptation. *Neural computation*, 15(11):2523–2564, 2003. (Cited on page 166.)

- Benucci, A. , Saleem, A. B. , and Carandini, M. . Adaptation maintains population homeostasis in primary visual cortex. *Nature Neuroscience*, 2013. (Cited on pages xvi, xvii, 28, 56, 57, 62, 74, 132, and 161.)
- Berkes, P. , Orbán, G. , Lengyel, M. , and Fiser, J. . Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science*, 331(6013), 2011. (Cited on pages 16, 83, 139, 141, 148, 155, 156, 159, 163, and 168.)
- Berry, D. A. . Bayesian clinical trials. *Nature reviews Drug discovery*, 2006. (Cited on page 15.)
- Berry, M. J. and Warland, D. K. . The structure and precision of retinal spike trains. In *Proceedings of the National Academy of Sciences*, 1997. (Cited on pages 71 and 124.)
- Bertram, R. , Sherman, A. , and Stanley, E. F. . Single-domain/bound calcium hypothesis of transmitter release and facilitation. *Journal of Neurophysiology*, 75(5):1919–1931, May 1996. (Cited on page 92.)
- Bhumbra, G. S. and Dyball, R. . Spike coding from the perspective of a neurone. *Cognitive Processing*, 2005. (Cited on page 106.)
- Bialek, W. . Physical limits to sensation and perception. *Annual Review of Biophysics and Biophysical Chemistry*, 1987. (Cited on pages 4 and 24.)
- Bialek, W. . Perspectives on theory at the interface of physics and biology. Dec. 2015. (Cited on page 2.)
- Bialek, W. and Setayeshgar, S. . Physical limits to biochemical signaling. *Proceedings of the National Academy of Sciences*, 2005. (Cited on pages 4 and 24.)
- Biernaskie, J. M. , Walker, S. C. , and Gegear, R. J. . Bumblebees learn to forage like Bayesians. *The American Naturalist*, 2009. (Cited on page 15.)
- Boerlin, M. , Machens, C. K. , and deneve, S. . Predictive Coding of Dynamical Variables in Balanced Spiking Networks. *PLoS Computational Biology*, 2013. (Cited on pages 29 and 33.)

- Bogacz, R. . Optimal decision-making theories: linking neurobiology with behaviour. *Trends In Cognitive Sciences*, 2007. (Cited on pages 3 and 15.)
- Borst, J. and Sakmann, B. . Calcium influx and transmitter release in a fast CNS synapse. *Nature*, 1996. (Cited on page 71.)
- Bouman, C. and Shapiro, M. . A multiscale random field model for Bayesian image segmentation. *Image Processing*, 1994. (Cited on page 115.)
- Branco, T. and Staras, K. . The probability of neurotransmitter release: variability and feedback control at single synapses. *Nature Reviews Neuroscience*, 10(5):373–383, May 2009. (Cited on page 80.)
- Bredt, D. S. and Snyder, S. H. . Nitric oxide, a novel neuronal messenger. *Neuron*, 1992. (Cited on page 145.)
- Brette, R. . Adaptive Exponential Integrate-and-Fire Model as an Effective Description of Neuronal Activity. *Journal of Neurophysiology*, 94(5):3637–3642, Nov. 2005. (Cited on page 106.)
- Burkitt, A. N. . A Review of the Integrate-and-fire Neuron Model: I. Homogeneous Synaptic Input. *Biological Cybernetics*, 95(1):1–19, Apr. 2006. (Cited on pages 19, 50, 67, and 106.)
- Burkitt, A. N. and Clark, G. M. . Analysis of integrate-and-fire neurons: synchronization of synaptic input and spike output. *Neural computation*, 1999. (Cited on page 20.)
- Caporale, N. and Dan, Y. . Spike Timing–Dependent Plasticity: A Hebbian Learning Rule. *Annu Rev Neurosci*, 31(1):25–46, July 2008. (Cited on pages 23 and 133.)
- Carlin, B. P. and Louis, T. A. . Bayes and empirical Bayes methods for data analysis. *Statistics and Computing*, 1997. (Cited on page 15.)
- Cash, S. and Yuste, R. . Linear summation of excitatory inputs by CA1 pyramidal neurons. *Neuron*, 1999. (Cited on pages 79 and 103.)
- Castellani, G. C. . A model of bidirectional synaptic plasticity: From signaling network to channel conductance. *Learning & Memory*, 12(4):423–432, July 2005. (Cited on page 133.)

- Castellani, G. C. , Quinlan, E. M. , Cooper, L. N. , and Shouval, H. Z. . A biophysical model of bidirectional synaptic plasticity: Dependence on AMPA and NMDA receptors. *Proceedings of the National Academy of Sciences*, 98(22):12772–12777, Oct. 2001. (Cited on page 133.)
- Catterall, W. A. , Leal, K. , and Nanou, E. . Calcium channels and short-term synaptic plasticity. *Journal of Biological Chemistry*, 2013. (Cited on page 87.)
- Chander, D. and Chichilnisky, E. J. . Adaptation to Temporal Contrast in Primate and Salamander Retina. *The Journal of Neuroscience*, 2001. (Cited on page 28.)
- Chen, N. , Chen, S. , Wu, Y. , and Wang, J. . The refractory periods and threshold potentials of sequential spikes measured by whole-cell recording. *Biochemical and biophysical research communications*, 2006. (Cited on page 115.)
- Cohen-Cory, S. . The developing synapse: construction and modulation of synaptic structures and circuits. *Science*, 2002. (Cited on page 64.)
- Connors, B. W. and Long, M. A. . Electrical synapses in the mammalian brain. *Annu Rev Neurosci*, 2004. (Cited on page 21.)
- Cox, M. A. , Schmid, M. C. , Peters, A. J. , Saunders, R. C. , Leopold, D. A. , and Maier, A. . Receptive field focus of visual area V4 neurons determines responses to illusory surfaces. *Proceedings of the National Academy of Sciences*, 110(42):17095–17100, Oct. 2013. (Cited on page 70.)
- Cramer, J. S. . The origins and development of the logit model. *Logit models from economics and other fields*, 2003. (Cited on page 74.)
- David, S. V. , Fritz, J. B. , and Shamma, S. A. . Task reward structure shapes rapid receptive field plasticity in auditory cortex. In *Proceedings of the National Academy of Sciences*, 2012. (Cited on page 70.)
- Arcangelis, L. , de, Perrone-Capano, C. , and Herrmann, H. J. . Self-organized criticality model for brain plasticity. *Physical Review Letters*, 2006. (Cited on page 16.)

- De La Rocha, J. , Doiron, B. , Shea-Brown, E. , and Josić, K. . Correlation between neural spike trains increases with firing rate. *Nature*, 2007. (Cited on page 71.)
- Dean, I. , Robinson, B. L. , and Harper, N. S. . Rapid neural adaptation to sound level statistics. *The Journal of Neuroscience*, 2008. (Cited on pages xvii, 61, and 62.)
- Deister, C. A. , Chan, C. S. , and Surmeier, D. J. . Calcium-activated SK channels influence voltage-gated ion channels to determine the precision of firing in globus pallidus neurons. *The Journal of Neuroscience*, 2009. (Cited on page 111.)
- Delaney, K. R. and Tank, D. W. . A quantitative measurement of the dependence of short-term synaptic enhancement on presynaptic residual calcium. *The Journal of Neuroscience*, 1994. (Cited on pages xx, 90, and 92.)
- Deneve, S. . Bayesian spiking neurons I: inference. *Neural computation*, 20(1):91–117, 2008. (Cited on pages 74, 104, 105, 106, and 159.)
- Derksen, H. E. and Verveen, A. A. . Fluctuations of resting neural membrane potential. *Science*, 1966. (Cited on page 25.)
- Díez-Sampedro, A. . Mechanism of Increased Open Probability by a Mutation of the BK Channel. *Journal of Neurophysiology*, 96(3):1507–1516, Jan. 2006. (Cited on pages xxi and 113.)
- Dobrunz, L. E. and Stevens, C. F. . Heterogeneity of release probability, facilitation, and depletion at central synapses. *Neuron*, 1997. (Cited on pages 81, 84, and 90.)
- Dodge, F. A. and Rahamimoff, R. . Cooperative action of calcium ions in transmitter release at the neuromuscular junction. *The Journal of physiology*, 1967. (Cited on pages 85 and 92.)
- Doidge, N. . The brain that changes itself: Stories of personal triumph from the frontiers of brain science, 2007. (Cited on page 22.)
- Down, T. A. , Rakyan, V. K. , Turner, D. J. , Flicek, P. , and Li, H. . A Bayesian deconvolution strategy for immunoprecipitation-based DNA methylome analysis. *Nature*, 2008. (Cited on page 15.)

- Drew, P. J. and Abbott, L. F. . Models and properties of power-law adaptation in neural systems. *Journal of Neurophysiology*, 2006. (Cited on page 68.)
- Duda, R. O. , Hart, P. E. , and Stork, D. G. . Pattern classification, 2012. (Cited on page 82.)
- Dudel, J. and Kuffler, S. W. . Mechanism of facilitation at the crayfish neuromuscular junction. *The Journal of physiology*, 1961. (Cited on page 86.)
- Eccles, J. C. . The physiology of synapses, 2013. (Cited on page 21.)
- Eccles, J. C. , Katz, B. , and Kuffler, S. W. . Nature of the "endplate potential" in curarized muscle. *Journal of Neurophysiology*, 1941. (Cited on page 85.)
- Ernst, M. O. and Banks, M. S. . Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870):429–433, Jan. 2002a. (Cited on page 3.)
- Ernst, M. O. and Banks, M. S. . Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 2002b. (Cited on pages 15 and 138.)
- Everitt, B. S. and Skrondal, A. . The Cambridge dictionary of statistics. *Cambridge*, 2002. (Cited on page 9.)
- Fairhall, A. . The receptive field is dead. Long live the receptive field? *Current Opinion in Neurobiology*, 2014. (Cited on page 71.)
- Fairhall, A. , Lewen, G. , Bialek, W. , and Steveninck, R. . Efficiency and ambiguity in an adaptive neural code. *Nature*, 412(6849):787–792, Aug. 2001. (Cited on pages 31 and 71.)
- Faisal, A. A. , White, J. A. , and Laughlin, S. B. . Ion-Channel Noise Places Limits on the Miniaturization of the Brain's Wiring. *Current Biology*, 15(12):1143–1149, June 2005. (Cited on pages 28 and 140.)
- Faisal, A. A. , Selen, L. P. J. , and Wolpert, D. M. . Noise in the nervous system. *Nature Reviews Neuroscience*, 9(4):292–303, Apr. 2008. (Cited on pages 3, 4, 24, 33, 140, and 148.)

- Feng, S. S. and Jaeger, D. . The role of SK calcium-dependent potassium currents in regulating the activity of deep cerebellar nucleus neurons: a dynamic clamp study. *The Cerebellum*, 2008. (Cited on page 65.)
- Feng, S. S. , Lin, R. , Gauck, V. , and Jaeger, D. . Gain Control of Synaptic Response Function in Cerebellar Nuclear Neurons by a Calcium-Activated Potassium Conductance. *The Cerebellum*, 12(5):692–706, Apr. 2013. (Cited on page 65.)
- Fiser, J. , Berkes, P. , Orban, G. , and Lengyel, M. . Statistically optimal perception and learning: from behavior to neural representations. *Trends In Cognitive Sciences*, 14(3):119–130, 2010. (Cited on page 83.)
- Fitzgerald, J. D. , Sincich, L. C. , and Sharpee, T. O. . Minimal models of multidimensional computations. 2011. (Cited on page 115.)
- Förster, J. , Higgins, E. T. , and Bianco, A. T. . Speed/accuracy decisions in task performance: Built-in trade-off or separate strategic concerns? *Organizational Behavior and Human Decision Processes*, 90(1):148–164, Jan. 2003. (Cited on page 154.)
- Frank, T. D. , Friedrich, R. , and Beek, P. J. . Stochastic order parameter equation of isometric force production revealed by drift-diffusion estimates. *Physical Review E*, 2006. (Cited on page 25.)
- Fuglevand, A. J. and Winter, D. A. . Models of recruitment and rate coding organization in motor-unit pools. *Journal of Neurophysiology*, 1993. (Cited on page 107.)
- Ganguli, D. and Simoncelli, E. P. . Implicit encoding of prior probabilities in optimal neural populations. *Advances in neural information processing systems (NIPS)*, 23:658–666, 2010. (Cited on pages 15, 28, 141, 146, 155, 159, and 162.)
- Gardos, G. . *The function of calcium in the potassium permeability of human erythrocytes*. *Biochimica et biophysica acta*, 1958. (Cited on page 111.)
- Garnier, J. , Gibrat, J. F. , and Robson, B. . GOR method for predicting protein secondary structure from amino acid sequence, 1996. (Cited on page 15.)

- Gautrais, J. and Thorpe, S. . Rate coding versus temporal order coding: a theoretical approach. *Biosystems*, 48(1-3):57–65, Nov. 1998. (Cited on page 108.)
- Geffen, M. N. , Broome, B. M. , Laurent, G. , and Meister, M. . Neural encoding of rapidly fluctuating odors. *Neuron*, 2009. (Cited on page 115.)
- Gelman, A. , Carlin, J. B. , Stern, H. S. , and Rubin, D. B. . Bayesian Data Analysis. 2003. (Cited on pages 34 and 96.)
- Golchi, S. and Lockhart, R. . A Bayesian Approach for Hypothesis Testing in Particle Physics. *arXiv.org*, 2015. (Cited on page 15.)
- Golding, N. L. , Kath, W. L. , and Spruston, N. . Dichotomy of Action-Potential Backpropagation in CA1 Pyramidal Neuron Dendrites. *Journal of Neurophysiology*, 86(6):2998–3010, Dec. 2001. (Cited on pages xxii, 129, and 130.)
- Green, P. E. and Frank, R. E. . Bayesian Statistics and Marketing Research. *Applied Statistics*, 1966. (Cited on page 15.)
- Gu, N. , Vervaeke, K. , and Storm, J. F. . BK potassium channels facilitate highfrequency firing and cause early spike frequency adaptation in rat CA1 hippocampal pyramidal cells. *The Journal of physiology*, 2007. (Cited on page 111.)
- Gurney, K. , Prescott, T. J. , and Redgrave, P. . A computational model of action selection in the basal ganglia. I. A new functional anatomy. *Biological Cybernetics*, 2001. (Cited on page 155.)
- Hansel, D. , Mato, G. , Meunier, C. , and Neltner, L. . On numerical simulations of integrate-and-fire neural networks. *Neural computation*, 1998. (Cited on page 20.)
- Harata, N. , Pyle, J. L. , Aravanis, A. M. , and Mozhayeva, M. . Limited numbers of recycling vesicles in small CNS nerve terminals: implications for neural signaling and vesicular cycling. *Trends in Neurosciences*, 2001. (Cited on page 90.)

- Harris, K. M. and Sultan, P. . Variation in the number, location and size of synaptic vesicles provides an anatomical basis for the nonuniform probability of release at hippocampal CA1 synapses. *Neuropharmacology*, 1995. (Cited on pages 81, 84, and 90.)
- Haruno, M. and Wolpert, D. M. . Optimal control of redundant muscles in step-tracking wrist movements. *Journal of Neurophysiology*, 2005. (Cited on page 26.)
- Hasson, R. and Swithenby, S. J. . A Bayesian test for the appropriateness of a model in the biomagnetic inverse problem. *Inverse problems*, 1999. (Cited on page 15.)
- Hebb, D. O. . *Distinctive features of learning in the higher animal*. Brain mechanisms and learning, 1961. (Cited on page 22.)
- Hodgkin, A. L. and Huxley, A. F. . A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of physiology*, 1952. (Cited on pages 17, 106, and 164.)
- Hosoya, T. , Baccus, S. A. , and Meister, M. . Dynamic predictive coding by the retina. *Nature*, 436(7047):71–77, July 2005. (Cited on pages 33 and 107.)
- Hughes, G. W. and Maffei, L. . Retinal ganglion cell response to sinusoidal light stimulation. *Journal of Neurophysiology*, 1966. (Cited on page 166.)
- Ikeda, K. and Bekkers, J. M. . Counting the number of releasable synaptic vesicles in a presynaptic terminal. *Proceedings of the National Academy of Sciences*, 2009. (Cited on page 90.)
- Jolivet, R. . Generalized Integrate-and-Fire Models of Neuronal Activity Approximate Spike Trains of a Detailed Model to a High Degree of Accuracy. *Journal of Neurophysiology*, 92(2):959–976, Mar. 2004. (Cited on page 20.)
- Jolivet, R. , Rauch, A. , Lüscher, H.-R. , and Gerstner, W. . Predicting spike timing of neocortical pyramidal neurons by simple threshold models. *Journal of Computational Neuroscience*, 21(1): 35–49, Apr. 2006. (Cited on page 106.)

- Jordan, M. I. . Why the logistic function? A tutorial discussion on probabilities and neural networks. 1995. (Cited on pages xviii, 76, and 78.)
- Kaardal, J. , Fitzgerald, J. D. , Berry, M. J. , and Sharpee, T. O. . Identifying functional bases for multidimensional neural computations. *Neural computation*, 2013. (Cited on page 115.)
- Kandel, E. R. , Schwartz, J. H. , and Jessell, T. M. . Principles of neural science, 2000. (Cited on page 25.)
- Kass, R. E. and Ventura, V. . A Spike-Train Probability Model. *Neural computation*, 13(8):1713–1720, Aug. 2001. (Cited on page 33.)
- Kato, S. , Xu, Y. , Cho, C. E. , Abbott, L. F. , and Bargmann, C. I. . Temporal responses of *C. elegans* chemosensory neurons are preserved in behavioral dynamics. *Neuron*, 81(3):616–628, Jan. 2014a. (Cited on page 66.)
- Kato, S. , Xu, Y. , Cho, C. E. , Abbott, L. F. , and Bargmann, C. I. . Temporal Responses of *C. elegans* Chemosensory Neurons Are Preserved in Behavioral Dynamics. *Neuron*, 81(3):616–628, Feb. 2014b. (Cited on page 129.)
- Katz, B. and Miledi, R. . The role of calcium in neuromuscular facilitation. *The Journal of physiology*, 1968. (Cited on page 87.)
- Katz, B. and Miledi, R. . Further study of the role of calcium in synaptic transmission. *The Journal of physiology*, 1970. (Cited on page 92.)
- Keat, J. , Reinagel, P. , Reid, R. C. , and Meister, M. . Predicting every spike: a model for the responses of visual neurons. *Neuron*, 2001. (Cited on page 115.)
- Kepecs, A. and Lisman, J. . Information encoding and computation with spikes and bursts. *Network: Computation in neural systems*, 2003. (Cited on page 70.)
- Kepecs, A. , Uchida, N. , Zariwala, H. A. , and Mainen, Z. F. . Neural correlates, computation and behavioural impact of decision confidence. *Nature*, 455(7210):227–231, Aug. 2008. (Cited on page 139.)

- Kim, K. J. and Rieke, F. . Temporal Contrast Adaptation in the Input and Output Signals of Salamander Retinal Ganglion Cells. *The Journal of Neuroscience*, 2001. (Cited on page 28.)
- Kleppe, I. C. and Robinson, H. . Correlation entropy of synaptic input-output dynamics. *Physical Review E*, 2006. (Cited on page 25.)
- Koehler, D. J. and James, G. . Probability matching in choice under uncertainty: Intuition versus deliberation. *Cognition*, 2009. (Cited on page 82.)
- Körding, K. P. and Wolpert, D. . Probabilistic inference in human sensorimotor processing. *Advances in neural information processing systems (NIPS)*, 2003. (Cited on pages 15 and 138.)
- Körding, K. P. and Wolpert, D. M. . Bayesian integration in sensorimotor learning. *Nature*, 427 (6971):244–247, 2004. (Cited on pages 15 and 138.)
- Laplace, P. S. . *Theorie analytique des probabilités*; translated as 'Concerning Probability'. The World of Mathematics, 1812. (Cited on page 10.)
- Laughlin, S. . Matching Coding to Scenes to Enhance Efficiency. In *Physical and biological processing of images*, pages 42–52. Springer Berlin Heidelberg, Berlin, Heidelberg, 1983. (Cited on page 141.)
- Laughlin, S. B. . A simple coding procedure enhances a neuron's information capacity. *Z Naturforsch*, 1981. (Cited on page 141.)
- Laughlin, S. B. . The role of sensory adaptation in the retina. *Journal of Experimental Biology*, 146 (1):39–62, Sept. 1989. (Cited on page 27.)
- Laughlin, S. B. , Steveninck, R. R. , de Ruyter van, and Anderson, J. C. . The metabolic cost of neural information. *Nature Neuroscience*, 1(1):36–41, May 1998. (Cited on pages 28, 72, and 107.)
- Lauria, E. . Bayesian Machine Learning. 2005. (Cited on page 15.)
- Lee, S.-H. , Kwan, A. C. , Zhang, S. , Phoumthippavong, V. , Flannery, J. G. , Masmanidis, S. C. , Taniguchi, H. , Huang, Z. J. , Zhang, F. , Boyden, E. S. , Deisseroth, K. , and Dan, Y. . Activation of specific interneurons improves V1 feature selectivity and visual perception. *Nature*, 488(7411):379–383, Apr. 2013. (Cited on page 106.)

- Legendy, C. R. and Salcman, M. . Bursts and recurrences of bursts in the spike trains of spontaneously active striate cortex neurons. *Journal of Neurophysiology*, 1985. (Cited on page 60.)
- Lennie, P. . The cost of cortical computation. *Current Biology*, 2003. (Cited on page 72.)
- Lindeberg, T. . A computational theory of visual receptive fields. *Biological Cybernetics*, 2013. (Cited on page 70.)
- Lindner, B. , Schimansky-Geier, L. , and Longtin, A. . Maximizing spike train coherence or incoherence in the leaky integrate-and-fire model. *Physical Review E*, 2002. (Cited on page 51.)
- Loewenstein, Y. and Sompolinsky, H. . Temporal integration by calcium dynamics in a model neuron. *Nature Neuroscience*, 6(9):961–967, Aug. 2003. (Cited on pages 65, 71, and 73.)
- London, M. and Häusser, M. . Dendritic Computation. *Annu Rev Neurosci*, 28(1):503–532, July 2005. (Cited on page 79.)
- Longtin, A. . Stochastic resonance in neuron models. *Journal of statistical physics*, 1993. (Cited on pages 4 and 140.)
- Longtin, A. , Bulsara, A. , and Moss, F. . Time-interval sequences in bistable systems and the noise-induced transmission of information by sensory neurons. *Physical Review Letters*, 1991. (Cited on page 26.)
- Louie, K. , Grattan, L. E. , and Glimcher, P. W. . Reward value-based gain control: divisive normalization in parietal cortex. *The Journal of Neuroscience*, 2011. (Cited on page 73.)
- Luo, X. , Gee, S. , Sohal, V. S. , and Small, D. S. . A Simple Probabilistic and Point-process Response Model for Predicting Every Spike in Optogenetics. *arXiv.org*, 2012. (Cited on page 115.)
- Lynch, M. A. . Long-term potentiation and memory. *Physiological reviews*, 2004. (Cited on page 22.)
- Ma, W. J. , Beck, J. M. , Latham, P. E. , and Pouget, A. . Bayesian inference with probabilistic population codes. *Nature Neuroscience*, 9(11):1432–1438, 2006. (Cited on pages 3, 139, and 153.)

- Maass, W. . Computing with spikes. *Special Issue on Foundations of Information Processing of TELEMATIK*, 2002. (Cited on page 64.)
- MacKay, D. G. . The problems of flexibility, fluency, and speed–accuracy trade-off in skilled behavior. *Psychological Review*, 1982. (Cited on page 154.)
- Magee, J. C. and Cook, E. P. . Somatic EPSP amplitude is independent of synapse location in hippocampal pyramidal neurons. *Nature Neuroscience*, 2000. (Cited on page 103.)
- Marella, S. , Fischler, W. , Kong, P. , Asgarian, S. , Rueckert, E. , and Scott, K. . Imaging Taste Responses in the Fly Brain Reveals a Functional Map of Taste Category and Behavior. *Neuron*, 49(2):285–295, Jan. 2006. (Cited on page 66.)
- Markram, H. , Gerstner, W. , and Sjöström, P. J. . Spike-timing-dependent plasticity: a comprehensive overview. *Frontiers in synaptic neuroscience*, 2012. (Cited on page 70.)
- Marr, D. and Poggio, T. . From understanding computation to understanding neural circuitry. 1976. (Cited on pages 132 and 158.)
- McCormick, T. H. , Raftery, A. E. , Madigan, D. , and Burd, R. S. . Dynamic logistic regression and dynamic model averaging for binary classification. *Biometrics*, 2012. (Cited on page 133.)
- McDonnell, M. D. and Abbott, D. . Stochastic Resonance, 2008. (Cited on pages 4 and 140.)
- Mihalaş, Ş. and Niebur, E. . A Generalized Linear Integrate-and-Fire Neural Model Produces Diverse Spiking Behaviors. *Neural computation*, 21(3):704–718, Mar. 2009. (Cited on pages 20 and 106.)
- Mittmann, W. , Wallace, D. J. , Czubayko, U. , and Herb, J. T. . Two-photon calcium imaging of evoked activity from L5 somatosensory neurons in vivo. *Nature*, 2011. (Cited on pages 71 and 166.)
- Moreno-Bote, R. and Parga, N. . Auto- and crosscorrelograms for the spike response of leaky integrate-and-fire neurons with slow synapses. *Physical Review Letters*, 2006. (Cited on page 51.)

- Moresco, E. M. Y. . Abl Family Nonreceptor Tyrosine Kinases Modulate Short-Term Synaptic Plasticity. *Journal of Neurophysiology*, 89(3):1678–1687, Nov. 2002. (Cited on pages xix, 89, and 91.)
- Moscovitch, M. . A Neuropsychological Approach to Perception and Memory in Normal and Pathological Aging. In *Aging and cognitive processes*, pages 55–78. Springer US, Boston, MA, 1982. (Cited on page 3.)
- Murray, R. F. , Patel, K. , and Yee, A. . Posterior Probability Matching and Human Perceptual Decision Making. *PLoS Computational Biology*, 2015. (Cited on page 82.)
- Murthy, V. N. , Sejnowski, T. J. , and Stevens, C. F. . Heterogeneous release properties of visualized individual hippocampal synapses. *Neuron*, 1997. (Cited on page 81.)
- Nedeljkov, M. and Oberguggenberger, M. . Ordinary differential equations with delta function terms. *Publications de l’Institut Mathematique*, 2012. (Cited on pages 47 and 63.)
- Neelakanta, P. S. , Sudhakar, R. , and DeGroff, D. . Langevin machine: a neural network based on stochastically justifiable sigmoidal function. *Biological Cybernetics*, 1991. (Cited on page 115.)
- Neher, E. and Sakaba, T. . Multiple roles of calcium ions in the regulation of neurotransmitter release. *Neuron*, 2008. (Cited on page 71.)
- Neiman, T. and Loewenstein, Y. . Reinforcement learning in professional basketball players. *Nature Communications*, 2:569, Dec. 2011. (Cited on page 15.)
- Neishabouri, A. and Faisal, A. A. . Energy constraints link structure and function in thin axons in the brain’s wiring. (Cited on page 4.)
- Neishabouri, A. and Faisal, A. A. . Energy constraints limit the sustainable firing rate of thin axons of the CNS and PNS. In *Ion Channels in health and disease*, 2013. (Cited on page 72.)
- Neishabouri, A. and Faisal, A. A. . Saltatory conduction in unmyelinated axons: clustering of Na⁺ channels on lipid rafts enables micro-saltatory conduction in C-fibers. *Frontiers in neuroanatomy*, 2014a. (Cited on page 159.)

- Neishabouri, A. and Faisal, A. A. . Axonal noise as a source of synaptic variability. 2014b. (Cited on page 4.)
- Neishabouri, A. M. and Faisal, A. A. . The metabolic efficiency of myelinated vs unmyelinated axons. *BMC Neuroscience*, 2011. (Cited on pages 72 and 159.)
- Nirenberg, S. and Latham, P. E. . Decoding neuronal spike trains: How important are correlations? In *Proceedings of the National Academy of Sciences*, 2003. (Cited on page 149.)
- Nonacs, P. and Soriano, J. L. . Patch sampling behaviour and future foraging expectations in Argentine ants, *Linepithema humile*. *Animal Behaviour*, 55(3):519–527, Mar. 1998. (Cited on page 15.)
- Olsen, S. R. , Bhandawat, V. , and Wilson, R. I. . Divisive Normalization in Olfactory Population Codes. *Neuron*, 66(2):287–299, Apr. 2010. (Cited on page 32.)
- Olshausen, B. A. and Field, D. J. . Natural image statistics and efficient coding*. *Network: Computation in neural systems*, 1996. (Cited on page 71.)
- Paninski, L. and Fellows, M. R. . Spatiotemporal tuning of motor cortical neurons for hand position and velocity. *Journal of Neurophysiology*, 2004. (Cited on page 106.)
- Penny, W. D. and Roberts, S. J. . Dynamic logistic regression. *Neural Networks*, 1999. (Cited on page 133.)
- Penrose, R. . Consciousness, the brain, and spacetime geometry: an addendum. *Annals of the New York Academy of Sciences*, 2001. (Cited on page 2.)
- Piazza, M. , Izard, V. , Pinel, P. , Le Bihan, D. , and Dehaene, S. . Tuning curves for approximate numerosity in the human intraparietal sulcus. *Neuron*, 2004. (Cited on page 71.)
- Pohlmann, K. C. . Principles of digital audio, 1985. (Cited on pages 4 and 26.)
- Prescott, S. A. and Sejnowski, T. J. . Spike-rate coding and spike-time coding are affected oppositely by different adaptation mechanisms. *The Journal of Neuroscience*, 2008. (Cited on page 71.)

- Rao, R. P. N. and Ballard, D. H. . Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1):79–87, Jan. 1999. (Cited on pages 33 and 72.)
- Rao-Mirotznik, R. , Harkins, A. B. , Buchsbaum, G. , and Sterling, P. . Mammalian rod terminal: architecture of a binary synapse. *Neuron*, 1995. (Cited on page 90.)
- Rayleigh, L. . XII. On our perception of sound direction. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 1907. (Cited on page 2.)
- Reich, D. S. , Victor, J. D. , and Knight, B. W. . Response variability and timing precision of neuronal spike trains in vivo. *Journal of Neurophysiology*, 1997. (Cited on page 71.)
- Roberts, L. G. . Picture coding using pseudo-random noise. *Information Theory*, 1962. (Cited on pages 4 and 26.)
- Roberts, P. D. and Bell, C. C. . Spike timing dependent synaptic plasticity in biological systems. *Biological Cybernetics*, 2002. (Cited on page 23.)
- Rossi, P. E. and Allenby, G. M. . Bayesian statistics and marketing. *Marketing Science*, 2003. (Cited on page 15.)
- Roxin, A. , Brunel, N. , Hansel, D. , Mongillo, G. , Vreeswijk, C. , van, Roxin, A. , Brunel, N. , Hansel, D. , Mongillo, G. , and Vreeswijk, C. , van. On the Distribution of Firing Rates in Networks of Cortical Neurons. *Journal of Neuroscience*, 31(45):16217–16226, Nov. 2011. (Cited on pages 35 and 115.)
- Rudolph, M. and Destexhe, A. . Analytical integrate-and-fire neuron models with conductance-based dynamics for event-driven simulation strategies. *Neural computation*, 2006. (Cited on page 20.)
- Saaty, T. L. and Vargas, L. G. . Diagnosis with dependent symptoms: Bayes theorem and the analytic hierarchy process. *Operations Research*, 1998. (Cited on page 15.)

- Sacerdote, L. and Giraudo, M. T. . Stochastic Integrate and Fire Models: a review on mathematical methods and their applications. *arXiv.org*, Jan. 2011. (Cited on page 106.)
- Sarpeshkar, R. . Analog Versus Digital: Extrapolating from Electronics to Neurobiology. *Neural computation*, 10(7):1601–1638, Oct. 1998. (Cited on page 3.)
- Schatz, C. J. . The developing brain. *Scientific American*, 1992. (Cited on page 22.)
- Schikorski, T. and Stevens, C. F. . Quantitative ultrastructural analysis of hippocampal excitatory synapses. *The Journal of Neuroscience*, 1997. (Cited on page 90.)
- Schikorski, T. and Stevens, C. F. . Quantitative fine-structural analysis of olfactory cortical synapses. In *Proceedings of the National Academy of Sciences*, 1999. (Cited on page 90.)
- Schiller, J. , Schiller, Y. , and Stuart, G. . Calcium action potentials restricted to distal apical dendrites of rat neocortical pyramidal neurons. *The Journal of physiology*, 1997. (Cited on pages xxii, 65, 129, and 130.)
- Schneggenburger, R. and Neher, E. . Intracellular calcium dependence of transmitter release rates at a fast central synapse. *Nature*, 406(6798):889–893, Aug. 2000. (Cited on pages 65 and 73.)
- Schneggenburger, R. and Rosenmund, C. . Molecular mechanisms governing Ca²⁺ regulation of evoked and spontaneous release. *Nature Neuroscience*, 18(7):935–941, June 2015. (Cited on pages xix, 85, and 86.)
- Schneggenburger, R. , Sakaba, T. , and Neher, E. . Vesicle pools and short-term synaptic depression: lessons from a large synapse. *Trends in Neurosciences*, 2002. (Cited on page 90.)
- Schrodinger, E. and Lewin. What is life?, 1967. (Cited on page 2.)
- Schum, D. A. . The evidential foundations of probabilistic reasoning, 2001. (Cited on page 11.)
- Sengupta, B. , Faisal, A. A. , Laughlin, S. B. , and Niven, J. E. . The effect of cell size and channel density on neuronal information encoding and energy efficiency. *Journal of Cerebral Blood Flow & Metabolism*, 33(9):1465–1473, June 2013. (Cited on pages 28 and 32.)

- Sernagor, E. , Eglén, S. J. , and Wong, R. . Development of retinal ganglion cell structure and function. *Progress in retinal and eye research*, 2001. (Cited on pages xv and 49.)
- Shadlen, M. N. and Newsome, W. T. . Noise, neural codes and cortical organization. *Current Opinion in Neurobiology*, 4(4):569–579, Aug. 1994. (Cited on pages 4, 26, 33, and 70.)
- Shepard, N. . Integrate and Fire Model. 2007. (Cited on page 73.)
- Smith, M. R. , Nelson, A. B. , and Lac, S. , du. Regulation of Firing Response Gain by Calcium-Dependent Mechanisms in Vestibular Nucleus Neurons. *Journal of Neurophysiology*, 87(4):2031–2042, Apr. 2002. (Cited on pages 50, 65, and 73.)
- Srinivasan, M. V. and Laughlin, S. B. . Predictive coding: a fresh view of inhibition in the retina. *Proceedings of the Royal Society of London*, 1982. (Cited on page 72.)
- Stein, R. B. , Gossen, E. R. , and Jones, K. E. . Neuronal variability: noise or part of the signal? *Nature Reviews Neuroscience*, 2005. (Cited on page 74.)
- Stockbridge, N. and Moore, J. W. . Dynamics of intracellular calcium and its possible relationship to phasic transmitter release and facilitation at the frog neuromuscular junction. 1984. (Cited on page 92.)
- Stockman, A. , Langendörfer, M. , and Smithson, H. E. . Human cone light adaptation: From behavioral measurements to molecular mechanisms. *Journal of Vision*, 2006. (Cited on page 27.)
- Suchard, M. A. and Redelings, B. D. . BAli-Phy: simultaneous Bayesian inference of alignment and phylogeny. *Bioinformatics*, 2006. (Cited on page 15.)
- Sullivan, D. A. and Holmqvist, M. H. . Characterization of gating and peptide block of mSlo, a cloned calcium-dependent potassium channel. *Journal of Neurophysiology*, 1997. (Cited on pages xxi, 111, and 113.)
- Tassinari, H. , Hudson, T. E. , and Landy, M. S. . Combining priors and noisy visual cues in a rapid pointing task. *The Journal of Neuroscience*, 26(40):10154–10163, 2006. (Cited on page 172.)

- Ticchi, A. and Faisal, A. A. . A model for neuronal learning and sensing in a noisy Arzigogoli environment through Bayesian Inference. In *Proceedings of Champalimaud Neuroscience Symposium*, 2012. (Cited on page 139.)
- Ticchi, A. and Faisal, A. A. . A Biophysical Model of Bayesian Inference and MCMC Sampling in Neural Circuits. In *Computational and Systems Neuroscience (Cosyne)*, 2013a. (Cited on page 139.)
- Ticchi, A. and Faisal, A. A. . A Neural Circuit Model of Bayesian Inference via MCMC exploiting Neuronal Noise. In *Proceedings of Society for Neuroscience annual meeting*, 2013b. (Cited on page 139.)
- Ticchi, A. and Faisal, A. A. . Bayes optimal spike adaptation predicts biophysical characteristics of neurons. In *Computational and Systems Neuroscience (Cosyne)*, 2014. (Cited on page 29.)
- Ticchi, A. and Faisal, A. A. . A normative model of neural adaptation via local approximate Bayesian inference predicts V1 response to dynamic stimuli. In *Computational and Systems Neuroscience (Cosyne)*, 2015. (Cited on page 52.)
- Tolhurst, D. J. , Movshon, J. A. , and Dean, A. F. . The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Research*, 1983. (Cited on pages 4 and 24.)
- Trussell, L. O. and Jackson, M. B. . Adenosine-activated potassium conductance in cultured striatal neurons. In *Proceedings of the National Academy of Sciences*, 1985. (Cited on page 111.)
- Turrigiano, G. . The self-tuning neuron: synaptic scaling of excitatory synapses. *Cell*, Jan. 2008. (Cited on page 32.)
- Turrigiano, G. , Abbott, L. F. , and Marder, E. . Activity-dependent changes in the intrinsic properties of cultured neurons. *Science AAAS Weekly Paper Edition*, 1994. (Cited on page 164.)
- Ulichney, R. . Dithering with blue noise. In *Proceedings of the IEEE*, 1988. (Cited on page 4.)
- Usher, M. and McClelland, J. L. . The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, 108(3):550–592, 2001. (Cited on page 154.)

- Vergara, C. , Latorre, R. , and Marrion, N. V. . Calcium-activated potassium channels. *Current Opinion in Neurobiology*, 1998. (Cited on page 111.)
- Wald, A. . Statistical decision functions. 1950. (Cited on page 10.)
- Wang, L. Y. and Kaczmarek, L. K. . High-frequency firing helps replenish the readily releasable pool of synaptic vesicles. *Nature*, 1998. (Cited on page 92.)
- Wark, B. , Lundstrom, B. N. , and Fairhall, A. . Sensory adaptation. *Current Opinion in Neurobiology*, 17(4):423–429, Aug. 2007. (Cited on page 32.)
- Warland, D. K. and Reinagel, P. . Decoding visual information from a population of retinal ganglion cells. *Journal of Neurophysiology*, 1997. (Cited on page 64.)
- Webster, M. A. , Kaping, D. , Mizokami, Y. , and Duhamel, P. . Adaptation to natural facial categories. *Nature*, 2004. (Cited on page 27.)
- Westerman, L. A. and Smith, R. L. . Rapid and short-term adaptation in auditory nerve responses. *Hearing research*, 1984. (Cited on pages xvii, 61, and 62.)
- Williams-García, R. V. , Moore, M. , Beggs, J. M. , and Ortiz, G. . Quasicritical brain dynamics on a nonequilibrium Widom line. *Physical Review E*, 2014. (Cited on page 16.)
- Wilson, R. I. , Turner, G. C. , and Laurent, G. . Transformation of olfactory representations in the *Drosophila* antennal lobe. *Science*, 2004. (Cited on page 71.)
- Witton, J. , Padmashri, R. , Zinyuk, L. E. , Popov, V. I. , Kraev, I. , Line, S. J. , Jensen, T. P. , Tedoldi, A. , Cummings, D. M. , Tybulewicz, V. L. J. , Fisher, E. M. C. , Bannerman, D. M. , Randall, A. D. , Brown, J. T. , Edwards, F. A. , Rusakov, D. A. , Stewart, M. G. , and Jones, M. W. . Hippocampal circuit dysfunction in the Tc1 mouse model of Down syndrome. *Nature Neuroscience*, 18(9):1291–1298, Aug. 2015. (Cited on pages xix, 89, and 91.)
- Wolpert, D. M. . Computational approaches to motor control. *Trends In Cognitive Sciences*, 1(6): 209–216, Sept. 1997. (Cited on pages 3 and 15.)

- Wolpert, D. M. , Ghahramani, Z. , and Jordan, M. I. . An internal model for sensorimotor integration. *Science AAAS Weekly Paper Edition*, 1995. (Cited on pages 15 and 138.)
- Wozny, D. R. , Beierholm, U. R. , and Shams, L. . Probability matching as a computational strategy used in perception. *PLoS Computational Biology*, 6(8):e1000871, 2010. (Cited on pages 3, 15, 82, 155, and 172.)
- Wu, L. G. and Borst, J. . The reduced release probability of releasable vesicles during recovery from short-term synaptic depression. *Neuron*, 1999. (Cited on page 92.)
- Yang, Z. . Maximum-likelihood estimation of phylogeny from DNA sequences when substitution rates differ over sites. *Molecular Biology and Evolution*, 1993. (Cited on page 15.)
- Yarrow, S. and Challis, E. . Fisher and Shannon Information in Finite Neural Populations. *Neural computation*, 2012. (Cited on page 106.)
- Yeomans, J. S. . The absolute refractory periods of self-stimulation neurons. *Physiology & Behavior*, 22(5):911–919, May 1979. (Cited on page 32.)
- Zador, A. , Koch, C. , and Brown, T. H. . Biophysical model of a Hebbian synapse. In *Proceedings of the National Academy of Sciences*, 1990. (Cited on page 64.)
- Zaghloul, K. A. , Boahen, K. , and Demb, J. B. . Contrast Adaptation in Subthreshold and Spiking Responses of Mammalian Y-Type Retinal Ganglion Cells. *The Journal of Neuroscience*, 2005. (Cited on page 28.)
- Zucker, R. S. . Calcium-and activity-dependent synaptic plasticity. *Current Opinion in Neurobiology*, 1999. (Cited on page 133.)
- Zucker, R. S. and Regehr, W. G. . Short-Term Synaptic Plasticity. *Annual review of physiology*, 64(1):355–405, Mar. 2002. (Cited on pages 85 and 90.)