

UNIVERSITY OF CYPRUS

DOCTORAL THESIS

**Understanding the neural code through
exploration of the causes of firing**

Author:

Achilleas Koutsou

Supervisor:

Dr. Chris Christodoulou

*A dissertation submitted in fulfilment of the requirements
for the degree of Doctor of Philosophy*

in the

Department of Computer Science

June, 2015

© Achilleas Koutsou, 2015

Achilleas Koutsou

APPROVAL PAGE

Doctor of Philosophy Dissertation

Understanding the neural code through exploration of the causes of firing

Presented by
Achilleas Koutsou

Research Supervisor _____
Chris Christodoulou

Committee Member _____
Christos N. Schizas

Committee Member _____
Constantinos Pattichis

Committee Member _____
Alessandro Villa

Committee Member _____
Panayiota Poirazi

Declaration of Doctoral Candidate

The present doctoral dissertation was submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy of the University of Cyprus. It is a product of original work of my own, unless otherwise mentioned through references, notes, or any other statements.

Achilleas Koutsou

Περίληψη

Τμήμα Πληροφορικής

Κατανόηση του νευρωνικού κώδικα μέσω της μελέτης των αιτιών της πυροδότησης

Αχιλλέας Κουτσού

Όταν αναφερόμαστε στον νευρωνικό κώδικα εννοούμε τους μηχανισμούς με τους οποίους οι νευρώνες και τα νευρωνικά δίκτυα μετατρέπουν πληροφορίες σε ακολουθίες από πυροδοτήσεις. Η ανακάλυψη και κατανόηση αυτών των μηχανισμών είναι θεμελιώδη βήματα προς την κατανόηση των μεθόδων με τις οποίες ο εγκέφαλος κωδικοποιεί, αποκωδικοποιεί και επεξεργάζεται πληροφορίες. Βασικός παράγοντας προς την επίλυση του προβλήματος του νευρωνικού κώδικα είναι η δυνατότητα να προσδιορίσουμε τον τρόπο λειτουργίας ενός νευρώνα: αν δηλαδή λειτουργεί ως χρονικός συναθροιστής δυναμικών (temporal integrator) ή ανιχνευτής ταυτόχρονων δυναμικών (coincidence detector). Γενικότερα, αυτό το πρόβλημα μπορεί να εκφραστεί με όρους χρονικών κλιμάκων της επεξεργασίας σημάτων από τον νευρώνα. Μπορούμε να το επιλύσουμε μετρώντας την χρονική ακρίβεια με την οποία ένας νευρώνας διακρίνει διαφορετικά ερεθίσματα. Αυτή η προσέγγιση γενικεύει την έννοια του τρόπου λειτουργίας του νευρώνα από τους δύο εναλλακτικούς τρόπους λειτουργίας σε ένα συνεχές φάσμα λειτουργιών, στα άκρα του οποίου βρίσκονται οι δύο ακραίοι τρόποι λειτουργίας: η χρονική συνάθροιση δυναμικών και η ανίχνευση ταυτόχρονων δυναμικών. Η παρούσα διατριβή προτείνει διάφορες μεθόδους για την επίλυση του προβλήματος του νευρωνικού κώδικα και συγκεκριμένα τις χρονικές κλίμακες με τις οποίες οι νευρώνες επεξεργάζονται πληροφορίες. Συγκεκριμένα, ο σκοπός της διατριβής αυτής είναι να παρουσιάσει μεθόδους με τις οποίες μπορούμε να υπολογίσουμε τον τρόπο λειτουργίας ενός νευρώνα μέσω προσομοιώσεων και μέσω των υπολογισμών αυτών να κατανοήσουμε τους τρόπους με τους οποίους ο τρόπος λειτουργίας ενός νευρώνα ορίζεται από τις ενδογενείς και εξωγενείς ιδιότητες του.

Η πρώτη μας συνεισφορά στην επίλυση του προβλήματος είναι η ανάπτυξη μιας μεθόδου που υπολογίζει τον τρόπο λειτουργίας ενός νευρώνα μέσω της κλίσης της εκπόλωσης του δυναμικού της μεμβράνης του, πριν την πυροδότηση ενός δυναμικού ενεργείας. Δείχνουμε ότι η μέθοδός μας μπορεί να προσδιορίσει το βαθμό συγχρονισμού που είναι υπεύθυνος για την πυροδότηση δυναμικών ενεργείας σε ένα απλό μοντέλο νευρώνα και

περιγράφουμε πώς η μέτρηση αυτή είναι ισοδύναμη με τον τρόπο λειτουργίας του νευρώνα. Χρησιμοποιώντας αυτή τη μέθοδο, δείχνουμε ότι ο τρόπος λειτουργίας ενός νευρώνα με μηχανισμό μερικής επαναπόλωσης, όταν πυροδοτεί ακανόνιστα σε πολύ υψηλές συχνότητες, είναι κατά κύριο λόγο χρονικός συναθροιστής δυναμικών με μικρή χρονική σταθερά διαρροής δυναμικού και χαμηλό κατώφλι. Επιπλέον, συμπεραίνουμε ότι σε πολύ υψηλές συχνότητες πυροδότησης, το συνεχές φάσμα λειτουργιών του νευρώνα μικραίνει, καθώς οι ορισμοί των δύο άκρων συγκλίνουν μέχρι τα δύο άκρα να γίνουν δυσδιάκριτα. Στη συνέχεια, διερευνούμε πώς ο συγχρονισμός στην είσοδο του νευρώνα συσχετίζεται με τον τρόπο λειτουργίας του και, κυριότερα, πώς αυτή η σχέση διαμορφώνεται από τις ιδιότητες του νευρώνα και των εισερχομένων ακολουθιών δυναμικών ενεργειών. Ανακαλύπτουμε ότι η σχέση μεταξύ συγχρονισμού στην είσοδο και τρόπου λειτουργίας καθορίζεται κυρίως από το πλάτος της εκπόλωσης που προκαλεί μία ενιαία, συγχρονισμένη άφιξη δυναμικών ενεργειών από τους προσυναπτικούς νευρώνες του κυττάρου και οποιαδήποτε άλλα χαρακτηριστικά των εισόδων, όπως ο ρυθμός άφιξης δυναμικών, δεν επηρεάζει τον τρόπο λειτουργίας.

Επίσης αναπτύσσουμε μία μέθοδο για τον υπολογισμό των παραμέτρων του σήματος εισόδου ενός απλού μοντέλου νευρώνα, κάτω από ορισμένες προϋποθέσεις. Μέσα από αυτές τις παραμέτρους, μπορούμε να ανασηματίσουμε το σήμα στην είσοδο του νευρώνα και να εξαγάγουμε συμπεράσματα για τον βαθμό συγχρονισμού των δυναμικών ενεργειών στην είσοδο και, κατ' επέκταση, τον τρόπο λειτουργίας. Αυτή η μέθοδος αναδιαμόρφωσης του σήματος στην είσοδο διαφέρει από την μέθοδο που βασίζεται στην κλίση του δυναμικού της μεμβράνης στον τρόπο με τον οποίο περιγράφεται η συμπεριφορά του νευρώνα: θεωρεί ότι το σήμα εισόδου αποτελείται από περιοδικές εκπυροσκοκροτήσεις συσχετισμένων δυναμικών ενεργειών, με μικρές τυχαίες αποκλίσεις στους χρόνους πυροδοτήσεων, το οποίο προσομοιώνουμε με τη χρήση ενός θορυβώδους ημιτονοειδούς κύματος. Δείχνουμε ότι όταν οι υποθέσεις μας ισχύουν, η μέθοδος μας μπορεί να υπολογίσει τις παραμέτρους του σήματος εισόδου με πολύ υψηλή ακρίβεια.

Τέλος, μελετάμε τη χρονική ακρίβεια της ανίχνευσης ταυτόχρονων δυναμικών σε ένα μορφολογικά ανακατασκευασμένο μοντέλο νευρώνα. Η έρευνα αυτή επικεντρώνεται στην ανακάλυψη των μηχανισμών μάθησης που επιτρέπουν στο κύτταρο να μάθει να διακρίνει ζεύγη δυναμικών ενεργειών με μικρές διαφορές στο χρόνο άφιξής τους στην είσοδο. Γενικότερα, διερευνούμε τα όρια της ικανότητας του νευρώνα να διακρίνει αυτού του είδους σήματα και τον τρόπο με τον οποίο αυτά τα όρια επηρεάζονται από τις διάφορες ιδιότητες των σημάτων στην είσοδο. Η ανάλυσή μας δείχνει ότι, για το μοντέλο που χρησιμοποιήθηκε στην έρευνά μας, η ακρίβεια δεν είναι αρκετά υψηλή ώστε να μπορεί ο νευρώνας να διακρίνει μικρές διαφορές (της τάξης των μερικών χιλιοστών του δευτερολέπτου) στις καθυστερήσεις μεταξύ ζευγών δυναμικών ενεργειών στην είσοδο. Δείχνουμε επίσης ότι η ευαισθησία του νευρώνα σε ζεύγη δυναμικών ενεργειών στην

είσοδο με ψηλή ακρίβεια στους χρόνους καθυστερήσεων μεταξύ τους, μειώνεται όταν η μεμβράνη είναι εκπολωμένη από ένα σταθερό ρεύμα.

Όπως φαίνεται από τα αποτελέσματα και τις συνεισφορές μας, δεν παρέχουμε οριστικές απαντήσεις στο ερώτημα του τρόπου λειτουργίας του νευρώνα. Αντί αυτού, φτάνουμε σε γενικά συμπεράσματα σχετικά με τους τρόπους με τους οποίους ο τρόπος λειτουργίας ενός νευρώνα ορίζεται από διάφορες ιδιότητες του νευρώνα και της συμπεριφοράς του.

Achilleas Koutsou

Abstract

Department of Computer Science

Understanding the neural code through exploration of the causes of firing

by Achilleas Koutsou

The neural code refers to the mechanisms with which single cells and networks of neurons transform information into sequences of spike trains. Discovering and understanding these mechanisms is fundamental to understanding information encoding, decoding and processing in the brain. A key aspect of solving the problem of the neural code is the ability to determine the operational mode of a single cell: whether it is a *temporal integrator* or a *coincidence detector*. More generally, this problem is a question of time scales of neural processing and in particular, it can be solved by identifying the temporal precision with which a neuron can distinguish between stimuli. This approach generalises the idea of a binary operational mode to a continuum which lies between the two extremes—temporal integration and coincidence detection.

This thesis proposes a number of methods for addressing the problem of the neural code and more specifically the time scales of neural processing, which define the operational mode of neurons. In particular, the purpose of this thesis is to provide methods with which the operational mode of simulated single neurons can be measured and through these measurements, understand the ways in which the operational mode is shaped by features of the neuron and its input. We present methods which we developed to identify the operational mode of single neurons, by observing their behaviour under simulated conditions. We aim to understand how the operational mode is defined by the the intrinsic and extrinsic properties of a cell being observed.

Our first contribution towards solving this problem is the development of a measure of the operational mode which depends on observations of the slope of the membrane potential of the neuron immediately prior to firing. We show that our measure can identify the degree of synchrony that is responsible for firing spikes in a simple neuron model and describe how this measurement is equivalent to the operational mode. Using this measure, we show that the operational mode of a neuron with partial somatic reset, when firing highly irregular spike trains at high rates, is primarily a temporal integrator with short leak time constant and low threshold. In addition, at very high rates, the continuum of operational modes shrinks as the definitions of the two extremes converge and they become indistinguishable.

We then investigate how input synchrony relates to the operational mode and more importantly, how this relationship is shaped by the properties of the neuron and the input spike trains. We find that the relationship between input synchrony and operational mode is primarily determined by the amplitude of depolarisation caused by a single, synchronous arrival of spikes from a pre-synaptic population and that other features of the input, such as spike rates, do not affect the operational mode.

We also developed a method for estimating the parameters of the input of a simple neuron model, under certain assumptions. Through these parameters, we can reconstruct the input signal driving the neuron and make inferences about the degree of synchrony driving the cell and by extension, the operational mode. This input reconstruction technique differs from the membrane potential slope-based measure in the way that the neuron's behaviour is formalised: the input is assumed to be composed of periodic bursts of correlated spike volleys, with small random deviations in spike times, which we simulate using a noisy sine wave. We show that when our assumptions hold, our method can determine the parameters of the input signal with very high accuracy.

Finally, we investigate the temporal precision of coincidence detection in a morphologically reconstructed neuron model. Our work focuses on discovering learning mechanisms that allow the cell to learn to distinguish input spike pairs with small differences in delays. More generally, we investigate the limits of the ability of the neuron to distinguish such signals and the way in which these limits are affected by various properties of the input. Our analysis revealed that, for the model used in our work, the precision is not high enough to distinguish between small differences (on the order of a few milliseconds) in delays between input spikes. We also show that the neuron's sensitivity to precise input spike pairs is reduced when the membrane is depolarised by a constant background current.

As seen by the results and our contributions, we did not actually provide a definitive answer to the question of which operational mode neurons employ. Instead, as indicated above, we arrived at general conclusions regarding the ways in which the operational mode is defined by various properties of the neuron and its behaviour.

Acknowledgements

I would like to express my appreciation and thanks to my advisor, Dr Chris Christodoulou, for his guidance and support throughout the entirety of my PhD research. His invaluable supervision is the reason I am able to call myself a research scientist. I am as grateful for the freedom he gave me to carve my own path and develop my ideas as I am of the times when he struggled to keep me focused and productive. I am especially thankful to him for introducing me to the field of Computational Neuroscience, which became my primary field of research.

I would also like to thank my main collaborators, Guido Bugmann and Jacob Kanev. Their ideas and comments, expressed during our long meetings, were instrumental in shaping my research and the content of this thesis.

I am grateful to Petr Lánský for inspiring discussions and for his suggestion to use the OU model. I am especially grateful for his kind hospitality during my one month stay at the Academy of Sciences of the Czech Republic in early 2012.

I gratefully acknowledge the support of the University of Cyprus for an Internal Research Project Grant and a scholarship under the programme “Young Researchers” (Νέοι Ερευνητές). I would also like to acknowledge the projects EGI – European Grid Infrastructure and “eMammoth – Compute and Store on Grids and Clouds infrastructure” (ANABATHMISI/06609/09), which is co-funded by the Republic of Cyprus and the European Regional Development Fund of the EU.

I am also grateful to my labmates (past and present), Margarita Zachariou, Vassilis Vassiliades, Michalis Agathocleous, and Petros Kountouris, for making the lab an enjoyable working environment as well as for stimulating casual discussions on research topics and helping me develop (while appropriately scrutinising) small and big ideas.

Finally, I am very thankful to my family for making everything possible and always supporting me in my educational, professional and academic choices, as well as my friends who supported me throughout the years and continue to provide emotional support (mostly through alcohol).

To my loving, caring wife, Gloria, for sharing with me the best of times and putting up with me during the worst of times. Her love and support is the reason I managed to keep my sanity.

Contents

Declaration of Doctoral Candidate	iii
Abstract in Greek	iv
Abstract	vii
Acknowledgements	ix
List of Figures	xiv
List of Tables	xvi
Abbreviations	xvii
Symbols	xviii
1 Introduction	1
1.1 Computational neuroscience	2
1.2 Thesis outline	5
2 Literature review	8
2.1 The neural code	8
2.2 Neural operational modes	12
2.3 Synchrony	17
2.4 Input reconstruction and inference	20
3 Determining the causes of firing using the membrane potential slope	22
3.1 Neuron models	23
3.2 LIF and LIFwPR neuron models	23
3.3 Input generation	24
3.4 Pre-spike membrane potential slope	26
3.5 Results for LIF	30
3.5.1 Upper- and lower-bound convergence at high firing rates	33
3.6 Results for LIFwPR	35
3.6.1 Compatibility of results with original analysis	36

3.7	Discussion	38
4	The relation between stimulus synchrony and the operational mode of a neuron	44
4.1	Methods	46
4.1.1	SPIKE-distance	47
4.1.2	Comparison	48
4.2	Results	49
4.2.1	Effect of jitter (σ_{in})	49
4.2.2	Effect of threshold-free potential (V_{∞})	51
4.2.3	Effect of volley peak potential (Δ_v)	52
4.2.4	Integration of coincident volleys	52
4.2.5	Coincidence detection	55
4.2.6	Effect of input parameter values ($N_{in}, f_{in}, \Delta V_s$)	55
4.2.7	Summary of results	56
4.3	Discussion	58
5	Input synchrony estimation for the Ornstein Uhlenbeck LIF	62
5.1	Operational mode of the Ornstein Uhlenbeck LIF	63
5.1.1	Comparison of LIF with OU model	64
5.2	Estimation of frequency	66
5.3	Estimation of μ	68
5.3.1	Estimation of μ_p	69
5.3.2	Estimation of μ_0	69
5.4	Estimation of σ	70
5.4.1	Estimation of σ_p and σ_0	71
5.5	Results	71
5.5.1	Frequency estimation results	72
5.5.2	$\hat{\mu}$ estimation results	74
5.5.3	$\hat{\sigma}$ estimation results	76
5.5.4	Frequency misestimation	76
5.6	Discussion	77
6	Can a biophysical model of a pyramidal neuron learn time delays between the spikes from different input neurons?	83
6.1	Methods	85
6.1.1	Overview	85
6.1.2	Model	87
6.1.3	Synaptic weights and scaling	88
6.1.4	Simulation procedure	91
6.2	Results	92
6.2.1	Approach 1: Backpropagating coincident EPSPs	92
6.2.2	Approach 2: Backpropagating action potential	94
6.2.2.1	Relative timing conditions for action potential generation	94
6.2.2.2	Action potential timing	96
6.2.2.3	Modelling the selectivity	97
6.2.2.4	Effect of background current on selectivity	99
6.2.2.5	Effect of background current on frequency of learning feedback	101

6.3	Discussion	103
7	General Discussion & Conclusions	105
7.1	NPSS: A measure of the operational mode	108
7.2	Input synchrony and the operational mode	109
7.3	Input synchrony of the Ornstein-Uhlenbeck LIF	111
7.4	Learning temporal delays between input spike pairs	112
7.5	Simulation, theory and experiment	116
7.5.1	Simulations	116
7.5.2	Time scales: Theory and experiment	117
7.5.3	Biological relevance of results and conclusions	118
7.6	General contributions	119
7.6.1	Contributions computational neuroscience	119
7.6.2	Contributions to computer science	120
7.7	Dissemination of PhD work	122
8	Future work	125
	Bibliography	129

List of Figures

2.1	Neural encoding/decoding illustration	9
2.2	Spike-triggered average schematic	11
2.3	Illustration of spike train binning	16
3.1	Sample input spike trains showing the effects of the synchrony parameters, S_{in} & σ_{in}	26
3.2	Illustration of coincidence window, slopes and bounds used in NPSS	29
3.3	NPSS measurements for 6 cases across entire range of synchrony parameters.	31
3.4	Correlation of NPSS with (a) S_{in} when $\sigma_{in} = 0$ ms and (b) σ_{in} when $S_{in} = 1$	32
3.5	Convergence of (a) LIF with perfect integrator and (b) slope bounds as a function of the ISI.	35
3.7	NPSS measurements for partial reset model firing highly irregular, high rate spike trains	39
4.1	Illustration of the values used in the calculation of the SPIKE-distance	48
4.2	NPSS vs SPIKE-distance with colour coded jitter σ_{in} . Subfigures separated based on value of σ_{in}	50
4.3	NPSS vs SPIKE-distance with colour coded asymptotic threshold-free potential V_{∞} . Subfigures separated based on value of V_{∞}	52
4.4	NPSS vs SPIKE-distance with colour coded volley peak Δ_v . Subfigures separated based on value of Δ_v	53
4.5	NPSS vs SPIKE-distance with colour coded jitter σ_{in} . Subfigures separated based on relation of V_{∞} and Δ_v to threshold.	54
4.6	NPSS vs SPIKE-distance for 6 cases across the entire range of synchrony parameters.	57
5.1	OU versus LIF comparison samples.	67
5.2	μ estimation procedure.	70
5.3	Frequency estimation results.	72
5.4	μ_0 and μ_a values where frequency estimation failed.	73
5.5	Firing ISI CV vs frequency estimation error and conditional CV distribution of failed frequency estimations.	74
5.6	μ estimation example: initial estimated values.	75
5.7	μ estimation example: period alignment.	76
5.8	μ estimation example: binned averages.	77
5.9	Estimated vs actual μ parameters.	78
5.10	Estimated vs actual σ parameters.	79
5.11	μ_p estimation error as a function of frequency estimation error.	80

6.1	Schematic of model and delay learning scenarios.	86
6.2	Synaptic locations on reconstructed model.	88
6.3	Synaptic weights after scaling.	90
6.4	Depolarisation peak voltage and peak time caused by EPSP at the soma.	90
6.5	NMDA conductance integrals.	92
6.6	NMDA conductance learning rule.	93
6.7	Firing domain: Range of delays that cause somatic firing, for a given background current amplitude.	95
6.8	Firing domain for background current amplitude of 0.222 nA.	96
6.9	AP firing delay for range of input delays.	98
6.10	Distribution of membrane fluctuations and firing probability of a neuron across a range of background current amplitudes.	102

Achilleas Koutsou

List of Tables

3.1	LIF and LIFwPR parameter symbols and value ranges.	27
4.1	Parameter symbols and value ranges	47
5.1	Parameter symbols and value ranges.	63
5.2	Input parameter estimation errors.	72
6.1	Normalised NMDA integral values at each ideal delay.	93
6.2	Model data and results.	99

Achilleas Koutsou

Abbreviations

STA	Spike Triggered Average
AP	Action Potential
CV	Coefficient of Variation
EPSP	Excitatory Post-Synaptic Potential
LV	Local Variation
ISI	Inter-Spike Interval
V-P	Victor-Purpura
LIF	Leaky Integrate and Fire
LIFwPR	Leaky Integrate and Fire with Partial Reset
PIF	Perfect Integrate and Fire
NPSS	Normalised Pre-Spike Slope
MAT	Multi-timescale Adaptive Threshold
LIFwVT	Leaky Integrate and Fire with Variable Threshold
OU	Ornstein Uhlenbeck
RMS	Root Mean Square
PSD	Power Spectral Density
AMPA	a-amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid
NMDA	N-methyl-D-aspartic acid
EPSC	Excitatory Post-Synaptic Current
IT	Information Technology

Symbols

Symbol	Name	Unit (base)
V	membrane potential	V
V_{rest}	resting potential	V
V_{reset}	reset potential	V
V_{th}	firing threshold	V
t_r	refractory period	s
τ_m	membrane leak time constant	s
β	reset parameter	—
N_{in}	number of input spike trains	—
f_{in}	input frequency (firing rate of input spike trains)	s
f_{out}	output frequency (firing rate of simulated neuron)	s
ΔV_s	spike weight — depolarisation caused by input spike	V
N_{th}	firing threshold in number of spikes	—
T	duration (simulation or experiment)	s
h	simulation time step	s
S_{in}	input synchrony degree	—
σ_{in}	input jitter	s
R	resistance	Ω
C	membrane capacitance	F
I	input current	A
I_V	input voltage (current times resistance)	V
w	coincidence window	s
m_i	measured pre-spike slope	V/s
U_i	slope measure upper bound	V/s

L_i	slope measure lower bound	V/s
M_i	normalised pre-spike slope (NPSS)	—
V_∞	asymptotic threshold-free membrane potential	V
Δ_v	volley peak potential	V
D_S	SPIKE-distance	—
$t_P^{(n)}$	time of the last spike before time t for spike train n	s
$t_F^{(n)}$	time of the first spike after time t for spike train n	s
$X_P^{(n)}$	time difference between t and $t_P^{(n)}$	s
$X_F^{(n)}$	time difference between t and $t_F^{(n)}$	s
$X_{ISI}^{(n)}$	duration of the inter-spike interval around t	s
$\Delta t_P^{(n)}$	time difference between $t_P^{(n)}$ and nearest spike in other spike trains	s
$\Delta t_F^{(n)}$	time difference between $t_F^{(n)}$ and nearest spike in other spike trains	s
μ	mean of stochastic input (drift)	V/s
μ_p	peak of sinusoidal drift	V/s
μ_a	amplitude of sinusoidal drift	V/s
μ_0	offset (baseline) of sinusoidal drift	V/s
σ	Variance of stochastic input (noise)	V/ \sqrt{s}
σ_p	peak of sinusoidal noise	V/ \sqrt{s}
σ_a	amplitude of sinusoidal noise	V/ \sqrt{s}
σ_0	offset (baseline) of sinusoidal noise	V/ \sqrt{s}
t_S	time of synchronous volley	s
N_P	number of Poisson input spike trains	—
N_S	number of synchronous input spike trains	—
ΔV_P	spike weight for Poisson inputs	V
ΔV_S	spike weight for synchronous inputs	V
D_V	maximum deviation (difference) between membrane potential traces	V
R_V	root mean square difference between membrane potential traces	V
Δt	input spike pair delay	s
g_{Na}	maximum sodium conductance	pS μm^{-2}
g_{Kv}	maximum voltage gated potassium conductance	pS μm^{-2}
g_{Ka}	maximum potassium conductance	pS μm^{-2}
g_{Km}	maximum muscarinic potassium conductance	pS μm^{-2}
g_{CaT}	maximum T-type calcium conductance	pS μm^{-2}

g_{Kca}	maximum calcium dependent potassium conductance	$\text{pS } \mu\text{m}^{-2}$
g_{Ca}	maximum calcium conductance	$\text{pS } \mu\text{m}^{-2}$
τ_{rise}	EPSC rise time	s
τ_{decay}	EPSC decay time	s
T_p	time-to-peak (TTP) – delay between onset and EPSP peak	s
T_w	EPSP peak width	s
d_{A_iB}	difference between TTP of EPSPs from A_i and B	s
d_S	somatic spike delay	s
I_0	baseline input current (when current is noisy)	A

Chapter 1

Introduction

Computational neuroscience aims to discover and understand the computational properties of the brain, from the single neuron, to the entire nervous system, by using techniques from computer science and mathematics in combination with the knowledge gained from the experimental areas of neuroscience [1]. The two fields coexist in a mutually beneficial relationship, as better understanding of the computational and functional aspects of neurons and their networks, advances both our understanding of the brain (benefiting *neuroscience*) and additionally, inspires new techniques for problem solving and artificial intelligence (AI) (which benefits *computer science*).

Neuroscience, in the most general terms, is the study of the brain, its structural organisation and its function [2]. Traditionally, such an endeavour alludes to biological sciences and as such, neuroscience appeared initially as an interdisciplinary specialty of molecular and cell biology, which later evolved into a stand-alone field [3]. While biological disciplines form the experimental basis for the production of knowledge in the field, abstraction is key to understanding the high level processes of such a complex system. The complexity of any nervous system and in particular, the human nervous system has often required the application of methods and techniques beyond the scope of experimental data analysis [4]. In the same way that fields such as physics require a separation into experimental and theoretical branches, neuroscience developed a requirement for theoretical frameworks to model and explain experimental data, as well as extrapolate beyond them [1, 4–6].

Computer science, electronic engineering, mathematics and physics form the computational and theoretical approaches to neuroscience and are concerned with the theoretical aspects of

information processing by neural systems. Such studies often involve single neuron or network modelling, from detailed descriptions of action potential (AP) generation [7] to simplified, abstract and analytically tractable models [8–13] which allow for mathematical and computational analysis of large, complex networks of neuron models. Modelling studies of neural behaviour have, in recent years, expanded to the level of entire artificial brain simulations [14] in an attempt to closely simulate and analyse the mammalian brain.

A major goal of modelling projects and in fact, one of the fundamental issues of the entire computational neuroscience field, is the problem of understanding the neural code [4, 15]. The neural code involves the mechanisms, both biological and computational, that the brain uses to represent information and carry out high level functions. We aim to build upon existing knowledge concerning the processing of information in the nervous system and address one of the key issues of the neural code, which is the role of the temporal structure of spike trains. The aim of this project is to use single neuron models in order to develop methods for identifying the time scales of single neuron processing. In particular, we aim to measure the effects of temporally synchronous inputs on the behaviour of a neuron and develop methods for identifying the causes of firing, with respect to the temporal structure of the input signal or spike trains. Additionally, we aim to measure the temporal precision of AP generation and the fidelity with which a cell is able to distinguish spike pairs. This will provide us with a better understanding of the temporal aspect of neural processing, which is an important part of understanding the neural code in general. Our models should be simple enough to allow for analytical and computational study and complex enough to make meaningful predictions about the processing capabilities of real neurons. In order for our conclusions to be useful, our underlying assumptions must be biologically plausible and experimentally confirmable.

1.1 Computational neuroscience

Computational neuroscience relies on methods of mathematics and computer science to investigate and understand the information processing functions of the nervous system [5]. However, there are two approaches to this goal: one is that mathematical and computational models are used to simulate single neuron and neural network behaviour in detail, as well as analyse related data. The other approach is a conceptual one, in that the brain is viewed and analysed as an abstract, computing device, providing researchers with the ability to describe neural

function within existing or new theoretical frameworks which allow alternative analytical approaches [4, 6].

In essence, this distinction is a matter of abstraction. The first approach mentioned above, deals with descriptive, *low-level* modelling and statistical analysis. Models that fall into this category often attempt to describe every (known) detail that relates to the neural function being studied. Such approaches however risk making the model as hard to understand as the nervous system itself. On the other hand, the second approach attempts to abstract only the important underlying principles responsible for computation, providing a theoretical framework for *high-level* analysis. In such cases, the analytical tractability of the models provides scalability and allows for more conceptual problems to be studied, such as the computational properties and limitations of large neural networks [5, 16].

Terms such as *theoretical neurobiology* [17, 18] and *theoretical neuroscience* [4] are also used to refer to work which falls in the same scope of computational neuroscience, therefore these terms are often considered to be synonymous. The term *computational neuroscience* did not appear until 1987 [1, 5, 19], where a symposium was held in order to discuss and define the field and in essence, the term itself. However, in retrospect, a number of scientific achievements have been recognised as the historical roots of this field. More specifically, the earliest and one of the most widely used models of a single neuron is the integrate-and-fire (I&F) model, introduced by Lapicque [8] in 1907, which models a neuron as an electric circuit consisting of a parallel capacitor and resistor. The simplicity and analytical tractability of the model, coupled with its ability to predict firing frequencies when the neuron is under constant stimulation, are the main reasons for its continued popularity. Additionally, the fact that the mechanisms for generating APs in neurons were not known at the time, exemplifies the ability of abstract models to capture high level operation, without requiring a full understanding of low level (e.g., molecular) mechanisms [8, 12].

A much greater understanding of neural operation and more specifically, AP generation came from the work of Hodgkin and Huxley [7] which described the generation and trajectory of an AP using a model consisting of four ordinary differential equations, describing the change in conductance of ion channels on the neuron's membrane [7]. This has been described as one of the most important breakthroughs in neuroscience, as it provided a framework for modelling work that continues to influence the field to this day [20].

The next historical breakthrough in the field of information processing in neural systems was the discovery and understanding of certain aspects of visual processing. In particular, Hubel and Wiesel [21] discovered that individual cells in the visual cortex respond to changes in the light intensity of specific areas of the visual field, called the neuron's *receptive field*. More importantly, they discovered that neurons that respond to motion of the same orientation, are organised in columns which span all layers of the cortex [21]. This discovery was a major indication that vertical columns of neurons in the cortex form discrete functional units, which hold neurons that share similar functionality and properties [4, 16].

Following these experimental discoveries, the field developed a theoretical foundation for the study of neurons' computational aspects [4, 17, 18, 22]. Researchers began integrating disciplines from mathematics and engineering such as dynamical systems [23, 24], information theory [15, 25–27] and stochastic processes [28–30].

As it is evident from the history of discoveries described above and the evolution of the field into a multidisciplinary area, the current state of research in the field takes advantage of both theoretical and experimental approaches, which should (ideally) complement each other. For example, the *Journal of Computational Neuroscience*¹ encourages the submission of combined experimental and theoretical work. It emphasises the requirement of biological relevance in theoretical work and inversely, relevance of computational function in experimental work. Similarly, other related journals also highlight their coverage of both theoretical and experimental work related to understanding information processing in the brain, such as *Biological Cybernetics*², which focuses on promoting the cooperation between life sciences and theoretical disciplines, and *Frontiers in Computational Neuroscience*³, which is primarily focused on modelling and theoretical studies but is open to experimental studies which relate to theoretical conclusions. Other publications which focus on the field and emphasise the integration of theoretical and experimental work, as well as the multidisciplinary nature of the field include *Network: Computation in Neural Systems*⁴ and *Neural Computation*⁵. *The Journal of Mathematical Neuroscience*⁶ focuses on using mathematics to understand the fundamental mechanisms responsible for experimentally observed behaviours in neuroscience at all scales. Therefore,

¹see <http://www.springer.com/biomed/neuroscience/journal/10827>

²subtitled "Advances in Computational Neuroscience"; <http://www.springer.com/biomed/neuroscience/journal/422>

³http://www.frontiersin.org/computational_neuroscience/about

⁴<http://informahealthcare.com/page/Description?journalCode=net&>

⁵<http://www.mitpressjournals.org/page/about/neco>

⁶<http://www.mathematical-neuroscience.com/>

although the field is primarily considered to have a theoretical focus, it also remains relevant to the experimental aspects of neuroscience.

The main contributions of this project focus on investigating the problem of the neural code by observing the behaviour of the membrane potential of a model neuron. This approach is based on several suggestions and observations made throughout the relevant literature. The membrane potential and its derivative provide information regarding synchronisation of firing [31, 32] and the trajectory of the membrane potential can be used to infer specific features of the input [33, 34]. It has been suggested that studying the reverse correlation of the membrane potential may be more informative than the spike triggered average (STA) stimulus for investigating neural operational modes [35]. In an intuitive sense, it can be argued that the trajectory of the intracellular, sub-threshold membrane potential in general should provide more information about the input, than the firing statistics (rate mean and variance) of the neuron. The former represents the neuron's internal state, which is directly affected by the input, while the latter is a discrete representation of super-threshold activity. More specifically, the neuron's internal state (the membrane potential) is determined by the neuron's physiology (cell type, ionic conductances, etc.) and pre-synaptic input. If we assume that the physiology of a neuron under study is known, it should be possible to infer specific properties of the input (e.g., synchrony), given the membrane potential time course of the neuron. In addition to the above, the thesis includes an investigation into the temporal precision of input spike pair discrimination in a detailed, biophysical compartmental model. This investigation aims to measure the smallest time difference between two spikes that are detectable by the neuron (i.e., that change a property of the response) and investigate ways in which the cell can *learn* to respond to specific time differences.

1.2 Thesis outline

Chapter 2 reviews the literature on the problem of the *neural code* (Section 2.1) and in particular the existing literature on the sub-problem of *neural operational modes* (Section 2.2) and *neural synchrony* (Section 2.3). Within the scope of neural synchrony, the development and use of *spike time distance* measures is also reviewed. The chapter concludes with a review of methods of *input reconstruction and inference*, which is relevant to our research into the causes of firing of neurons (Section 2.4).

Chapter 3 presents the development and application of a method for estimating the degree of synchrony driving a neuron using measurements of the membrane potential and in particular, the slope of the membrane potential shortly prior to threshold crossing. Through this estimation of the causes of firing, inferences can be made regarding the operational mode of the neuron under study. The chapter concludes with a discussion of the results, the implications of the measure regarding the problem of the neural code, as well as limitations and assumptions of the measure. The work presented in this chapter was the topic of a paper published in *Neural Computation* [36].

Chapter 4 describes a comparative study between the measure presented in Chapter 3 and a measure of *spike time distance*. Spike time distances measure the similarity (i.e., synchrony) between a pair or set of spike trains, while the measure presented in the previous chapter determines the degree of synchrony that is relevant to firing, which is an indicator of the operational mode. In this chapter, the synchrony between the input spike trains driving a neuron is calculated and the results are compared with the input synchrony estimated by the measure presented in the previous chapter. The comparison allows us to analyse how different properties of the neuron and its inputs shape the relationship between absolute input synchrony and operational mode. The chapter concludes with a discussion on the significance of the results and the value of the relationship between the measures. The work presented in this chapter is the topic of a paper submitted to a special issue on the *Application of mathematics in neuroscience* in the *Mathematical Biosciences and Engineering* and is currently under review.

Chapter 5 presents a method for inferring the input parameters of a simple neuron model driven by periodic input, simulating regular cycles of increased synchrony, by observing the firing behaviour of the cell. As with the previous measure which relied on the slope of the membrane potential to estimate input synchrony, this method aims to infer the degree of synchrony driving a neuron, without observing the input spike trains directly. The chapter concludes with a discussion on the accuracy of the results and a description of the limitations and assumptions of the method. The work presented in this chapter was the topic of a paper published in *Brain Research* [37].

Chapter 6 describes a case study on the use of coincidence detection and input correlations in synaptic plasticity and input selectivity. We simulated a reconstructed layer 5 pyramidal neuron model, driven by spike pairs of fixed inter-spike delay, which were regarded as single events from a pair of correlated spike trains. This setup allowed us to study the effects of

correlated inputs that temporally coincide on a specific part of the neuron's dendrite or soma, after propagating along the membrane. This chapter deviates from the previous work on identifying input synchrony and measuring the prevalence of one operational mode over another. The relevance of this chapter to the rest of the thesis is in the evaluation of a realistic model's ability to distinguish between coincident spike pairs with millisecond temporal differences. The chapter concludes with an analysis of the neuron's ability to distinguish temporally precise inter-spike delays, the implications of these limits on the temporal precision of the neural code, and the potential for learning such fine temporal differences. The work presented in this chapter was the topic of a paper which has been recently accepted to *BioSystems* subject to minor revisions [38].

The main part of the thesis concludes in Chapter 7, with an overview of the results, the conclusions and their implications for the production of new knowledge, with respect to the initial objectives of the thesis project.

Finally, Chapter 8 presents future extensions and adaptations of the presented methods, in order to make them more generally applicable and overcome limitations. It also discusses potential new avenues of research that could build upon the work described in the thesis.

Chapter 2

Literature review

2.1 The neural code

The problem of the neural code can be summarised by the following two questions, illustrated in Figure 2.1:

- (i) How is information (e.g., environmental stimulus, behaviour, thoughts) encoded into spike trains? This is the problem of neural *encoding*.
- (ii) How does the organism “read” spike trains in order to extract the contained information? This is the problem of neural *decoding*.

More specifically, the first question addresses the importance of understanding how the information processed by neurons is transformed into spike trains and propagated to the higher layers of the central nervous system (CNS) and more importantly, which aspects of the spike trains encode that information. By understanding this mechanism, one could predict the resulting neural activity given a specific stimulus, or generate artificial neural activity in order to represent a stimulus (e.g., for an artificial sensory prosthetic). The second question addresses the opposite process, i.e., inferring the stimulus or information from the resulting spike trains. This would allow us to read spike trains in order to fully understand not only the relevant external stimuli that caused them, but higher brain functions such as abstract thought, language and memory. These two problems are opposite sides of the same coin, since both result in a detailed understanding of the brain’s information processing mechanisms. In particular, both questions involve the discovery of the relationship between specific attributes of spike trains




FIGURE 2.1: Neural encoding illustrated by the images from left to right (a-d). (a) An object in the environment (b) is observed which causes (c) neurons in the brain to fire (d) a series of spikes, which represent the object, or a certain property of the object. Neural decoding is the problem of taking (d) the spike train and using knowledge of how (c) neurons operate to infer (b) the observation of (a) the objects.

(e.g., spike rate) and the corresponding attributes of the information they encode (e.g., light intensity).

The earliest and most commonly accepted proposed encoding mechanism for neurons is the rate code [39–44]. According to the rate coding model, neurons encode all the available information on the average rate of firing. This concept was introduced as early as 1926 [39, 40] when the discharge rate of motor and sensory neurons was found to be correlated with the force exerted by the relevant muscle (for motor neurons) and the intensity of the stimulus (for sensory neurons) [40, 41]. In many cases repeated representations of the same stimulus results in different patterns of spike trains being produced by the same neurons, while the average firing rates remain relatively constant [45–47]. This variability is often considered to be *noise*, i.e., random activity of neurons that is not associated with encoding. This noise may result from stochastic properties of spike generation, stochastic synaptic transmission or other sources of background randomness [10, 48]. A rate code is considered to be robust to such noise sources, since temporal jitter can alter the relative timing of individual spikes, but not the overall firing rate.

Contrary to the rate code model, there exists a substantial body of work which supports the hypothesis that variability in the discharge patterns of neurons is functionally significant, as opposed to being simply noise [49–52]. Theoretically, temporal codes which rely on precise

temporal spike patterns or multi-neuron correlations are much more efficient encoders of information, since in such cases, information can be encoded on various attributes of the spike train's timing and does not rely solely on the spike rate [53, 54]. Conversely, when spike trains are irregular and vary across repetitions of the same stimulus, the average rate of firing cannot be effectively extracted during any isolated trial [42, 55]. In other words, spike train analyses that rely on inter-trial averages tend to ignore the fact that the organism does not have access to such statistics and must make decisions based on a single trial's response.

Complementary to single neuron codes are population codes which take into account the combined activity of multiple cells [56–59]. Multi-neuron activity has been proposed as a plausible explanation to several issues surrounding the neural code. The *binding problem*, as it is often called, refers to the problem of relating specific perceived features in a stimulus environment to the corresponding objects. Coordinated multi-neuron activity can explain how feature binding is achieved [60, 61]. Additionally, studies have shown that groups of motor neurons use a voting strategy such that the resulting movement is a weighted sum of the activity of a neural ensemble in which each neuron has a preferred direction [56, 57].

Theoretical analysis of the neural code and spike trains in particular relies heavily on the fact that the firing of a spike is a stereotyped all-or-none event which, as mentioned above, means that spike trains can be analysed as a series of indistinguishable events. This allows spike trains to be treated and analysed as stochastic point processes [28, 29] and all the related mathematical formalisms can be applied. For instance, the firing rate can be described as the first passage time of the membrane potential to the threshold. Spike trains can be compared in terms of their firing rate, inter-spike interval distributions and correlated event times.

Of notable interest is the approach of the spike-triggered average (STA) stimulus (also known as the reverse correlation between spike train and stimulus) [15, 26, 42, 62]. This is a relatively direct approach to the second question posed above, i.e., the problem of decoding. Generally, the STA is the average stimulus in all time windows of a fixed length preceding each spike in a trial (Figure 2.2). In other words, it is the average of all inputs that caused a response spike and can be interpreted as the *meaning* of a single spike, in terms of the sensory stimulus, or the stimulus current generated by pre-synaptic firing. Additionally, such an analysis can indicate the stimulus selectivity of a neuron, as any deviation of the STA from the overall average stimulus shows a preference of the neuron towards the specific stimulus [42].




FIGURE 2.2: Schematic representation of the calculation of the spike-triggered average (STA) stimulus. The waveform (top row) represents the time varying stimulus which is aligned with the response spikes (second row) fired by a neuron driven by that stimulus. The graph at the bottom of the figure is the average of all stimulus waveforms within a short time period prior to each response spike. These time periods are represented as dashed boxes on the stimulus, which are aligned with the fired response spikes. Image taken from Gerstner and Kistler [16].

An interesting approach to the problem of the neural code is analysing spike trains in terms of their information content using information theory [25]. In general, information theory can be used to determine the amount of information flowing through a communication channel. In the case of neural systems, the communication channel is a network of neurons. Such analysis relies on specific assumptions about the neural code which can be confirmed, or refuted, when theoretical results accurately predict, or support experimental behaviour. For instance, in the motion-sensitive neuron of the fly, time-dependent input signals produce variable spike trains which can carry twice the information of a constant rate, noisy response to constant stimuli [46]. Deco and Schürmann [63, 64] used a simple neuron model to generate spike trains in response to various signals. The study focused on demonstrating how information theory can be applied to analyse the amount of spikes required to distinguish between signal inputs of varying similarity. Such a method can be used to measure the time-scales of processing in neurons, by estimating the time required to process and convey enough information to distinguish between different stimuli. With respect to the problem of the neural code, in certain studies it has been noted that information theory does not confirm the existence of a temporal code [27, 65] while in others, rate codes proved insufficient in carrying the information contained in the known stimulus [52, 66]. However, it should be noted that estimates of information depend greatly on the experimenter's choice of stimuli and selecting the appropriate stimulus set is not trivial [27, 67, 68]. Regardless of the difficulties, information theory has proven to be a very useful means of investigating the information processing capabilities and limitations of biological neural systems and conversely, its application to neuroscience has helped information theory evolve by providing it with new directions of research [69].

2.2 Neural operational modes

Coincidence detection & temporal integration

It is evident from Section 2.1 that there is a great debate within the computational neuroscience community concerning what is perhaps the most fundamental aspect of the field: the mechanisms used by neurons to encode information into spike trains. In 1982 it was proposed that neurons in the cortex may act as coincidence detectors rather than temporal integrators [70]. *Temporal integration* implies that neurons operate as integrate-and-fire devices, firing a response after a certain number of pre-synaptic spikes (on average) arrive at the neuron's input [43, 71]. Under this assumption, the temporal structure of pre-synaptic spike trains is lost in the integration process of the post-synaptic neuron. The firing rate of the post-synaptic neuron encodes the intensity (i.e., the rate) of the joint pre-synaptic activity. On the other hand, *coincidence detection* implies that the firing of a neuron is a result of synchronous arrival of multiple input spikes and that temporally dispersed spike trains are unable (or less likely) to cause a neuron to fire [72, 73]. While it is relatively straightforward to show that a passive model neuron can be excited by both temporally dispersed and precisely coincident activity [31, 74, 75], the existence of the latter has great implications for the nature of the neural code [35, 72, 73, 76]. In particular, if coincidence detection was indeed found to be the dominant mode of operation in certain neurons, it would indicate an increased significance of precise timing in spike trains, which greatly reinforces and supports the existence of temporal codes, as opposed to rate codes [72, 73, 76].

The debate surrounding the matter of the two encoding schemes, coincidence detection and temporal integration, was fuelled in 1992 and 1993 when William Softky and Christof Koch published an analysis of cortical cell recordings which showed that the particular neurons fire highly irregular spike trains in response to constant visual stimuli [72, 73]. While the irregularity found in spike trains was not an original discovery on its own [39–42], the (near Poissonian) high levels of variability, indicated by a coefficient of variation (CV) between 0.5 and 1.0, at rates up to 300 Hz, contradicted the notion that cortical cells operate as integrate-and-fire devices. More specifically, an integrator firing at such high rates, produces much more regular spike trains, a result that is evident by simulations of leaky integrate-and-fire (LIF) model neurons with biologically realistic parameter values (and in particular, a membrane leak time constant within the range of realistic values, 10 to 50 ms) [35, 73, 77].

From this analysis it follows that either cortical cells are selective of synchronous (i.e., coincident) activity at their input [73], or that a number of the underlying assumptions about neural operation and connectivity are inaccurate, or both. Certain explanations for the observed irregularity focused on the single neuron and in particular, weak re-polarisation (i.e., partial reset) [35, 77, 78] and variable threshold [79, 80], while others considered strong inhibition in cortical cells [43, 71, 81, 82], network effects [83–87] or the temporal structure of excitatory input [88, 89].

Additionally, the discovery of such highly variable spiking at high rates inspired a number of publications regarding methods of measuring the level of variability [see 90, for a review]. The reason for this arose from the realisation that the classical measurement, the coefficient of variation (CV), is very sensitive to rate changes and may overestimate the variability in cases where the underlying rate fluctuates [47, 48, 90, 91]. Initially, Holt et al. [48] proposed the localised coefficient of variation (CV_2), which measures the variability between adjacent interspike intervals (ISIs). This makes the metric robust to firing rate changes in long recordings. Other localised measures of variability, which also rely on the variability between adjacent ISIs are the local variation (LV) [92], the IR [93] and the spiking irregularity measure SI [94]. In particular, the LV metric was developed in order to distinguish between multiple motor neurons based on their spiking irregularity, measured by the metric. The LV was shown to be a reliable identifier of individual neurons [92]. The IR metric was developed for a similar purpose as the LV and was also successful in identifying specific types of neurons based on their spiking irregularity [93]. These authors also mentioned that the IR overestimates the irregularity of a spike train when there are very fast rate variations, but is robust to slower rate fluctuations [93]. The SI metric is used under the assumption that the intervals of a spike train follow a gamma distribution. The metric is used to estimate the shape parameter of the distribution, which corresponds to the spike train's irregularity [94]. In a review article where these local measures of irregularity were compared by measuring the irregularity of a group of spike trains with known fluctuating rates, the most reliable metric was shown to be the CV_2 [90].

Christodoulou and Bugmann [77] compared some of the proposed solutions to the problem of high irregularity at high rates. They criticised the lack of proof of exact Poisson type firing of other studies [43, 78, 81, 82, 89, 95, 96] and proved that only the partial somatic reset mechanism [35] reproduced the experimental firing statistics adequately. However, in their critique, these authors made some strong assumptions about the experimental data and the

related statistics [77]. The authors tested alternative mechanisms and assumptions based on their ability to fit Poisson-like statistics in three ways: (i) CV vs mean inter-spike interval (ISI) curve that fits the corresponding theoretical curve for a Poisson spike train with refractory period, (ii) exponentially distributed ISIs, and (iii) independence between ISIs (shown using an autocorrelation test). While these tests are indeed required to prove that a stochastic point process is of Poisson type, the original experimental data did not meet these statistical requirements either [72, 73]. In particular, Christodoulou and Bugmann [77] rejected the idea of strong inhibition being an explanation of highly variable firing at high rates [43, 71, 97] based on tests (i) and (iii) above. However, Softky and Koch [73] showed only that experimental data is “approximately consistent with a description of spiking as a Poisson process” (i.e., the experimental data is close to the theoretical CV vs mean ISI curve for a Poisson process with refractory period) and did not attempt to show independence of ISIs [72, 73]. Furthermore, it has often been noted that cortical spike trains are not completely described by a Poisson process [98–102]. Therefore, a hypothesis for the causes of highly variable firing at high rates cannot be rejected for not fitting Poisson statistics *exactly*.

During the discussion regarding the causes of high firing variability at high rates, the importance of the existence of such high variability received much less attention. Discovering the mechanisms of simple single neuron models that are able to reproduce experimental statistics became more important than the function of the phenomenon itself. While it is clearly desirable to discover the mechanism underlying the behaviour, since it could lead to a better understanding of its function, the implications of its existence were given less attention. Additionally, it should be reiterated that Softky and Koch [72, 73] discovered that the variability was only observable *in vivo*, while *in vitro* stimulation always produced regular spiking. It is curious then how a great proportion of the studies of this phenomenon focused on single neuron mechanisms [35, 77–80] and its input [43, 71, 81, 82], while network effects were considered less frequently [83–86]. Correlated, even synchronous, activity has been observed throughout the nervous system [50, 103–110] while Softky and Koch [73] assumed that the input to cortical cells is always of Poisson type and that high rate small excitatory post-synaptic potentials (EPSPs) create a near DC current input to cortical neurons.

A very interesting view on the matter of neural spiking activity is that the two operational modes, coincidence detection and temporal integration, are opposite extremes on the same continuum [31, 51, 74, 75, 111]. This notion supports that temporal integration occurs when a neuron receives temporally dispersed input spike trains and coincidence detection occurs

when input spikes are highly synchronised. The ability of neurons to spike in response to both synchronous and asynchronous activity suggests that the underlying operational mode is not dependent on an attribute of the neuron itself (such as selectivity to coincidences), but rather it is a result of the level of synchrony at the input. This can be understood by considering the quantitative differences between a rate (or spike count) code and a precise temporal code. While the former relies on averaging spike rates across relatively long periods of activity, or counting the number of spikes in long temporal windows, the latter relies on the same process but with a temporal window being only wide enough to contain a single spike, effectively increasing the importance of the precise time of firing of each individual spike (Figure 2.3).

The implication of this point of view is that it is counterproductive to speak of temporal integration and coincidence detection and by extension, rate and temporal codes, as clearly distinguishable alternatives. While there is a clear conceptual distinction between the two modes, it is important to keep in mind that between them lies a continuum of encodings based on the temporal precision and accuracy of spike generation. Therefore, it is more useful to speak of temporal precision in terms of the time scales required to distinguish responses to different external stimuli [112–116] (see Section 2.3 subsection on *Spike train distances*).

It has been identified experimentally and subsequently studied theoretically, that in the auditory system, the temporal precision is of the order of microseconds [117–122]. This means that neurons in the auditory system operate primarily as coincidence detectors, with a high sensitivity to precisely timed synchronous arrival of input spikes. For neurons with membrane time constants between 10 to 20 ms, as is found in the visual cortex, the temporal precision is considered to be between 1 to 3 ms [117].

Additionally, it is useful to define the neural modes of operation and encoding schemes in terms of the temporal characteristics of the input [31, 74, 75, 123]. Of particular interest is the concept of *pulse packets* [74, 124, 125], which provides a way to formalise the degree of synchrony at the input of a neuron and directly affects the resulting operational mode. More specifically, a pulse packet describes the combined activity of a number of neurons in terms of two parameters: (i) the number of neurons that fired and contributed to the packet (i.e., the total number of spikes in the packet) and (ii) the temporal dispersion of the spiking activity. Using this formalism, a low temporal dispersion (i.e., high synchrony) results in a high reliability of response [31, 75], which means that highly synchronised spike trains are more likely to propagate their precise temporal structure to the next layer. Therefore, highly synchronised




FIGURE 2.3: Two spike trains (A and B) with the same number of spikes (9 in total) may be considered identical or different depending on the size and positions of the bins of a *binning* procedure. The first bin size is equal to the length of the spike trains, which makes the spike trains appear identical. With smaller bin sizes, the difference between the two spike trains becomes more apparent until the bin size is only large enough to contain a single spike, in which case the spike trains can be represented by binary strings, where 1 represents the occurrence of a spike and 0 the absence.

input which defines coincidence detection, raises the significance and reliability of the temporal precision of spike trains. In other words, the existence of synchrony in the input suggests temporal codes.

2.3 Synchrony

The importance of synchrony in neural processing is a heavily studied topic. Synchronous activity can reliably propagate through neural layers [123, 126–128], which suggests that information may be encoded on more precise time scales than the general rate codes. Recently, the concept of *synchrony receptive fields* was introduced theoretically [129], which is defined as a set of stimuli that cause a pair (or group) of neurons to fire synchronously. Such a mechanism is closely related to theories of the role of synchrony in neural binding [60, 61]. Detecting convergent synchronous activity (i.e., synchronous firing with a common, post-synaptic cell or population) and measuring it can be very useful in discovering the prevalence and effects of such a phenomenon.

A great number of experimental studies confirm the existence of synchronous activity in various parts of the brain, such as the olfactory bulb [106], hippocampus [104], somatosensory cortex [105], auditory cortex [107], visual cortex [50, 109], and retina [103, 108, 110]. As we argued at the end of Section 2.2, synchronous activity between neurons in a network causes downstream neurons to behave as coincidence detectors, assuming the downstream neurons are driven by the synchronous activity directly, which provides strong evidence of temporal coding mechanisms. Therefore, any search for spike synchrony in a network of neurons is synonymous with a search for temporal encoding schemes and the temporal precision of the encoding performed by any given neuron can be expressed in terms of its input spike train synchrony and vice versa.

Identifying synchrony and more importantly, quantifying the levels of synchrony in neural circuits is not a trivial task however [130–135]. Given any pair or group of active neurons, there is always an amount of observable coincident spiking due to random chance. In other words, even under the assumption that inter-neuron firing times are independent (i.e., no correlated activity), there is always an expected amount of synchronous activity. This expected

level of synchrony naturally increases as the firing rates of the neurons under observation increase [134]. High amounts of observable synchrony between large ensembles of neurons with high activity is therefore not necessarily of any significance.

However, there is a point to be made on the subtle difference between a neural ensemble spiking in synchrony and coincidence detection at the single neuron level. In particular, the former implies the latter if the synchronous activity, or at least part of it, converges onto a neuron which can cause a repetition of spike patterns in subsequent layers, as in the so called *synfire chain* [70, 74, 123, 130, 136]. A synfire chain defines a specific network structure. It is a feedforward network of layered neurons, where each neuron in a given layer connects only to neurons in the next. Activity in synfire chains can propagate with tight synchrony throughout subsequent layers, depending on the size of the layer and the connectivity of the network [74, 130]. If precise spike patterns are used to encode information, the information is preserved by the precise repetition of patterns in subsequent layers. As with synchronous activity in general, identifying synfire chain activity embedded in large networks still poses a challenge. The existence of repeating patterns of activity does help to support the arguments for the existence of a feed-forward, strongly connected structure, but even if such activity is significantly more than expected by chance, it does not prove the existence of the particular connectivity associated with synfire chains [130, 137].

As mentioned above, synfire chain activity is hard to detect and even harder to prove [130, 137]. Precise temporal spike trains are not alone sufficient to prove coincidence detection or a temporal code. In order to address the matter of precise spike train propagation and successfully identify coincidence detection, the detailed interplay between input spike train statistics, somatic integration and spike generation mechanisms need to be studied at the intracellular level.

Spike train distance

Measures of spike train distance are a family of methods for calculating the dissimilarity (distance) between two or more spike trains [113, 114, 138–144]. Such measurements for spike trains are useful for quantifying the differences and similarities of spike timings across a number of trials with common stimuli [114], or measuring the spike time synchrony within a neural sub-population [139, 140, 142, 143]. The *distance* measured by these methods is equivalent to the inverse of the synchrony of the spike times. In most cases, these measures operate on two spike trains, calculating the difference between the times of individual spikes between the pair.

They evaluate to 0 when a pair of spike trains is identical and the value increases as the spike times in each spike train diverge.

The Victor-Purpura spike train distance (V-P), is a measure of the dissimilarity between two spike trains expressed as the cost of converting one spike train to another, in a similar fashion to the more general *edit-length distance* for sequences of characters or symbols [141]. The measure assigns a cost of 1 for adding or removing a spike and a parameter cost q (with units s^{-1}) for shifting a spike along the time axis. The algorithm computes the minimum cost for transforming one spike train into another. Removal and addition of a spike is less costly than shifting, when the time difference between two spikes (Δt) is greater than $2\Delta tq$. The parameter q therefore controls the temporal precision assumed by the measure. For $q = 0s^{-1}$, the V-P distance becomes a simple spike count measure, since shifting spikes is free and therefore the cost of converting a spike train into any other is defined simply by the cost of adding or removing spikes. As q increases, the measure becomes increasingly sensitive to precise spike timing.

The van Rossum distance has a similar nature to the V-P distance but uses a more straightforward algorithm to compute the distance [113]. The algorithm to compute the van Rossum distance between two spike trains is simply composed of a convolution of the spike trains (where a spike train is represented as a sequence or sum of Dirac delta functions, eqn. 2.1) with an exponential kernel with time constant τ_c (eqn. 2.2). The purpose of the kernel is to resemble the shape of a post-synaptic potential or current. The final distance between the spike trains is the integral of the difference between the two convolved functions (eqn. 2.3).

$$f^{orig}(t) = \sum_i^M \delta(t - t_i) \quad (2.1)$$

$$f(t) = \sum_i^M H(t - t_i) e^{-(t-t_i)/\tau_c} \quad (2.2)$$

$$D(f, g) = \frac{1}{\tau_c} \int_0^{\infty} [f(t) - g(t)]^2 dt \quad (2.3)$$

The parameter τ_c serves the same function as the q parameter of the V-P distance, though with inverse effect. In other words, for large values ($\tau_c \rightarrow \infty$) the measure computes the difference in spike counts, and for small values ($\tau_c \rightarrow 0$) the distance increases as spike trains differ in precise spike timing.

The SPIKE-distance, a time resolved and time-scale independent metric, relies on instantaneous differences between nearby spikes from two spike trains and constructs a temporal profile of differences between the spike sequences [142]. Single valued differences between a pair of spike trains are calculated as the integral of the temporal profile. The SPIKE-distance is a parameter-free measure, which is advantageous when no assumptions can be made about the temporal precision of the spike trains being analysed.

More recently, Rusu and Florian [144] introduced the modulus-metric, a spike train metric inspired by the *Pompeiu-Hausdorff distance* [145, 146] between two nonempty, compact sets. The modulus metric computes the distances between two spike trains using the difference between the time-resolved distances between all times t and the closest spike in each spike train for the duration of the spike trains [144]. The behaviour of the modulus-metric is similar to the SPIKE-distance since they are both parameter-free and both construct a temporal profile of spike coincidences and distances, the integral of which is computed as a spike-train distance.

When analysing multiple spike trains, the aforementioned measures can be used to calculate the average pairwise distance. Pairwise calculations however scale fast (on the order of N^2 , where N is the number of spike trains), making such calculations infeasible for even a few dozen spike trains. A variant of the SPIKE-distance for multiple spike trains, the *multivariate SPIKE-distance*, can be calculated using the standard deviation of spike times following or preceding a point in time, which scales linearly with N [142].

2.4 Input reconstruction and inference

A neuron's behaviour is generally thought to be a function of its intrinsic properties and the input that is driving it. In modelling studies, the intrinsic properties are represented by the neural parameters (e.g., membrane leak time constant), which often have direct biological interpretation. Input parameters depend on the modelling technique used for the input. Most commonly, inputs are modelled as either a time-varying input current or a sum of Dirac delta pulses [16]. Reconstructing a neuron's input involves observing the neuron's behaviour (either the spiking activity or the intracellular voltage), knowing the values of the intrinsic neural parameters and estimating the parameters or the waveform of the input.

The STA, mentioned in Section 2.1 can be used as part of an input reconstruction method. After calculating the STA waveform of a neuron based on the firing activity and the known input

signal, an estimate of unknown input signals can be constructed by convolving the spikes fired by the same neuron at a different time, with the calculated STA waveform [15, 16].

Other methods however, based on maximum likelihood estimation, do not assume any known input and use the membrane voltage to estimate input parameters [34, 147–149]. These studies use a stochastic diffusion model, based on the Ornstein-Uhlenbeck process [150] and estimate the drift (μ) and variance/noise (σ) parameters of the stochastic process that models the input signal.

Lansky [147] developed a maximum likelihood estimator for the Ornstein-Uhlenbeck model, assuming constant values for the parameters. Kobayashi et al. [34] developed a similar method for estimating the time-varying input parameters ($\mu(t)$ and $\sigma(t)$) and achieved high accuracy, as long as the variation of the parameters was slow. Bibbona et al. [149] measured the bias of input parameter estimators when applied to single samples of observed trajectories (i.e., a single inter-spike interval) of the membrane potential. This single-shot estimation procedure assumes that the values of the parameters, μ and σ are constant for the duration of the ISI. Iolov et al. [151] assumed a sinusoidal input signal and described a method for estimating the parameters of the input by binning and averaging estimates at appropriate time shifts, such that intervals that shared similar parameter values were averaged together.

Generally, input estimation methods require simplifying assumptions for the model and the nature of the input. The LIF or one of its variants is used as a model of neural behaviour, which allows for analytical treatment. Input estimation, i.e., reversing the integration process performed by the model, becomes feasible due to the passive integration performed by the model, especially when intrinsic neural parameters are known. Conversely, such estimations are less successful on more complex models, where active conductances and the interplay between various ionic channels create nonlinear interactions between input signals and membrane trajectories. Despite the simplicity of the LIF and its variants, it is widely regarded as a useful tool for analysing neural behaviour and input estimation methods can help us develop methods for decoding the stimulus, after observing the resulting behaviour of the internal state of the neuron.

Chapter 3

Determining the causes of firing using the membrane potential slope

One of the general aims of this thesis is to develop methods for analysing single neuron activity, with the specific goal of discovering the temporal precision of the neural code. More specifically, it aims at distinguishing between the two operational modes, temporal integration and coincidence detection, as they were presented in Section 2.2. This follows from the work of Bugmann et al. [35], who used the STA stimulus of a simple model neuron, the LIF with partial somatic reset (LIFwPR), in order to determine the dominant operational mode. After failing to provide conclusive evidence towards either one of the proposed mechanisms, these authors suggested using reverse correlations of the spike train and membrane potential of the specific neuron, which is where this thesis begins.

The realisation that the operational mode of a single neuron is a function of the temporal statistics of its input, rather than an intrinsic neural mechanism (see Section 2.2), and the suggestion that the time course of the membrane potential may hold evidence towards determining the operational mode of the single neuron, drove the project towards looking at evidence of input parameters projected onto the internal state of the neuron itself and particularly, the intracellular membrane potential. Previous studies have investigated the correlation between input synchrony and membrane potential slope [31]. A similar, analytical study also identified the dependence of a neuron's firing rate on the derivative of the membrane potential as a mechanism for inter-neuron synchronisation [32]. We used this knowledge to identify evidence of

high levels of input synchrony at a neuron's input, by looking at the shape and rise time of the membrane potential prior to the firing of a spike.

This quantification required defining the two operational modes in terms of input synchrony and in effect, in terms of the membrane potential slope. This way, the slope of the membrane potential resulting from completely synchronous (denoting coincidence detection) and dispersed input (denoting temporal integration) could be determined. Any subsequent pre-spike membrane potential slope, resulting from partially synchronous activity, could be normalised and expressed as a value within a bounded range, where 0 denotes temporal integration and 1 denotes coincidence detection.

3.1 Neuron models

Despite the criticism described in Section 2.2 concerning the possibly incorrect assumptions of the Poisson-like nature of neural firing, we still consider the partial reset variant of the LIF a very accurate spike generator model, because it fits accurately the experimental CV vs mean firing ISI curve for high firing rates [35, 73, 77]. Therefore, the method was to be used to quantify the relative contribution of each operational mode to the firing of highly irregular spikes at high rates, as analysed and presented by Softky and Koch [72, 73, see also Section 2.2 of this thesis]. The standard LIF (with total reset) is used as a baseline model for developing and generalising the methodology and for further investigations into the operational mode and the effects of pre-synaptic synchrony.

3.2 LIF and LIFwPR neuron models

The LIF model is described by the following differential equation:

$$\tau_m \frac{dV}{dt} = -V(t) + V_{rest} + RI(t) \quad (3.1)$$

where τ_m is the membrane leak time constant, V_{rest} is the resting potential, R is the resistance and $I(t)$ is the time-dependent input. The input is usually modelled either as a time-varying function or a sum of spike trains that cause instantaneous jumps. In the former case, $I(t)$ is modelled as an input current and in the latter, R is removed and $I(t)$ is composed of a sum of

N_{in} pre-synaptic spike trains, which fire spikes ($I_n(t)$) that cause instantaneous jumps of size ΔV_s in the membrane potential (eqn. 3.2).

$$I(t) = \sum_{n=1}^{N_{in}} \Delta V_s I_n(t) \quad (3.2)$$

In our work, we use the latter formalism of multiple super-imposed spike trains (see Section 3.3). AP firing is modelled explicitly by resetting the membrane to a fixed value V_{reset} when the potential $V(t)$ reaches a fixed threshold V_{th} . A reset parameter β is used to control the value of V_{reset} which also controls the model variant we use [35] (eqn. 3.3).

$$V_{reset} = \beta(V_{th} - V_{rest}) + V_{rest} \quad (3.3)$$

For $\beta = 0$, we have $V_{reset} = V_{rest}$ which makes the model behave as the standard (total reset) LIF. Setting $0 < \beta < 1$, causes $V_{rest} < V_{reset} < V_{th}$ which means that the membrane potential is set to a higher value than rest after a spike is fired. We call the second variant the *partial reset* model (LIFwPR) [35]. After reset, the integration of inputs continues, but the threshold (and firing mechanism) is disabled for a refractory period t_r . This method of simulating the refractory period, as well as the value for the parameter (2 ms), were taken from Bugmann et al. [35] for compatibility with their results and analysis regarding the highly irregular firing at high rates.

3.3 Input generation

Synaptic inputs were modelled in all cases as realisations of a Poisson process (i.e., intervals were exponentially distributed). The input population was characterised by five parameters, two of which relate to synchrony. The non-synchrony parameters are (i) the number of input spike trains (N_{in}), (ii) the average rate of the inputs (f_{in}) and (iii) the level of depolarisation caused by each spike on the membrane potential (ΔV_s , the synaptic *weight*). In addition, the two parameters which define the level of synchrony are (iv) S_{in} , which denotes the proportion of spike trains which are synchronous and (v) σ_{in} , which is the standard deviation of a normally distributed random variable that is used to apply Gaussian jitter to each individual spike in each of the identical spike trains. The parameter ranges were chosen such that our measure is investigated in the entire operational mode range of the LIF model. More specifically, spike trains are generated as follows:

1. Generate one Poisson spike train, with rate f_{in} for the length of the simulation T .
2. Copy the generated spike train $S_{in}N_{in} - 1$ times, giving a total of $S_{in}N_{in}$ identical spike trains (where $S_{in}N_{in}$ is rounded to the nearest integer).
3. For each spike in all spike trains generated so far, shift its time by a random variate drawn from a normal distribution $X \sim \mathcal{N}(0, \sigma_{in}^2)$.
4. Generate $(1 - S_{in})N_{in}$ Poisson spike trains, giving a total of N_{in} input spike trains.

Figure 3.1 shows four sample input cases. The raster plots show the effect of the two variables S_{in} and σ_{in} on the overall synchrony of the spike trains. During a simulation, each input spike causes an instantaneous jump of ΔV_s in the post-synaptic neuron's membrane potential.

The maximum value for σ_{in} of 4 ms (see Table 3.1 for ranges and values for all parameters) was chosen such that it is high enough to reduce synchrony significantly, even for cases where $S_{in} = 1$.

Note that, while various input parameter ranges were investigated for the LIF model, the parameters of the LIFwPR model are constant (Table 3.1). These values were taken from Bugmann et al. [35] who investigated the LIFwPR model and determined the parameter values which cause highly irregular high rate firing. Therefore, an investigation of the parameters of the LIFwPR model is not necessary, as we employ the specific model solely to investigate its sub-threshold membrane potential trajectories in the highly irregular high firing rate regime.

Our work in this chapter focused exclusively on excitatory inputs. This simplifying assumption allowed us to define more clearly the effects of synchronous activity on the membrane potential trajectory in a more predictable fashion. Even though it has been demonstrated [152] that increasing inhibition leads to greater membrane potential fluctuations in addition to reducing the mean membrane potential, the effects of inhibition on the slope of the membrane potential were not studied in this part of the thesis.

Table 3.1 gives a list of all the parameter values and ranges used for the LIF and LIFwPR simulations.




FIGURE 3.1: Four sample input cases showing the effects of the two synchrony parameters S_{in} & σ_{in} (see text for details) on the overall temporal structure of the input spike trains. For all three cases, $N_{in} = 50$ and $T = 1$ s. The first raster plot (a) shows a mostly random set of spike trains, with only 20% of the spike trains being completely synchronised ($S_{in} = 0.2$). The second plot (b) shows the effects of jitter ($\sigma_{in} = 3$ ms) on a set of spike trains with low synchrony ($S_{in} = 0.2$). The third plot (c) shows a much higher degree of synchrony with 80% of the spike trains being identical ($S_{in} = 0.8$). The fourth plot (d) shows the effects of high jitter ($\sigma_{in} = 3$ ms) on spike trains with 80% synchrony ($S_{in} = 0.8$). Comparing (c) to (d), while it is apparent by the vertical columns of aligned spikes that there is a high amount of synchrony in both, the existence of Gaussian jitter in (d) makes the overall spike trains more noisy and the columns are less pronounced.

3.4 Pre-spike membrane potential slope

The definition of the operational modes had to take into account the partial somatic reset mechanism, which only partially re-polarises the membrane potential after a spike is fired. Since this mechanism affects the internal state of the model neuron (i.e., the membrane potential), our method should account for arbitrary post-spike reset levels.

Our method should produce a normalised value within a bounded range. We designed the measure such that the range of values is $[0, 1]$. The value is calculated by normalising the slope of the membrane potential preceding each spike. The normalisation requires defining the bounds

Symbol	Description	Value for LIF	Value for LIFwPR	Units
V_{th}	Firing threshold	15	15	mV
V_{rest}	Neuron resting potential	0	0	mV
V_{reset}	Neuron reset potential	0	0	mV
R	Membrane resistance	10	10	k Ω
τ_m	Membrane leak time constant	10	10	mV
t_r	Refractory period	2	2	ms
ΔV_s	Synaptic weight	0.1 to 2.0	0.16	mV
N_{in}	Number of input spike trains	30 to 200	50	—
f_{in}	Input spike frequency	20 to 700	150 to 300	Hz
S_{in}	Input synchrony degree	0 to 1	0 to 1	—
σ_{in}	Gaussian jitter	0 to 4	0 to 4	ms

TABLE 3.1: Parameter symbols and value ranges for the models used throughout this chapter. Many of the parameters share a common value; we list them explicitly for completeness. The ranges of the input rates were chosen accordingly to study the entire obtainable firing frequency range.

of the membrane potential slope associated with each operational mode, which in turn requires precise definitions of the operational modes themselves. We define the upper and lower slope bounds for the normalisation as the slopes that would result from highly synchronised and highly dispersed input spikes respectively. This follows from previous work which indicates that precisely synchronised input results in the neuron operating as a coincidence detector, while completely dispersed input is associated with temporal integration [31, 74, 75, 111].

The following equations are used to calculate the bounds of the membrane potential slope associated with completely synchronised (eqn. 3.4) and completely dispersed (eqn. 3.5) input spike trains respectively:

$$U_i = [V_{th} - (V_{rest} + (V_{reset} - V_{rest})E)] w^{-1} \quad (3.4)$$

$$L_i = [V_{th} - (V_{rest} + I_V (1 - E))] w^{-1} \quad (3.5)$$

$$E = e^{-\frac{\Delta t_i - w}{\tau_m}} \quad (3.6)$$

$$I_V = \frac{V_{th} - V_{reset}}{1 - e^{-\frac{\Delta t_i}{\tau_m}}} \quad (3.7)$$

where w is a method parameter called the *coincidence window*, Δt_i is the duration of the ISI preceding the i^{th} output spike ($\Delta t_i = t_i - t_{i-1}$), and I_V is the constant input (expressed as a voltage – current times resistance is implied) required to fire at the end of the ISI, starting from V_{rest} . E is a common term that is used to calculate the effect of the membrane leak from

the start of an ISI up to the start of the coincidence window, $\Delta t_i - w$.

The coincidence window w defines the maximum temporal distance between a pair of spikes that are considered to be coincident. Smaller values of w increase the temporal precision of the assumed neural code. For instance, for $w = 5$ ms, any number of spikes which arrive within 5 ms prior to the firing of a spike are considered to have arrived synchronously. This will be reflected in the slope value associated with coincidence detection, U_i (eqn. 3.4), which will regard any membrane potential changes within 5 ms prior to a spike as instantaneous. The value of the membrane potential slope associated with temporal integration, L_i (eqn. 3.5), assumes that under conditions where there is no synchronous activity, the membrane potential of the neuron rises almost steadily from the reset potential (V_{reset}) to spike threshold (V_{th}), during the entire duration of the inter-spike interval (Δt_i).

The requirement for a coincidence window parameter is twofold. First, since inputs to the LIF model are modelled as instantaneous jumps, the trajectory of the membrane potential is discontinuous. The coincidence window allows us to calculate the rate of change of the membrane potential from the start of the window to the end, ignoring the discontinuity of the trajectory. The second reason is somewhat conceptual. Coinciding or synchronous events almost never occur simultaneously (to an arbitrarily high temporal precision), therefore the precision of temporal coincidence is different for each case. For instance, synchrony can be defined as two events within at most 2 ms of each other if time-scales are small and processing is considered to be fast. In other cases this temporal difference can be larger, e.g., 5 ms or even 10 ms. The w parameter allows us to measure synchrony using different definitions of temporal precision. The choice of value depends on the temporal precision of a given neuron, which is usually considered to depend on the leak time constant (shorter time constants make a neuron more sensitive to coincidences which implies a higher precision). As mentioned in Section 2.2, the temporal precision for neurons with membrane time constants between 10 to 20 ms is considered to be between 1 to 3 ms [117]. We use a membrane leak time constant of 10 ms, as our study originated from data and models of neurons of the visual system [35, 73] and therefore we use a temporal window of length 2 ms.

The bounds are used to linearly normalise the slope of the measured pre-spike membrane potential. Using the coincidence window defined for the bound definitions w , we calculate the average rate of change of the membrane potential between t_i and $t_i - w$, which is the slope of




FIGURE 3.2: Example membrane potential trace $V(t)$ for the LIFwPR model (blue solid line). The pre-spike windows (w) are shown as red rectangles. The related secant lines are shown as dashed blue lines starting from $w = 2$ ms before each spike ($t_i - w$) and ending at the threshold (V_{th} , represented by the dashed black line) at the time of each spike (t_i). The two grey curves represent the upper and lower bounds. The lower bound (low slope) corresponds to the curve starting at the post-spike reset potential (in this case, $V_{reset} = 13.65$ mV) and ending at the point where the potential crossed the threshold ($V_{th} = 15$ mV). The upper bound (high slope) corresponds to the curve which decays from the post-spike reset potential for the duration of the ISI. The two bounds correspond to the theoretical trajectory of the membrane in the presence of constant input for the lower bound and completely synchronised inputs, with no background activity, for the upper bound.

the secant line that intersects the membrane potential curve at the start and end of the window w (Figure 3.2), as shown in eqn. 3.8.

The slope of the secant line of the membrane potential prior to each spike is calculated as the difference of membrane potentials between the start and end of the coincidence window, over the duration of the window (Figure 3.2, eqn. 3.8):

$$m_i = \frac{V(t_i) - V(t_i - w)}{w}, \quad (3.8)$$

where $V(t)$ denotes the membrane potential at an arbitrary time t . Note that $V(t_i)$ denotes the membrane potential at the time of firing of the i -th spike, in other words, during threshold

crossing and therefore $V(t_i) = V_{th}$ always. The resulting value is subsequently normalised between U_i and L_i :

$$M_i = \frac{m_i - L_i}{U_i - L_i} \quad (3.9)$$

Finally, the mean M_i of an entire simulation is calculated to characterise the simulation's behaviour as a whole:

$$\bar{M} = \frac{1}{N_{sp}} \sum_{i=1}^{N_{sp}} M_i, \quad (3.10)$$

where N_{sp} is the total number of spikes fired in the given simulation run. Since the bound calculations involve the preceding ISI of each spike, the value of M for the first spike is calculated using the start of the simulation ($t_0 = 0$ ms), assuming the initial voltage is equal to the reset potential ($V(0) = V_{reset}$) and the input spike trains begin driving the neuron from the start of the simulation (stimulus onset is at t_0), both of which hold for our simulations. Alternatively, the calculation may consider only the $N_{sp} - 1$ intervals between all fired spikes.

We refer to the method described in this section, as well as the value it measures, as the *Normalised Pre-Spike Slope*, or NPSS for short.

3.5 Results for LIF

The six images in Figure 3.3 show results of simulations of the LIF neuron for various parameter combinations (see figure caption for details). The parameters were chosen to demonstrate how the NPSS behaves under sub- and supra-threshold stimulation. The plots show the mean NPSS of the membrane potential (\bar{M}) for all combinations of S_{in} and σ_{in} within the value ranges specified in Table 3.1. Each \bar{M} value represents the mean M for all spikes fired during $T = 10$ s of simulated time.

\bar{M} reaches the maximal value of 1 in the lower right hand corner, which corresponds to completely synchronised input spike trains ($S_{in} = 1$) with no jitter ($\sigma_{in} = 0$ ms). As expected, less synchronised spike trains (lower S_{in} values, i.e., moving left on the plot) decrease the value of \bar{M} . The correlation between S_{in} and \bar{M} , when $\sigma_{in} = 0$ ms, shows a near perfect positive linear relationship, with a correlation coefficient of $\rho_{S,M} = 0.99$ (clearly shown in Figure 3.4a). Similarly, increasing the amount of jitter (higher σ_{in} values, i.e., moving up on the image) also decreases the value of \bar{M} . The correlation coefficient between σ_{in} and \bar{M} , when $S_{in} = 1$, is




FIGURE 3.3: The normalised pre-spike membrane potential slope (NPSS, \bar{M}) for a LIF neuron model with total reset. For each plot, the firing rate of the LIF neuron is kept constant by calibrating the rate of the input spike trains at each data point. All the input spike trains were calibrated simultaneously and always shared the same mean rate. The parameters for

the plots were as follows:

- (a) $N_{in} = 100, \Delta V_s = 0.1 \text{ mV}, f_{out} = 5 \text{ Hz}$
- (b) $N_{in} = 50, \Delta V_s = 0.2 \text{ mV}, f_{out} = 100 \text{ Hz}$
- (c) $N_{in} = 60, \Delta V_s = 0.3 \text{ mV}, f_{out} = 10 \text{ Hz}$
- (d) $N_{in} = 60, \Delta V_s = 0.5 \text{ mV}, f_{out} = 70 \text{ Hz}$
- (e) $N_{in} = 200, \Delta V_s = 0.1 \text{ mV}, f_{out} = 10 \text{ Hz}$
- (f) $N_{in} = 60, \Delta V_s = 0.5 \text{ mV}, f_{out} = 400 \text{ Hz}$

The horizontal axis shows the proportion of synchronised spike trains ($S_{in} \in [0, 1]$), while the vertical axis shows the amount of jitter applied to the synchronous spikes ($\sigma_{in} \in [0, 4] \text{ ms}$). The colour indicates the value of \bar{M} for the simulation, as shown by the colour-bar on the right hand side.




FIGURE 3.4: Normalised pre-spike membrane potential slope (\bar{M}) for (a) $\sigma_{in} = 0$ plotted against the full range of S_{in} values and (b) $S_{in} = 1$ plotted against full range of σ_{in} . The dots and blue line represent measured data points for the LIF neuron firing at 70 Hz (corresponding to Figure 3.3d), while the black dashed line represents perfect linear correlation for comparison.

$\rho_{\sigma,M} = -0.95$ indicating a very high, negative linear relationship (Figure 3.4b). These values correspond to a desired $f_{out} = 70$ Hz (corresponding to Figure 3.3d).

Deviations of \bar{M} from the trend in the measurements of the NPSS (Figures 3.3 and 3.4) are due to randomness in the input spike trains, which cause random fluctuations of the membrane potential. This in turn causes pre-spike membrane potential slopes that do not follow the typical relationship between the slope and the parameters used to generate the synchronous inputs. However, the slope is still correlated with the degree of input synchrony; the randomness affects the correlation between the input synchrony parameters (S_{in} & σ_{in}) and the degree of synchrony in the generated input spike trains. Although the averaging is used to remove such variations (Eqn. 3.10), deviations may occur frequently and appear in the final results.

The results shown in Figures 3.3 and 3.4 as well as the aforementioned strong correlation between the input parameters and \bar{M} indicate that the NPSS can reliably detect and measure the input synchrony which was relevant to the firing of response spikes. The method maintained a high reliability for a wide range of the input parameter values, i.e., the number of spike trains (N_{in}), the desired firing rate (f_{out}) and the membrane potential rise per spike (ΔV_s). However, the robustness of this correlation depends on the input regime, i.e., the strength of the input volleys with respect to the firing threshold. In particular, this correlation between input synchrony and measured \bar{M} is robust as long as the synchronous volleys are supra-threshold, $N_{in}\Delta V_s \geq (V_{th} - V_{rest})$. The relationship between input synchrony and \bar{M} , in other words, the relationship between input synchrony and operational mode, and the way in which it is

affected by various properties of the neuron and the neuron's activity, is investigated in depth in Chapter 4.

3.5.1 Upper- and lower-bound convergence at high firing rates

Of particular interest is the case where the LIF neuron is driven by high rate inputs causing it to fire at extremely high rates. Figure 3.3f shows the measured synchrony of a LIF neuron with $f_{out} = 400$ Hz. Comparing this plot to the others of Figure 3.3, it is evident that the value of \bar{M} remains very high for all configurations.

This behaviour (Figure 3.3f) is due to the mean firing ISI approaching (in duration) to the coincidence window. More specifically, if an ISI is equal to the coincidence window $\Delta t_i = w$, then the values of the two bounds (eqns. 3.4 and 3.5) become equal. This occurs because in such circumstances, the firing of a spike due to integration of inputs within a period equal to the ISI is equivalent to firing solely from input spikes arriving within a period w . We can investigate the divergent behaviour between the two modes as a function of the ISI (Δt_i) and coincidence window (w).

In order to grasp the intuition behind this phenomenon, let us first consider the case where a firing ISI is equal to the coincidence window, i.e., $\Delta t_i = w$. In this case, it is clear from eqns. 3.4 and 3.5 that $U_i = L_i$, i.e., the two operational modes are described by the exact same slope value and are therefore identical and indistinguishable.

More generally, at very high firing rates the ISIs are much shorter than the membrane leak time constant ($\Delta t_i \ll \tau_m$). In such cases, the solution of eqn. 3.1 for constant input given below

$$V(t) = V_{rest} + IR \left(1 - e^{-\frac{t-t_0}{\tau_m}} \right), \quad (3.11)$$

can have its leak term replaced by an approximation of the term's Taylor series expansion as shown in eqn. 3.12.

$$e^{-\frac{t-t_0}{\tau_m}} \approx 1 - \frac{t-t_0}{\tau_m} \quad (3.12)$$

Therefore, from eqns. 3.11 and 3.12, the membrane potential equation of the LIF model is simplified and approximated by eqn. 3.13.

$$V'(t) = V_{rest} + \frac{I(t-t_0)}{C} \quad (3.13)$$

where $V'(t)$ signifies the membrane potential of the approximating model and C is the capacitance of the membrane.

This approximating model is the perfect (i.e., non-leaky) Integrate-and-Fire neuron model (PIF), which simply integrates post-synaptic inputs, without losing any of its charge over time. We then use eqns. 3.11 and 3.13 to calculate the relative difference between the two models as a function of the ISI (Δt_i). The relative difference is measured at the beginning of the pre-spike coincidence window, because the membrane potential at this time determines the slope of the secant line associated with that specific spike (see eqn. 3.8 and Figure 3.2a). Therefore, the relative difference for any given ISI is calculated as the difference between the two models at the beginning of the coincidence window, i.e., $t_i - w$ (eqn. 3.14).

$$d_{t_i, w} = \frac{|V(t_i - w) - V'(t_i - w)|}{V(t_i - w)} \quad (3.14)$$

where $V(t)$ and $V'(t)$ are given by eqns. 3.11 and 3.13 respectively. The relative difference is used as a measure of dissimilarity between the two models and by extension, it measures the distinguishability between the two operational modes. Therefore, the relative difference d represents the level of divergence between the LIF and PIF.

Alternatively, we can look at the behaviour of the two bounds U and L as a function of the ISI and observe the rate of their convergence as the interval is decreased. More specifically, the ratio of the lower bound eqn. 3.5 to the upper bound eqn. 3.4, $\frac{L}{U}$, can show us how fast the lower bound approaches the upper bound for decreasing ISI.

Figure 3.5a shows the relative difference d (eqn. 3.14), as a function of the ISI (Δt_i) at high firing rates, for a coincidence window $w = 2$ ms. Similarly, Figure 3.5b shows the ratio of the bounds $\frac{L}{U}$, for the same scenario. Note that changing the value of w shifts both curves along the horizontal axis, i.e., increasing the value moves the curve to the right and decreasing the value moves the curve to the left. The figures show that at high firing rates, the LIF and perfect integrator converge, as d gets closer to 0 and therefore the two models become indistinguishable. Similarly, the ratio between the two bounds approaches 1, as described in the beginning of this section, indicating that the two operational modes are also indistinguishable.




FIGURE 3.5: (a) Relative difference d between the LIF and the perfect integrator models and ratio of lower bound to upper bound ($\frac{L}{U}$) as a function of firing ISI (Δt_i), for window $w = 2$ ms (solid lines). For details on the derivation of the difference d see eqn. 3.14. The three points marked in each graph correspond to the firing rates of the simulations which produced Figure 3.3b, (100 Hz), Figure 3.3d (70 Hz), Figure 3.3f (400 Hz). As we are interested in the relative difference and convergence of the bounds at high firing rates, the graph does not show values corresponding to firing rates lower than 50 Hz. Note that changing the value w shifts both curves along the horizontal axis, i.e., increasing the value moves the curves to the right and decreasing the value moves the curve to the left.

3.6 Results for LIFwPR

We also measured the normalised pre-spike membrane potential slope of a model neuron exhibiting highly irregular firing at high rates. We used the LIFwPR, with neuron and input parameter values identical to the model by Bugmann et al. [35]. The inputs to the neuron consisted of 50 Poisson spike trains and each input spike caused a depolarisation of the neuron's membrane potential by $\Delta V_s = 0.16$ mV. We used a reset parameter value of $\beta = 0.91$ as it has been shown to be the only value that can produce purely temporally irregular firing (with no bursting activity that can increase the firing variability) [35, 77]. This is compatible with the high firing irregularity at high rates observed in cortical neurons [72, 73].

The results (Figure 3.7) show the value of \bar{M} being below 0.35 for the entire range of firing rates and below 0.2 for rates below 300 Hz. Each firing rate was achieved by varying input rates within physiological ranges. These results suggest that neurons firing highly irregularly at high rates operate mainly as temporal integrators.

In Bugmann et al. [35], it was suggested that temporal integration and fluctuation detection (i.e., coincidence detection) can coexist and cause irregular firing, which was indicated by the ISI of a LIFwPR neuron driven by a fluctuating input current being significantly shorter than the ISI of the same neuron driven by a constant input current (of the same average value).

The current results however indicate that there is a strong dominance of temporal integration and the relatively small contribution of coincidence detection in the firing of spikes is not sufficiently high to be distinguishable from the effects of the temporal integration process. This indicates that coincidence detection is not necessary for producing highly irregular firing at high rates (as was suggested by Softky and Koch [72, 73]) and that temporal integration on its own is sufficient for such a purpose, provided the neuron does not completely repolarise.

3.6.1 Compatibility of results with original analysis

Our results are not incompatible with the analysis by Softky and Koch [73]. In their work, these authors express the threshold in number of input pulses, N_{th} , necessary to raise the neuron's membrane potential from rest to discharge. We can use the definition of N_{th} in the same way and describe our model in terms of the difference in potential between threshold and rest, divided by the depolarisation per spike:

$$N_{th} = \left\lceil \frac{V_{th} - V_{rest}}{\Delta V_s} \right\rceil \quad (3.15)$$

where $\lceil x \rceil$ denotes the ceiling of x (i.e., the smallest integer not less than x). This is done in order to make our results comparable with their analysis, which showed how the coefficient of variation (CV) varied as a function of the time constant τ_m and N_{th} , as seen in Figure 3.6.

Assuming that $V_{reset} \geq V_{rest}$ (which holds for any model with a reset parameter $\beta \geq 0$), for any given time where $V(t) \geq V_{reset}$, the model neuron can be expressed in terms of an equivalent model with effective resting potential $V'_{rest} = V_{reset}$ and effective time constant τ'_m (the prime signifies a parameter or variable of the equivalent model). From this, it follows that $N'_{th} \leq N_{th}$, since $V'_{rest} \geq V_{rest}$:

$$N'_{th} = \left\lceil \frac{V_{th} - V'_{rest}}{\Delta V_s} \right\rceil = \lceil N_{th}(1 - \beta) \rceil \quad (3.16)$$

Substituting for the parameter values used in our simulations, i.e., $V_{th} = 15$ mV, $V_{rest} = 0$ mV and $\Delta V_s = 0.16$ mV, the original value of N_{th} is 94. For $V'_{rest} = 13.65$ mV however, which is the reset value of the LIFwPR model ($\beta = 0.91$), the effective number of input spikes required to fire a spike N'_{th} is reduced to just 9.




FIGURE 3.6: Contour plot of CV for leaky integrator for discrete values of τ_m (shown as τ in the figure) and N_{th} , when the mean output ISI is 5 ms. Image taken from Softky and Koch [73]

The effective time constant's value τ'_m should be such that (assuming the inputs are the same for both models) the rate of change of the membrane potential for the equivalent model should be equal to that of the original model, $\frac{dV}{dt} = \frac{dV'}{dt}$. Since the two models share the same input, we can calculate τ'_m by ignoring the input terms of the two models and equating the leak term of eqn. 3.1 with the leak term of the equivalent model (eqn. 3.17):

$$\tau'_m \frac{dV'}{dt} = -V(t) + V'_{rest} \quad (3.17)$$

Comparing eqns. 3.1 and 3.17, replacing the effective resting value V'_{rest} with the original reset value V_{reset} and solving for τ'_m gives eqn. 3.18:

$$\tau'_m = \tau_m \frac{V(t) - V_{reset}}{V(t) - V_{rest}} \quad (3.18)$$

Therefore the value of the effective time constant τ'_m changes as a function of the membrane potential $V(t)$. For our simulations, we can calculate the range of values that τ'_m takes, first by substituting the parameter values we used, i.e., $\tau_m = 10$ ms, $V_{rest} = 0$ mV and $V_{reset} = 13.65$ mV and then by calculating τ'_m for the known range of $V(t)$ using eqn. 3.18. Thus, for our simulations $\tau'_m = 10 \text{ ms} (V(t) - 13.65 \text{ mV}) / V(t)$.

Since the effects discussed here are relevant for membrane potential levels above the reset potential, we calculate the range of τ'_m for membrane potential values between $V(t) \in [V_{reset}, V_{th}] = [13.65 \text{ mV}, 15 \text{ mV}]$ giving respective $\tau'_m \in [0 \text{ ms}, 0.9 \text{ ms}]$.

Therefore, for the LIFwPR, when the input is constantly high enough to keep the membrane potential above the reset potential, the effective number of spikes required to reach threshold becomes $N_{th} = 9$ and the effective membrane time constant becomes a function of the membrane potential and fluctuates within the range $\tau_m \in [0, 0.9]$ (with the parameter values used to get the results shown in Figure 3.7). This result describes theoretically how a model behaving primarily as an integrator, such as the LIFwPR, can fire at such high variability at high rates. With the partial reset mechanism, the LIFwPR neuron's membrane potential stays very close to the spike threshold between successive firings (i.e., during an ISI), assuming the neuron is spiking at high enough rates. More importantly, the membrane potential is almost always above the reset potential V_{reset} after the first spike is fired. With this in mind, we have shown that when the LIFwPR neuron is driven by sufficiently frequent arriving inputs, it operates equivalently to a neuron with an effective resting potential V'_{rest} equal to the reset potential V_{reset} , a very short effective time constant ($\tau'_m \leq 0.9 \text{ ms}$), and a very low effective threshold ($N'_{th} = 9$).

Our results are thus in accord with the analysis by Softky and Koch [73] who showed that for low N_{th} values and sub-millisecond membrane time constant τ_m , a LIF neuron operating as a temporal integrator can fire highly irregularly at high rates (Figure 3.6) [73]. From the above analysis, we can therefore conclude that the LIFwPR model, which models the incomplete post-spike re-polarisation of a neuron, can be used for (i) reducing the effective number of input spikes N'_{th} required to cause a spike and (ii) decreasing the effective membrane leak time constant τ'_m such that a neuron can fire highly irregularly at high rates, in accordance with experimental recordings.

3.7 Discussion

This chapter establishes the correlation between input synchrony and the slope of the membrane potential prior to firing. This depends on normalising the slope between two bounds. Our results suggest a strong correlation between pre-spike membrane potential slope and pre-synaptic synchrony levels, that allows us to infer the degree of response-relevant input




FIGURE 3.7: The normalised pre-spike membrane potential slope (\overline{M}) for a LIF neuron model with partial somatic reset firing highly irregularly at rates up to ~ 470 Hz ($\overline{ISI} \approx 2.1$ ms).

synchrony under certain assumptions, namely the existence of excitation only and of supra-threshold volleys. In a theoretical study, Stein [10] showed that the slope of the membrane potential is inversely proportional to the variance of the firing ISIs, for a neuron driven by Poisson inputs. Goedeke and Diesmann [32] showed that the membrane potential, as well as its derivative, define the response of a LIF model to synchronised inputs. They analytically studied the dynamics of the behaviour of a LIF neuron, both in isolation and in homogeneous networks and concluded that the synchronisation between neurons depends both on the membrane potential and its derivative. While these studies prove the existence of a correlation between membrane potential and firing statistics of both individual neurons and networks, our own work establishes a specific correlation measure between membrane potential and input statistics. As such, the two results may be considered complimentary.

However, a potential correlation between firing and input statistics is most likely not as straightforward to investigate, since both the membrane potential fluctuations and the firing ISI distribution are affected by multiple parameters of the stimulus. Synchrony in the pre-synaptic activity of a neuron can affect its firing rate [153] and irregularity [154]. However, these effects are not consistent and depend heavily on the state of the neuron. More precisely, the output firing rate is a non-monotonic function of the correlation among excitatory inputs [153]. Additionally, the firing variability depends heavily on other factors besides the degree of input

correlations [154]. The NPSS relies on the assumption that changes in input parameters are reflected in the trajectory of the membrane potential, while similar changes may not affect the distribution of firing ISIs in a consistent manner.

Our work is more closely related to Kisley and Gerstein [31] and more recently to DeWeese and Zador [33] and Kobayashi et al. [34], in that we establish a relationship between membrane potential properties and properties of the input spike trains, in order to infer the latter from measurements of the former. DeWeese and Zador [33] analysed membrane potential dynamics to infer properties of the input population. Similarly, Kobayashi et al. [34] developed an algorithm to estimate the time-varying input rates of the pre-synaptic population by studying the membrane potential of the neuron. The correlation between membrane potential slope and input synchrony was studied by Kisley and Gerstein [31]. The work presented in this chapter relies on this correlation to provide a measure of the response-relevant input synchrony, which relates to the operational mode of the neuron. In particular, the NPSS provides a measure of the relative contribution of temporal integration and coincidence detection to the firing of a spike, or the operation of a neuron in general.

The choice of the duration of the coincidence window, i.e., the value of w , is an important aspect of the NPSS. It has to be noted that the effect of this variable on the temporal precision of firing has also been the subject of theoretical study [120, 122]. As mentioned in Section 3.4, the value of this parameter should reflect the time that is regarded as the maximum temporal distance between two events that are considered to be coincident (2 ms in our case). The only limit for the duration of the coincidence window is the time step of the simulation (here 0.1 ms) or more generally, the temporal resolution of the data being analysed. However, the smaller the value of w , the stricter the definition of coincident activity becomes, which in turn produces lower \overline{M} values, unless the input spike trains are completely synchronised. This provides a degree of flexibility for the NPSS that allows it to be adapted to various levels of temporal resolution.

For our simulations, we set the coincidence window to $w = 2$ ms since, as mentioned in Section 2.2, for neurons with membrane time constants within the range of 10 to 20 ms, a temporal code with accuracy in the range 1 to 3 ms is theoretically possible [117]. By setting the width of the coincidence window to 2 ms, we effectively attempt to measure the level of synchrony of the input spikes that caused each response, under the assumption that the temporal precision of coincidence detection is 2 ms. Higher precision has also been identified in experiments on a CA1 hippocampal pyramidal neuron [155]. In their experiments, Losonczy and Magee [155]

observed that the cell is sensitive to coincidences with a precision of up to 0.1 ms. As mentioned above, setting w to such a small value would greatly affect our results, by significantly reducing the values measured by the NPSS. Such a strict definition of coincidences would naturally translate to a very strict definition of *coincidence detection*, making it much less likely that the neuron is operating in an operational mode other than *temporal integration*, for even the slightest amount of jitter.

The meaning of the value for the duration of the coincidence window can be intuitively understood in terms of the cost parameter found in spike train distance measures [see 112, 114, 142, as well as the relevant subsection of Section 2.3 on spike train distances], which controls the sensitivity of the metric to spike count and spike timing, i.e., the assumed resolution of the temporal code. Spike train distances measure the distance between two spike trains by calculating the minimum cost of transforming one spike train into the other by adding, removing or shifting spikes. By manipulating the cost parameter, one can control the measured distance between two given spike trains. For instance, with a small cost parameter value, two very different spike trains will be measured as having a small distance, i.e., they are considered similar by the metric due to the low cost of shifting spikes. Conversely, with a high cost parameter value, two similar spike trains will be measured as having a large distance, i.e., they are considered dissimilar by the metric, due to the high cost of shifting spikes. While the NPSS measures the response-relevant synchrony of the input spike trains of a neuron and by extension, the operational mode of that neuron, the spike train distances measure the distance, or similarity between a pair or group of spike trains directly. However, both types of metrics can be used in different circumstances to measure the temporal precision of the neural code.

While similar work exists on measuring spike train correlations and synchrony, either by directly observing the spikes fired from a population of neurons [133, 135], or by identifying synchronous activity in local field potentials [156], the NPSS differs in that it only responds to such correlations between spike trains converging into a single neuron, when they are responsible for the triggering of response spikes. In particular, the NPSS explicitly calculates the degree of input synchrony directly preceding a response spike and implicitly considers any previous activity by taking into account the potential at the start of the coincidence window w in the calculation. The higher the potential of the neuron's membrane at the start of the coincidence window, the lower the relative contribution of the synchronous spike trains within the coincidence window would be to the response. Consequently, the slope of the membrane potential within the coincidence window is low, denoting a higher contribution of temporal

integration. In this way, the NPSS is only concerned with the input statistics that affect the neuron's own spiking, in other words, it is sensitive to the response-relevant statistics of the input. It is this particular feature which links our measure's estimation of response-relevant input synchrony to the underlying operational mode. The operational mode of a neuron is not defined solely by the synchrony of the spike trains it receives, but also by whether or not that synchronous activity causes firing. We refer to this distinction as the difference between *actual synchrony* and *utilised synchrony*. The actual synchrony is the synchrony between a set of spike trains, determined by the amount of coincident spikes in relation to the total amount of spikes (coincident versus non-coincident, or synchronous versus random). The utilised synchrony is determined by the way in which the actual synchrony affects the neuron's spiking, whether coincident spikes are able to cause firing, or non-coincident spikes are also required. Chapter 4, studies this relationship in more depth.

The simulation method for the absolute refractory period described in Section 3.2 differs from the traditional method of clamping the membrane potential to V_{rest} or V_{reset} for the duration t_r . In our simulations, the membrane potential is allowed to evolve freely during t_r , but the comparison with the firing threshold is not performed, effectively disabling the firing mechanism. This method is used in models with a partially resetting [35] or non-resetting [80] membrane potential following a spike. This has several implications for the firing behaviour of the model. It allows for ISI lengths of exactly t_r , which occurs when the membrane potential reaches V_{th} during the t_r period following a spike and remains above threshold until the end of the refractory period. When this occurs, the neuron fires immediately when the threshold is reactivated at the end of the refractory period. With the traditional simulation method, the smallest ISI is determined by the absolute refractory period, plus a small duration required for the membrane potential to reach threshold from V_{rest} , once the membrane potential integration is reactivated (i.e., the clamping is disabled). This rise time delay is generally determined by the membrane leak time constant (τ_m). Moreover, using this method of simulating the refractory period results in a model that behaves in the same way as a model with no refractory period, as long as the smallest ISI is $\geq t_r$. When this holds, the effect of the refractory period (i.e., the disabling of the threshold) has no impact on the firing of the neuron, in addition to the refractory period having no effect on the behaviour of the membrane potential. Our choice of parameter value for the refractory period ($t_r = 2$ ms) follows from our usage of the LIFwPR model [35] and our goal of reproducing and analysing the behaviour of a model neuron firing highly irregular spike trains at very high rates. The duration and simulation method of the

refractory period has little impact on our results, however. The slope bounds defined for the NPSS can very easily be adjusted to account for a refractory period, assuming it is known, by subtracting its duration (t_r) from the duration of the ISI (Δt_i) during the bound calculations (see eqns. 3.4–3.7).

One of the motivations for the work presented in this chapter was to determine the operational mode responsible for the firing of the highly irregular spike trains at very high rates. In particular, we were interested in the operational mode of the LIF model with partial reset (LIFwPR), which adds a simple modification to the spiking mechanism of a simple model and has been shown to accurately reproduce the firing statistics of experimentally recorded data [35, 77]. Similarly, models which implement a variable threshold, such as the Multi-timescale Adaptive Threshold (MAT) model [80], can reproduce a wide variety of firing characteristics and accurately predict spike times. While studying the LIFwPR model, Bugmann et al. [35] showed that the partial reset mechanism is functionally equivalent to a model with a totally resetting membrane potential and a time varying threshold (LIFwVT), where the threshold decreases immediately following a spike and increases to its original value over time. This differs from the behaviour of the MAT model, which has a non-resetting membrane potential and a variable threshold which increases immediately following a spike and decays to its original value over time. The three models (LIFwPR, MAT and LIFwVT) can behave equivalently when the immediate change in threshold is set accordingly across all models and the decay constant of the threshold is equal to the membrane leak time constant. This implies that our use of the LIFwPR model and slope bound equations can be used to study the dynamics of all of these models, with appropriate modifications. However, it should be noted that the MAT model allows for modification of its firing dynamics by manipulating the decay rate of the threshold, which is an extra parameter that has no equivalency in the LIFwPR model. More specifically, when the threshold decay constant of the MAT model is constant and equal to the membrane leak time constant and the reset parameter of the LIFwPR corresponds to the increase in threshold of the MAT model, the two models behave very similarly. The equivalence between the two models does not hold when the threshold decay constant is not equal to the membrane leak time constant.

The methods and results presented in this chapter were published in *Neural Computation* in 2012 [36].

Chapter 4

The relation between stimulus synchrony and the operational mode of a neuron

In this chapter, we investigate the relation between the operational mode, as determined by the NPSS and the spike time distance of the spike trains driving the neuron, as measured by the SPIKE-distance (see Section 2.3). Our goal is to determine the circumstances under which input synchrony directly determines operational mode (i.e., high synchrony causes coincidence detection while low synchrony causes temporal integration) and more importantly, to investigate how the properties of the input (number of inputs, rate), combined with the properties of the neuron itself (synaptic weight, membrane leak time constant and threshold) affect this relationship.

We chose to use the SPIKE-distance measure [142] for the reasons described in Section 2.3. More specifically, the SPIKE-distance and in particular the multivariate variant of the measure, is more suitable than other spike time distance measures for measuring the distance between multiple (i.e., more than two) spike trains. Spike time distance measures such as the Victor-Purpura (V-P) distance [141], the van Rossum distance [113], and the modulus metric [144], operate on spike train pairs. Calculating the distance between multiple spike trains using these measures involves taking the average of all pairwise distances, which scales with the square of the number of spike trains (number of pairs between N spike trains = $N(N - 1)/2$). This is true for the original (bivariate) version of the SPIKE-distance as well. Calculating the spike

distance between N spike trains becomes infeasible when N is on the order of hundreds of spike trains, as used in this chapter. The multivariate SPIKE-distance scales linearly, since it only uses the standard deviation of the times of spikes around a certain time t in its calculation (see Section 4.1.1 below and in particular eqn. 4.1).

In addition to the lower computational complexity, the SPIKE-distance is parameter free, as opposed to the V-P and van Rossum measures. The V-P measure is parameterised by a cost parameter which defines the cost (distance) between nearby spikes (in units of cost per second). The van Rossum measure is parameterised by a time constant which affects the exponential convolution window used in the calculation of the distance. These two parameters are inversely related and in both cases they are used to control the temporal sensitivity of the measures. In other words, the values of the parameters control whether the measures assume a rate code or a temporal code. In this chapter, we aim to understand the relationship between the operational mode and the input synchrony, across the whole continuum of operational modes. The SPIKE-distance measure is more appropriate for this case since it makes no assumptions about the underlying operational mode, or neural code.

Finally, the SPIKE-distance is time-scale independent. Calculation of the distance involves scaling the measurements using the instantaneous mean ISI around a certain point in time, which results in the measure removing any dependence on the time scales or firing rates of the underlying spike trains. This is more compatible with the way the NPSS calculates the operational mode. The calculation of the membrane slope bounds accounts for the different bounds that are theoretically achievable given the length of a specific ISI (see eqns. 3.4 and 3.5). The subsequent normalisation of the slope using the bounds (eqn. 3.8) therefore removes the dependence of the slope measurement on the ISI and by extension, the spike rate.

Generally, the meaning of the two measures can be interpreted as *actual synchrony* for the SPIKE-distance versus *utilised synchrony* for the NPSS. In the former case, the SPIKE-distance is derived directly from the spike trains and is a measure of the synchrony between them, unaffected by spike rates (it is time-scale independent) or the properties of the neuron. On the other hand, the NPSS is a measure of how the neuron utilises the synchrony in the input, whether synchronous spike trains are solely responsible for causing response spikes, no coincidences are responsible for causing response spikes, or any combination in between. An important distinction is that the SPIKE-distance is measured on the input spike trains directly, i.e., the spike times are used as inputs into the calculation. The SPIKE-distance measures actual synchrony:

it is the degree synchrony that is measured by applying a method on the spike trains. On the other hand, the NPSS is measured on the membrane potential of the neuron, i.e., it analyses the behaviour of the membrane potential while the neuron is being driven by the input spike trains. The NPSS measures utilised synchrony: it is the degree of synchrony that affects the trajectory of the membrane potential after it has been integrated (utilised) by the neuron. The term *utilised* in this context should not be confused with the idea of utility in any higher level function. It is simply meant to refer to the degree of synchrony that is responsible for causing response spikes. In other words, we are referring to the degree of input synchrony that is effective in causing responses, as opposed to the degree of synchrony that appears at the input. The degree of utilised synchrony measured by the NPSS defines the neural mode: full synchrony utilisation defines coincidence detection, while no synchrony utilisation defines temporal integration and intermediate values define the operational mode continuum. The purpose of this chapter is to investigate how the input and neuron parameters determine the degree of utilised synchrony, i.e., the operational mode, in relation to the actual input synchrony.

4.1 Methods

The overall methodology used is as follows:

1. Generate spike trains with parameter values (number of spike trains N_{in} , spike frequency f_{in} , fraction of synchronous spike trains S_{in} and jitter σ_{in}) randomised within the ranges defined in Table 4.1.
2. Use the spike trains to drive a neuron model, causing it to fire. The weight (depolarisation caused per spike, ΔV_s) is also randomised and shown in Table 4.1.
3. Calculate the NPSS as described in Section 3.4 [36].
4. Calculate the multivariate SPIKE-distance (D_S) [142] between the generated (input) spike trains using the method described in Section 4.1.1 below.
5. Compare the NPSS with D_S . See Section 4.1.2.

For step 1, inputs were generated as in Section 3.3. Table 4.1 shows input related parameter values and ranges ($N_{in}, f_{in}, S_{in}, \sigma_{in}$). In step 2, we use the leaky integrate-and-fire (LIF) model, as described in Section 3.2, with total reset ($V_{rest} = V_{reset}$). The neuron parameters are

Symbol	Description	Value or range	Units
N_{in}	Number of input spike trains	30 to 400	—
f_{in}	Input spike frequency	50 to 100	Hz
ΔV_s	Synaptic weight	0.1 to 1.0	mV
S_{in}	Input synchrony degree	0 to 1	—
σ_{in}	Gaussian jitter	0 to 4	ms
V_{th}	Firing threshold	15	mV
V_{rest}	Neuron resting and reset potential	0	mV
τ_m	Membrane leak time constant	10	ms

TABLE 4.1: Parameter symbols and value ranges

duplicated in Table 4.1 for completeness. In step 3, we calculate the NPSS after running the LIF neuron with the generated inputs from step 1 for 5 s of simulated time.

4.1.1 SPIKE-distance

The multivariate SPIKE-distance (step 4) is calculated using the method described in [142] and Section 2.1.3 of the cited paper in particular. For any time t , the spike distance is:

$$S^m(t) = \frac{\sigma \left[t_P^{(n)}(t) \right]_n \langle X_F^{(n)}(t) \rangle_n + \sigma \left[t_F^{(n)}(t) \right]_n \langle X_P^{(n)}(t) \rangle_n}{\langle X_{ISI}^{(n)}(t) \rangle_n^2}, \quad (4.1)$$

Where $\langle X_P^{(n)}(t) \rangle_n$ and $\langle X_F^{(n)}(t) \rangle_n$ are the average intervals to the previous and following spikes respectively, across all N_{in} spike trains. $\sigma \left[t_P^{(n)}(t) \right]_n$ and $\sigma \left[t_F^{(n)}(t) \right]_n$ are the standard deviations of the spike times of the previous and following spikes respectively, across all N_{in} spike trains. Finally, $\langle X_{ISI}^{(n)}(t) \rangle_n$ is the average ISI across the N_{in} spike trains, around time t .

$$t_P^{(n)} = \max(t_i^{(n)} | t_i^{(n)} \leq t) \quad (4.2)$$

$$t_F^{(n)} = \min(t_i^{(n)} | t_i^{(n)} > t) \quad (4.3)$$

$$X_P^{(n)} = t - t_P^{(n)} \quad (4.4)$$

$$X_F^{(n)} = t_F^{(n)} - t \quad (4.5)$$

$$X_{ISI}^{(n)} = t_F^{(n)} - t_P^{(n)} \quad (4.6)$$

Figure 4.1 shows an illustration of the quantities involved in the calculation of the SPIKE-distance and eqns. 4.2–4.6.




FIGURE 4.1: Illustration of the values used in the calculation of the SPIKE-distance. The figure shows two spike trains (1: blue and 2: red) with four and two spikes respectively. The dashed black line represents the time (t) for which the local synchrony is being calculated. $t_P^{(n)}$ is the time of the last spike preceding time t for spike train n . $t_F^{(n)}$ is the time of the first spike following time t for spike train n . $X_P^{(n)}$ is the time difference between t and $t_P^{(n)}$ and $X_F^{(n)}$ is the time difference between t and $t_F^{(n)}$. $X_{ISI}^{(n)}$ is the duration of the interval around t , i.e., it is the time interval between $t_P^{(n)}$ and $t_F^{(n)}$. $\Delta t_P^{(n)}$ is the time difference between $t_P^{(n)}$ and the nearest spike in the other spike train and similarly for $\Delta t_F^{(n)}$. These last two values are not used in the calculation of the multivariate SPIKE-distance, which we use in this chapter.

Image taken from Kreuz [157].

4.1.2 Comparison

Comparisons between the NPSS and the SPIKE-distance are made on averages across each simulation. More specifically, for the NPSS we use the arithmetic mean (eqn. 3.10) and for the SPIKE-distance, the integral over time,

$$D_S = \frac{1}{T} \int_{t=0}^T S^m(t) dt. \quad (4.7)$$

Numerically, the SPIKE-distance was computed at intervals for t of 1 ms and the integral was calculated using the trapezoidal rule. The interval was chosen based on trials where the SPIKE-distance was calculated using decreasing values of dt , on the same set of spike trains. During

these trials, the D_S followed an asymptotic trajectory, with minimal change for values of dt below several ms. The value of 1 ms was also chosen to be small enough so that no interval can contain two spikes and it is large enough so that computation was fast, as smaller values increase the number of computations required.

4.2 Results

Our results initially focus on the relationship between the NPSS and the SPIKE-distance for all simulations. We look at the effect of several parameters and properties of the simulation on the relationship between the two measures and discuss the cause and implications of the value of the parameter in question. We then generalise our conclusions regarding the effect of synchrony on the operational mode.

For the following analysis, we introduce two features of the model: the asymptotic threshold-free potential (V_∞) and the volley peak potential (Δ_v). The asymptotic threshold-free potential is the asymptotic temporal mean of the membrane potential in the absence of a spike threshold and is calculated as follows:

$$V_\infty = N_{in} f_{in} \Delta V_s \tau_m + V_{rest}. \quad (4.8)$$

The volley peak potential is the increase in potential at the membrane of the neuron when a completely synchronous volley arrives from all input spike trains and is calculated as follows:

$$\Delta_v = N_{in} \Delta V_s. \quad (4.9)$$

4.2.1 Effect of jitter (σ_{in})

Figure 4.2a shows the relationship between the NPSS and the multivariate SPIKE-distance for each simulation. The NPSS (\bar{M}) indicates the operational mode by measuring the degree of utilised synchrony and ranges from 0 to 1. $\bar{M} = 0$ indicates that there is no synchrony responsible for the firing of the neuron, while $\bar{M} = 1$ means that the neuron fires solely due to synchronous spike trains. The SPIKE-distance (D_S) indicates the average spike distance between the input spike trains, which is an inverse measure of synchrony. $D_S = 0$ indicates that all spike trains are synchronous (identical) while larger numbers indicate less synchrony. In




FIGURE 4.2: NPSS (\bar{M}) vs SPIKE-distance (D_S) with the colour of each point showing the degree of applied jitter (σ_{in}). High values mean high degrees of applied jitter (up to 4 ms) which makes spike volleys wider, while low values indicate that volleys are tighter and more synchronous, with 0 ms indicating that all spikes in a volley are simultaneous. (a) shows all points while the three smaller figures split the point into subsets depending on the value of σ_{in} . (b) $\sigma_{in} = 0$ ms. (c) $0 \text{ ms} < \sigma_{in} \leq 2$ ms. (d) $\sigma_{in} \geq 2$ ms.

our simulations, the maximum value of D_S approached 0.5, which is in accordance with the values shown in Kreuz et al. [142]. The data consists of 253 simulations where no jitter was added to synchronous spike trains ($\sigma_{in} = 0$ ms) and 1012 simulations with jitter falling within the value ranges shown in Table 4.1 (1265 total). Figures 4.2b and 4.2d show the same data, split into three cases, depending on the degree of jitter. For Figure 4.2b, $\sigma_{in} = 0$ ms, for Figure 4.2c, $0 \text{ ms} < \sigma_{in} \leq 2$ ms and for Figure 4.2d, $2 \text{ ms} \leq \sigma_{in}$ (up to 4 ms). In all three figures, the colour of each point represents the value of jitter applied, denoted by the colour bar on the right-hand side.

In Chapter 3, we noted that increasing jitter from 0 to 4 ms, causes the operational mode to shift from coincidence detection to temporal integration. This is especially apparent in cases where volleys are supra-threshold ($\Delta_v > V_{th} - V_{rest}$, see below) and pre-synaptic firing is completely synchronised ($S_{in} = 1$). In such cases, the value of σ_{in} can cause the neuron to

operate as a pure coincidence detector (if $\sigma_{in} = 0$ ms, $\bar{M} = 1$) or a pure temporal integrator (if $\sigma_{in} \geq 4$ ms, $\bar{M} = 0$). The data presented in Figure 4.2 generalises this conclusion, since as the jitter increases, the data points tend to stay closer to the upper left part, where the SPIKE-distance is high ($D_S > 0.2$) and the NPSS is low ($\bar{M} < 0.6$). This is expected, since lower synchrony, caused by higher jitter, causes the neuron to operate as a temporal integrator, regardless of any other behaviour or parameter. It is especially apparent in Figure 4.2d, which shows the results with high jitter. The effect of increasing the amount of jitter is the reduction of points in the lower right corner, i.e., high synchrony and coincidence detection are not possible. This is also accompanied by an increase in points in the upper left corner or even the middle-left side, which indicates that the high amount of jitter causes the NPSS to measure temporal integration, even though there is some degree of synchrony measured by the SPIKE-distance.

4.2.2 Effect of threshold-free potential (V_∞)

Figure 4.3 shows the same data as Figure 4.2, with the colour of each point representing the asymptotic, threshold-free potential. This term is the asymptotic value of the neuron's membrane potential, assuming an absence of a spiking threshold. If it is higher than the spike threshold, response spikes are generated almost surely, while if it is sub-threshold, response spikes are generated only if there is enough noise (either in the input, or intrinsic to the neuron) causing random fluctuations that could drive the potential above threshold. The three smaller subfigures, Figures 4.3b and 4.3c, separate the points into three categories, based on the relationship between V_∞ and the neuron's spike threshold V_{th} . For Figure 4.3b, $V_\infty < V_{th}$, for Figure 4.3c, $V_{th} \leq V_\infty < 2V_{th}$ and for Figure 4.3d, $V_\infty \geq 2V_{th}$.

It is interesting to note the tighter concentration of points in Figure 4.3d. This subfigure shows the relationship between \bar{M} and D_S for very high values of the asymptotic membrane potential. This indicates that the input drive is very high, causing the neuron to fire at high firing rates. As mentioned in Section 3.5.1, as firing rates increase, the range of slope values shrinks, making the two extremes indistinguishable once the firing rate becomes very high. What appears in Figure 4.3d is the trend towards overestimation of input synchrony by the NPSS, i.e., the points start concentrating in the upper right hand corner, where the spike distance is high (input synchrony is low) yet the NPSS measures high \bar{M} value, close to coincidence detection. This is the phenomenon that occurs in Figure 3.3f.




FIGURE 4.3: NPSS (\bar{M}) vs SPIKE-distance (D_s) with the colour of each point showing the asymptotic threshold-free potential (V_∞ , see eqn. 4.8). (a) shows all points while the three smaller figures split the point into subsets depending on the value of V_∞ with respect to the spike threshold (V_{th}). (b) $V_\infty < V_{th}$. (c) $V_{th} \leq V_\infty < 2V_{th}$. (d) $V_\infty \geq 2V_{th}$.

4.2.3 Effect of volley peak potential (Δ_v)

Similarly, Figure 4.4 shows the same data again, but with the colour of each point representing the peak voltage change caused by a synchronous volley (when no jitter is applied). The three smaller subfigures, Figures 4.4b–4.4d, separate the points into three categories, based on the relationship between Δ_v and the neuron's spike threshold V_{th} . For Figure 4.4b, $\Delta_v < V_{th} - V_{rest}$, for Figure 4.4c, $V_{th} \leq \Delta_v < 2(V_{th} - V_{rest})$ and for Figure 4.4d, $\Delta_v \geq 2(V_{th} - V_{rest})$.

4.2.4 Integration of coincident volleys

The data points in Figure 4.2b that do not follow the trend seen in the other points of the same subfigure (points where $\bar{M} \approx 0.5$ and $D_s < 2$) are of particular interest. This deviation is caused by sub-threshold values for the asymptotic potential (V_∞) and more importantly the volley peak (Δ_v). This is evident in Figures 4.3b, 4.3c, 4.4b and 4.4c, which show the behaviour




FIGURE 4.4: NPSS (\bar{M}) vs SPIKE-distance (D_S) with the colour of each point showing the peak depolarisation of a fully synchronous volley (Δ_v , see eqn. 4.9). (a) shows all points while the three smaller figures split the point into subsets depending on the value of Δ_v with respect to the spike threshold (V_{th}). (b) $\Delta_v < V_{th} - V_{rest}$. (c) $V_{th} \leq \Delta_v < 2(V_{th} - V_{rest})$. (d) $\Delta_v \geq 2(V_{th} - V_{rest})$.

occurring for low values of V_∞ and Δ_v . Even more clearly, the phenomenon is exemplified in Figures 4.5a, 4.6a and 4.6b, which show how the operational mode is determined by varying degrees of input synchrony when both the asymptotic potential (V_∞) and the volley peak (Δ_v) are below threshold. More specifically, when $\Delta_v < V_{th} - V_{rest}$, coincidence detection cannot occur, even for $S_{in} = 1$, since a single spike volley cannot bring the neuron to fire from rest. Instead, the neuron may *integrate* multiple spike volleys in order to reach the firing threshold and in that case, the NPSS measures a balance of integration and coincidence detection ($\bar{M} \approx 0.5$). The points that appear at $\bar{M} \approx 0.5$ and $D_S = 0$ on the figures are examples of this occurrence: the input spike trains were completely synchronous ($S_{in} = 1$) and there was no jitter ($\sigma_{in} = 0$ ms), which results in a SPIKE-distance of zero, but the neuron behaves in a manner signified by the middle of the operational mode continuum, where neither pure coincidence detection nor temporal integration is taking place. In this case, we describe the behaviour of the neuron as *integrating coincidences*, to denote that high degrees of synchrony are driving the




FIGURE 4.5: NPSS (\bar{M}) vs SPIKE-distance (D_S) with the colour of each point showing the degree of applied jitter (σ_{in}). Each subplot is separated based on the relationship of the asymptotic threshold-free potential (V_∞) and the volley peak (Δ_v) to the spike threshold (V_{th}). (a) Sub-threshold asymptotic potential and sub-threshold peak: $V_\infty < V_{th}$ & $\Delta_v < V_{th} - V_{rest}$. (b) Sub-threshold asymptotic potential and supra-threshold peak: $V_\infty < V_{th}$ & $\Delta_v \geq V_{th} - V_{rest}$. (c) Supra-threshold asymptotic potential and supra-threshold peak: $V_\infty \geq V_{th}$ & $\Delta_v \geq V_{th} - V_{rest}$.

neuron, but multiple volleys are required to fire. In terms of the aforementioned *actual* versus *utilised* synchrony, the input spike trains have very high degrees of actual synchrony, but the neuron utilises this synchrony in a different way. Instead, the high synchrony in a single volley is only utilised to bring the membrane potential up to a certain level below threshold, before subsequent volleys are integrated to raise the membrane potential above threshold and trigger a response.

4.2.5 Coincidence detection

Our results also show the necessary circumstances for a neuron to operate as a coincidence detector. Coincidence detection is measured by the NPSS when $\bar{M} \approx 1$. We know that coincidence detection only occurs for high values of S_{in} and in fact, pure coincidence detection only occurs when $S_{in} = 1$. This is fundamental to our interpretation of the continuum of operational modes. For any value $S_{in} < 1$, there are $N_{in}(1 - S_{in})$ random spike trains that contribute to the depolarisation of the membrane across each inter-spike interval. Therefore, a response spike cannot be fired as a result of coincident inputs only; there is always a degree of depolarisation that occurred before the coincidence window w and is part of the interval's integration period. By looking at each figure, we can determine how each of the other parameters and variables (σ_{in} , V_{∞} , Δ_v) affect the possibility of achieving pure coincidence detection. From Figure 4.2 it is apparent that coincidence detection only occurs when $\sigma_{in} = 0$ ms (see, Figure 4.2b), while Figure 4.3 shows that the asymptotic potential, V_{∞} , has no effect on the ability of a neuron to operate as a coincidence detector. Finally, Figure 4.4 shows that coincidence detection requires supra-threshold volley peaks, $\Delta_v \geq V_{th} - V_{rest}$. The above is summarised in Figure 4.5, which shows that coincidence detection occurs when $\Delta_v \geq V_{th} - V_{rest}$ and $\sigma_{in} = 0$ ms (dark blue points in lower-right corner), but V_{∞} can be either supra- or sub-threshold (Figures 4.5b and 4.5c). Examples of pure coincidence detection occurring can also be seen in the sample configurations in Figure 4.6. Figures 4.6c and 4.6d show two cases where the asymptotic potential is sub-threshold and the volley peak is supra-threshold ($V_{\infty} = 10$ mV, $\Delta_v = 20$ mV for Figure 4.6c and $V_{\infty} = 12.5$ mV, $\Delta_v = 25$ mV for Figure 4.6d) and in both cases, $\bar{M} = 1$ when $S_{in} = 1$, $D_S = 0$ and $\sigma_{in} = 0$ ms. The same holds for Figures 4.6e and 4.6f, which show two cases where the asymptotic potential and the volley peak are supra-threshold ($V_{\infty} = 25$ mV, $\Delta_v = 25$ mV for Figure 4.6e and $V_{\infty} = 20$ mV, $\Delta_v = 40$ mV for Figure 4.6f).

4.2.6 Effect of input parameter values (N_{in} , f_{in} , ΔV_s)

The figures for separate parameter configurations in Figure 4.6 illustrate how the individual parameters of the neuron and the input have little effect on the relationship between the NPSS and the SPIKE-distance and it is in fact the volley peak's relation to the spike threshold that mostly defines the relationship. The relationship of the asymptotic potential to the threshold mainly determines the number of spikes that are fired when there is little or no synchrony.

Cases where the asymptotic potential is sub-threshold and there is no synchrony simply produce very few spikes or none at all. Cases where no spikes were fired were discarded and therefore do not appear in any of the results. Each pair of figures in Figure 4.6 shows very similar behaviour between them, suggesting that specific parameter values (N_{in} , f_{in} , ΔV_s), or specific values for V_∞ & Δ_v are not important for the interaction between SPIKE-distance and NPSS, but the relative value of Δ_v towards the threshold is the most important factor.

4.2.7 Summary of results

In all figures, the horizontal axis represents the value of the average NPSS (\bar{M}) and the vertical axis represents the value of the integral of the SPIKE-distance over the entire duration of each simulation (D_S). The location of each point on a figure tells us the degree of synchrony in the input and how that synchrony is being utilised by the neuron. Points in the upper-left part of the plot denote cases where the input is mostly random (D_S is high) and the neuron is simply integrating random input spikes (\bar{M} is low). This is the most common case, since it occurs both when S_{in} is low and when S_{in} and σ_{in} are both high.

Points in the upper-right part denote cases where the input is mostly random (D_S is high) but the neuron is performing coincidence detection. This case is rare and is caused by the SPIKE-distance measuring low spike train synchrony due to a large number of random spike trains, but the NPSS measures coincidence detection because the neuron is responding primarily to coincidences. A clear example of this occurrence can be seen in Figure 4.6c, where $\bar{M} \approx 0.9$ and $D_S \approx 0.4$. Similar occurrences can also be seen in Figures 4.6d and 4.6f. In such cases, D_S is high due to the large number of random spike trains in the input increasing the spike time distance. Even though $S_{in} = 0.8$ & $\sigma_{in} = 0$ ms in the case of the aforementioned point in Figure 4.6c, which means that 80 % of spike trains are identical, the large value of $N_{in} = 200$ creates enough random spike trains for the SPIKE-distance to measure high overall spike distance. On the other hand, the NPSS measures a dominance of coincidence detection, due to the high synchrony driving the neuron. As mentioned above, this can also occur when the firing rate is very high and the two operational modes begin to converge, making them indistinguishable.

Even less common are points in the lower-left part of the plot. This area denotes cases where the input is highly synchronous (D_S is high) but the neuron performs mostly temporal integration (\bar{M} is low). Such cases can occur, theoretically, when the inputs are highly synchronised




FIGURE 4.6: NPSS (\bar{M}) vs SPIKE-distance (D_S) for specific configurations of N_{in} , f_{in} and ΔV_s , across the full range of synchrony parameters, $S_{in} \in [0, 1]$ (0.1 step size) and $\sigma_{in} \in 0$ to 4 ms (1 ms step size). The colour of each point shows the degree of applied jitter (σ_{in}). The dashed lines connect points which share the same value of σ_{in} and are ordered based on the value of S_{in} . Each subsequent point on a given line, starting from the top left, has a higher value of S_{in} . The values of the asymptotic potential and the volley peak for each subfigure are:

- (a) $V_\infty = 10$ mV and $\Delta_v = 10$ mV,
- (b) $V_\infty = 10$ mV and $\Delta_v = 10$ mV,
- (c) $V_\infty = 10$ mV and $\Delta_v = 20$ mV,
- (d) $V_\infty = 12.5$ mV and $\Delta_v = 25$ mV,
- (e) $V_\infty = 25$ mV and $\Delta_v = 25$ mV,
- (f) $V_\infty = 20$ mV and $\Delta_v = 40$ mV.

Therefore, each pair of figures correspond to one of the three cases shown in Figure 4.5: (a) and (b) correspond to the first case (Figure 4.5a) where both the asymptotic potential and the volley peak are sub-threshold, (c) and (d) correspond to the second case (Figure 4.5b) where the asymptotic potential is sub-threshold and the volley peak is supra-threshold, and (e) and (f) correspond to the third case (Figure 4.5c) where both the asymptotic potential and the volley peak are supra-threshold.

($S_{in} \approx 1$ & $\sigma_{in} \approx 0$), the volley peak is very low ($\Delta_v \ll V_{th} - V_{rest}$), but the input rate of each spike train (f_{in}) is very high. The result of this, is an extreme case of the aforementioned behaviour of *integrating coincidences*, where the coincidences (synchronous volleys) cause very small depolarisation (low peak), but they are frequent enough to drive the membrane potential above threshold and cause response spikes. In terms of operational mode, this is equivalent to temporal integration, even though the spike trains at the input are highly synchronised. The results in this chapter do not show occurrences of this case (e.g., points in the lower-left corner) since that would require extremely high (biologically implausible) input spike rates.

The fourth and final case is represented by the lower-right part of the plot, where the input spike trains are highly synchronised (D_S is low) and the neuron is performing coincidence detection (\bar{M} is high). As mentioned in Section 4.2.5, high input synchrony, which occurs when S_{in} is high and σ_{in} is low, causes the neuron to operate as a coincidence detector only when the volley peak is supra-threshold ($\Delta_v > V_{th} - V_{rest}$).

The following section (Section 4.3) discusses these cases further and relates them back to the results in the previous chapter (Chapter 3).

4.3 Discussion

This section serves as a discussion on the results of the NPSS as a whole, which includes the development and evaluation of the measure's behaviour in relation to the input synchrony parameters—presented in Chapter 3—and the results of the analysis of the measure's relationship to the SPIKE-distance—presented in this chapter.

The work presented in this thesis has so far focused solely on excitatory inputs which caused the neuron to fire a response, in order to infer the degree of response-relevant synchrony specifically. This simplifying choice was made to establish the viability of the slope of the membrane potential in inferring pre-synaptic synchrony. The potential inclusion of inhibitory inputs in our models would require our methods, namely the slope bound calculations, to account for the effects of inhibition on the range of potentials the membrane can acquire. More specifically, inhibition can drive the membrane potential below V_{rest} , which can cause pre-spike membrane potential slopes with higher values than the upper bound (see eqn. 3.4). Furthermore, slope values *smaller* than the lower bound are possible, albeit rare. This can occur in cases where the membrane potential increases rapidly, early during the ISI and remains very

close to the threshold until firing occurs. In such a case, the membrane potential at the start of the coincidence window, $V(t-w)$, could potentially be higher than the value calculated by the lower bound (see eqn. 3.5), since the lower bound assumes a *reverse exponential shaped curve* (i.e., $-e^{-x}$) that increases from V_{reset} at $t = 0$ s and hits V_{th} at $t = \Delta t_i$. For short ISIs, this value can be quite low compared to the threshold at $t = \Delta t_i - w$. In the cases where this occurred in our analysis, the normalised slope (which resulted in a negative number) was bound to 0, i.e., a *clipping* procedure was introduced after the values are rescaled by the normalisation procedure.

Generally, a neuron can operate in a sub- or supra-threshold input regime. The different input regimes are generally defined in terms of the asymptotic time-averaged membrane potential V_∞ (eqn. 4.8) in the absence of a threshold. If $V_\infty < V_{th}$, the neuron is operating in a sub-threshold regime and spikes are caused by fluctuations which can briefly drive the membrane potential above threshold. Conversely, if $V_\infty > V_{th}$, the neuron is operating in a supra-threshold regime and spikes are fired quite regularly and inevitably by the integration of inputs [16]. However, we can also define the total contribution of a fully synchronous ($S_{in} = 1, \sigma_{in}$) volley as ΔV_v (eqn. 4.9). Given these two properties, we can define four conditions, which based on their relationship with the firing threshold and separate 4 cases:

1. Case where $\Delta V_v < V_{th} - V_{rest}$ and $V_\infty < V_{th}$.

This is analogous to a true sub-threshold regime where spikes are fired only in cases where two or more volleys arrive close enough for their combined contribution to reach the threshold. In other words, output spikes are caused by fluctuations in the arrival times of sub-threshold volleys. When these relations hold, the NPSS will not produce a value of $M = 1$, even if the inputs are highly synchronous ($D_S = 0$), due to the contribution of each individual volley being sub-threshold as well as the lack of general supra-threshold input drive. Spikes in such cases are caused with very low probability and depend on the timing of individual volleys and spikes, i.e., the fluctuations in the input. Figures 3.3a, 4.6a and 4.6b corresponds to this case.

2. Case where $\Delta V_v > V_{th} - V_{rest}$ and $V_\infty < V_{th}$.

In this case, the NPSS will be able to achieve a value of 1 (if $D_S = 0$), regardless of the fact that $V_\infty < V_{th}$. Although this may correspond to a sub-threshold regime, with respect to the asymptotic potential, the presence of supra-threshold volleys makes firing of spikes a certainty,

as long as σ_{in} is small enough and we can therefore refer to it as a *supra-threshold volley regime*. Figures 3.3c, 3.3e, 4.6c and 4.6d correspond to this case.

3. Case where $\Delta V_v < V_{th} - V_{rest}$ and $V_\infty > V_{th}$.

In this interesting case, spikes are fired *almost surely* due to the neuron being in a supra-threshold regime, in the general sense (the asymptotic potential is supra-threshold), but the NPSS will never achieve $M = 1$, due to the contribution of each individual volley being sub-threshold. Therefore, although the input is supra-threshold, in terms of the contribution of single volleys it is operating in a *sub-threshold volley regime*. While the spikes within a single volley may coincide (highly synchronous volley, i.e., $D_S \approx 0$), the total dispersion between all the spikes that caused the neuron to fire is high. Instances of this case are rare, since the combination of a low volley peak and a supra-threshold asymptotic potential implies a very high rate of input (f_{in}), outside the range of biologically plausible values. However, the relationship between the NPSS and the SPIKE-distance in this case, is very similar to the behaviour of case 1, since the volley peak potential is the most important factor in determining the relationship between the two measures. Figure 3.3b corresponds to this case.

4. Case where $\Delta V_v > V_{th} - V_{rest}$ and $V_\infty > V_{th}$.

In this case, the mean drive is very high as well as the depolarisation caused by individual volleys. The behaviour of the NPSS is the same as for case 2, as the mean drive has little effect on the measure as long as ΔV_v is high enough to consistently cause a response. The NPSS however will behave unpredictably when the asymptotic potential is strong enough to cause very high firing rates, as has been already discussed (see Section 3.5.1). Figures 3.3d, 3.3f, 4.6e and 4.6f correspond to this case.

In summary, the relationship between input synchrony (D_S) and operational mode (\bar{M}) is more dependent on the relationship between ΔV_v and V_{th} , and is only slightly affected by the mean drive V_∞ . When volleys have a total contribution which is sub-threshold (cases 1 and 2) then $\bar{M} < 1$ even when $S_{in} = 1$, $\sigma_{in} = 0$ ms and therefore $D_S = 0$. In such cases, we use the term *integration of coincidence* to describe the behaviour of the neuron, which is an operational mode that is characterised by a balance of temporal integration and coincidence detection and lies somewhere in the middle of the operational mode continuum. The value of M (for a single inter-spike interval) reflects the total dispersion between all the contributing spikes, across multiple spike volleys, not the dispersion between spikes within a single volley. This reflects

the fact that the response can be caused by a number of volleys, each of which consists of completely synchronised spikes, but whose total inter-synchrony is much lower. This emphasises the difference between input synchrony in the traditional sense, which we measure with the SPIKE-distance, and the response-relevant input synchrony, which we measure with the NPSS. This reflects the same separation between *actual* and *utilised* synchrony, discussed in Section 3.7. The operational mode of a neuron is determined by the temporal dispersion of all the spikes that were responsible for the neuron's firing [31, 75].

The methods and results presented in this chapter have been submitted to a special issue on the *Application of mathematics in neuroscience* in the *Mathematical Biosciences and Engineering* journal and the manuscript is pending review.

Chapter 5

Input synchrony estimation for the Ornstein Uhlenbeck LIF

The Ornstein-Uhlenbeck (OU) model is a continuous approximation and a stochastic variant of the LIF [158–162], where the input is characterised by a stochastic Ornstein-Uhlenbeck process [150]. The behaviour of the membrane potential is described by the following stochastic differential equation:

$$dV(t) = \left(-\frac{V(t) - V_{rest}}{\tau_m} + \mu \right) dt + \sigma dW \quad (5.1)$$

where W is a standard Wiener process and μ and σ characterise the input. The input parameters μ and σ represent the mean and variance of the stochastic input respectively. The noisy input of the model is the diffusion approximation of many (tens of thousands) uncorrelated, Poissonian spike trains, where each spike contributes a small amount to the membrane depolarisation [163].

While this model does not add complexity over the LIF model, it allowed us to reformulate the stimulus to account for inhibition, which was lacking from the work up to this point and limited the applicability of the methods. In addition, the previous methods required defining a temporal window within which the pre-spike slope was averaged for each response spike. This was necessary due to the trajectory of the LIF's membrane potential being discontinuous. The length of the window affected the interpretation of the results, since any spikes arriving within that time prior to a response spike were regarded as completely synchronous by the method.

Symbol	Description	Value or range	Units
τ_m	Membrane leak time constant	10	ms
V_{th}	Firing threshold	10	mV
V_{rest}	Neuron resting and reset potential	0	mV
μ_a	Amplitude of $\mu(t)$ signal	0.2 to 2.0	mV/ms
μ_0	$\mu(t)$ offset (baseline)	0.2 to 2.0	mV/ms
σ_a	Amplitude of $\sigma(t)$	0 to 1.0	mV/ $\sqrt{\text{ms}}$
σ_0	$\sigma(t)$ offset (baseline)	0 to 1.0	mV/ $\sqrt{\text{ms}}$
f	Frequency of sine waves	5 to 20	Hz
h	Simulation time step	0.1	ms

TABLE 5.1: Parameter symbols and value ranges.

5.1 Operational mode of the Ornstein Uhlenbeck LIF

Due to the nature of the input of the OU model, simulating synchronous volleys of varying intensity required that we define μ and σ as time-dependent functions that replicate such input. The input was therefore modelled using sine waves to describe both $\mu(t)$ and $\sigma(t)$ (eqns. 5.2 and 5.3).

$$\mu(t) = \mu_0 + \mu_a \sin(2\pi t f) \quad (5.2)$$

$$\sigma(t) = \sigma_0 + \sigma_a \sin(2\pi t f) \quad (5.3)$$

where μ_0 and σ_0 are the sine wave offsets, μ_a and σ_a are the amplitudes and f is the frequency of the sine wave envelope. The final model has the following form:

$$dV(t) = -\frac{(V(t) - V_{rest})}{\tau_m} dt + \mu(t) dt + \sigma(t) dW \quad (5.4)$$

This way, by controlling the amplitude, frequency and offset of the waves, we can approximate the effect of periodic volleys of increased activity, which arise from synchronous volleys of spikes (see Section 5.1.1). The peaks of $\mu(t)$ represent the increased mean activity of the input, while the peaks of $\sigma(t)$ represent the increased variance of the input signal which accompanies a spike volley. The two waves are always synchronised and the offset of sigma must always be greater or equal to the amplitude ($\sigma_0 \geq \sigma_a$) in order to avoid negative $\sigma(t)$ values, as that would produce high noise at the low $\mu(t)$ phases, due to σ acting as an absolute factor on the Wiener process (i.e., positive and negative σ values have the same effect). Parameter descriptions and values used in this work can be found in Table 5.1.

It is useful to define two ranges for the values of the input signal $\mu(t)$, with respect to the threshold V_{th} . *Supra-threshold* input refers to the input activity when $\mu(t)\tau_m \geq V_{th}$, at a given time t . At these values, the neuron will fire even in the absence of noise. On the other hand, *sub-threshold* input refers to the input activity when $\mu(t)\tau_m < V_{th}$. Even though in these cases the neuron never fires in the absence of noise, noise-induced firing may occur.

Our approach begins by determining the frequency of the oscillations of the input rates. Knowing the frequency of the input sine wave allows us to identify the intervals of low and high activity, which represent the background and synchronous input intervals respectively. By applying initially existing input estimation methods [149], we obtain a discrete, yet noisy, reconstruction of the input signal and noise. The values at each interval of the period of the input can then be more accurately determined by calculating the mean of the estimated values that fall within the same interval of the input signal period.

5.1.1 Comparison of LIF with OU model

We initially compared the behaviour of the two models—the LIF used in previous chapters and the OU used in this chapter—to determine how well the sinusoidal input of the OU approximates the synchronous input spike trains of the LIF. In order to determine the range of parameters and conditions in which the OU model with sinusoidal inputs can reproduce the behaviour of the LIF driven by synchronous input volleys, we derived the relationship between the models' parameters and validated our results through simulations.

Due to the nature of the inputs used in this chapter, i.e., sinusoidal drive and noise, the corresponding inputs to the LIF needed to contain periodic synchronous volleys with a wide dispersion (high σ_{in} , see eqn. 5.5). This high value was necessary to create a (short) overlap between the tails of the distributions that are centred on each volley peak. To create a continuous sequence of volleys aligned with the peaks of the corresponding sine wave, the synchronous volley times (t_S) were calculated based on the input sine wave frequency (eqn. 5.6). The input spike trains needed to be separated into two groups, the Poisson inputs that drive the membrane potential to a baseline voltage and the synchronous inputs that fire in volleys. In order to control the relative contribution of these two groups, we can manipulate two parameters: (i) the number of spike trains in each group (N_P for the Poisson group and N_S for the synchronous group) or (ii) the depolarisation per spike (i.e., the synaptic weight) from each group (ΔV_P and ΔV_S respectively). For simplification, we chose to keep the number of spike trains

constant and manipulate only the weights. This choice was made to allow for inhibitory inputs, where necessary, which we would achieve with negative weights. The Poisson input group acts as a background voltage which keeps the membrane almost constantly depolarised (or hyperpolarised). The background input equivalent in the OU model is the minimum value of the sinusoidal drive ($\mu_0 - \mu_a$) and so the weight of the Poisson inputs (ΔV_P) is analogous to this value (eqn. 5.7). The rest of the input spike trains (ΔV_S) make up the synchronous inputs and their weight is analogous to the sine wave amplitude (eqn. 5.8).

$$\sigma_{in} = (5f)^{-1} \quad (5.5)$$

$$t_S = \frac{n + 0.25}{f} : n \in \mathbb{Z} \quad (5.6)$$

$$\Delta V_P = \frac{\mu_0 - \mu_a}{N_P f} \quad (5.7)$$

$$\Delta V_S = \frac{\mu_a}{N_S f} \quad (5.8)$$

All other parameters ($V_{rest}, V_{reset}, \tau_m$) are common between models and were naturally set to the same values in each case.

The validation procedure was performed in two steps. First we removed the threshold (i.e., the neurons did not spike) and we determined the equivalence of the two models by measuring the difference between the membrane voltage traces. Looking at the membrane potential in the absence of spiking was necessary in order to avoid large membrane potential deviations caused by mismatched spike times. The difference between the two traces was calculated using two values, the maximum deviation (D_V , eqn. 5.9) and the root mean square (RMS) of the difference at each simulated time step (R_V , eqn. 5.10).

$$D_V = \max_h |V_{LIF}(h) - V_{OU}(h)| \quad (5.9)$$

$$R_V = \sqrt{\frac{1}{H} \sum_{h=1}^H (V_{LIF}(h) - V_{OU}(h))^2} \quad (5.10)$$

where h indicates the iteration through simulation time steps and H is the total number of time steps.

For the second step we set the threshold to several different values (below the maximum value of the membrane potential) and the SPIKE-distance was measured, using the spike train distance measure described in Chapter 4.

The results of our validation indicate that the derived relationships between the input parameters are accurate, as long as the sinusoidal noise is not very high ($\sigma(t) < 1.0 \text{ mV}/\sqrt{\text{ms}}$). When the sinusoidal noise was low, the maximum RMS difference was 1.6 mV and the maximum deviation from all simulations was 4 mV (approximate values). Note that the maximum depolarisation of the membrane in these cases was 35 mV (for the max RMS) and 25 mV (for the max deviation). When the threshold was enabled, causing spiking, the SPIKE-distance was no larger than 0.14, which is very low when compared to the ranges studied in Chapter 4. Higher values of $\sigma(t)$, caused by either large σ_a or σ_0 , cause the simulations to deviate substantially, since the nature of the inputs we use for the LIF does not allow for controlling of the noise in the same manner as in the OU model. It is theoretically possible to approximate this type of noise by reducing the number of inputs to the LIF (N_P and N_S), which would result in an increase in input weights (eqns. 5.7 and 5.8). With this change, individual input spikes would cause larger fluctuations of the membrane potential, which should produce behaviour similar to the potential of the OU model with high noise parameter values. This high degree of accurate reproduction is beyond the scope of the current section however. We only aim to show that an approximate reproduction of the behaviour of the LIF driven by synchronous inputs and Poisson background spike trains is possible using an OU model driven by sinusoidal mean drive and noise.

Figure 5.1 shows two examples of the similarity between the two models. For each example, we show the membrane potential trace with and without the threshold enabled, totalling four comparisons. One example was chosen for having high similarity (small differences), by keeping the noise in the OU model low, while the other was chosen to show the effects of high noise on the similarity between the two models. The parameter values of each case and the measured errors (D_V , R_V , and D_S) are listed in the figure's caption.

5.2 Estimation of frequency

The frequency of the sine wave f can be estimated using the power spectral density (PSD) of the spike train [164]. The PSD is calculated as the Fourier transform of the binned autocorrelation of the spike times. The frequency component of the PSD with the highest power (the spectral peak) provides estimates of the underlying frequency of the sine wave.




FIGURE 5.1: Membrane potential traces for two configuration pairs. Each configuration pair consists of one simulated run without a firing threshold ((a) & (c)), in order to compare the voltage trace directly, and a simulated run with the firing threshold enabled ((b) & (d)), in order to compare the spike train distance. In all figures, the blue trace represents the voltage of the OU model and the green trace represents the voltage of the LIF model. The parameters of the LIF model were derived from the parameters of the OU using the equations described in Section 5.1.1.

(a) Parameters: $\mu_a = 0.5$ mV/ms, $\mu_0 = 1.0$ mV/ms $\sigma_a = 0.1$ mV/ $\sqrt{\text{ms}}$, $\sigma_0 = 0.1$ mV/ $\sqrt{\text{ms}}$
 $f = 10$ Hz, no threshold.

Voltage trace differences: $D_V = 1.4$ mV, $R_V = 0.5$ mV

(b) Parameters are the same as (a) but with $V_{th} = 10$ mV.

Spike time distance: $D_S = 0.035$

(c) Parameters: $\mu_a = 0.5$ mV/ms, $\mu_0 = 1.0$ mV/ms $\sigma_a = 0.1$ mV/ $\sqrt{\text{ms}}$, $\sigma_0 = 0.5$ mV/ $\sqrt{\text{ms}}$
 $f = 20$ Hz, no threshold.

Voltage trace differences: $D_V = 4.4$ mV, $R_V = 1.2$ mV

(d) Parameters are the same as (c) but with $V_{th} = 10$ mV.

Spike time distance: $D_S = 0.06$

The first configuration pair ((a) & (b)) has low noise (σ) values and the differences between the voltage traces as well as the spike times are therefore very small. On the other hand, for the second configuration pair ((c) & (d)) the differences between the voltage traces and spike times are relatively high, due to the high σ values used in the OU model.

The estimation of the frequency assumes that the input is periodic, with constant frequency throughout the duration of the observed data. In addition, the binning introduced during the calculation of the autocorrelation introduces an implicit rounding in the estimated frequency values. A side effect of this rounding is the removal of any small errors that may have occurred in the frequency estimation had the precision been finer. For the results in this paper, we used a bin width of 1 ms, which made bins small enough to include at most one spike per bin.

5.3 Estimation of μ

Figure 5.2 shows the steps of the estimation of $\mu(t)$. The noisy input signal (Figure 5.2a) produces spikes at the output of the neuron (Figure 5.2b). From the spike train and the discrete observations of the membrane potential, we can obtain estimates of the input signal $\mu(t)$ (Figure 5.2c, eqn. 5.11). The estimated values are aligned to a single period (Figure 5.2d, eqn. 5.12) and subsequently binned and averaged (Figure 5.2e, eqn. 5.13), to produce the final estimates of the input signal.

The estimation of the parameters for $\mu(t)$ relies on estimation methods found in the literature. More specifically, we use the μ estimation method in Bibbona et al. [149]:

$$\hat{\mu}_i = \frac{V_{th}}{\tau_m K_i (1 - e^{-\frac{h}{\tau_m}})} + \frac{1}{\tau_m K_i} \sum_{k=1}^{K_i-1} V_k^i \quad (5.11)$$

where $K_i = T_i/h$, T_i is the i th ISI and h is the discretisation step (i.e., the simulation time step). V_k^i is the discrete membrane potential trace during the i th ISI.

By applying eqn. 5.11 to the entire simulation's membrane potential and spike train, we can calculate $\hat{\mu}_i$ for each ISI during a simulation, which results in a sequence of tuples $X = \{(t_i, \hat{\mu}_i)\}$, where t_i is the time of the i th spike and μ_i is the average membrane potential during the time interval of the i th ISI (Figure 5.2c). The sequence X represents a discrete reconstruction of $\mu(t)$, with discretisation bins having the length of each respective ISI.

Using the estimated frequency of the input sine wave \hat{f} , each $\hat{\mu}$ is aligned to the start of their respective period:

$$G = \{(\phi_i, \hat{\mu}_i) | \phi_i \equiv t_i \pmod{\hat{f}^{-1}} \wedge (t_i, \hat{\mu}_i) \in X\}. \quad (5.12)$$

The set G groups all estimated values into a single period length, where ϕ_i is the period-aligned spike time (Figure 5.2d). The period is separated into N bins and the estimated values in each bin are averaged to produce a sequence of average binned estimated values:

$$\hat{\mu}_b = \frac{\sum_i \hat{\mu}_i [b/(\hat{f}N) < \phi_i \leq (b+1)/(\hat{f}N)]}{\sum_i [b/(\hat{f}N) < \phi_i \leq (b+1)/(\hat{f}N)]} \quad (5.13)$$

where $b \in \{0, 1, \dots, N-1\}$ and the square brackets are Iverson brackets [165] yielding 1 if the condition in square brackets is satisfied and 0 otherwise.

The binning and averaging removes estimation inaccuracies arising from the input noise $\sigma(t)$ (Figure 5.2e).

5.3.1 Estimation of μ_p

The sequence of $\hat{\mu}_b$ values represents a noiseless, partial reconstruction of a single period of the input signal $\mu(t)$. From this, it is straightforward to extract the peak of the input signal $\mu_p = \mu_0 + \mu_a$ by taking the maximum of the $\hat{\mu}_b$ values.

5.3.2 Estimation of μ_0

As the aim of this work is to estimate the degree of input synchrony, it is important to know how high the peak activity μ_p is, compared to the baseline value. In other words, the amplitude of the signal μ_a represents the increase in input spike rates during a synchronous volley, or a brief synchronised increase in input rates.

For supra-threshold baseline values ($\mu_0\tau_m > V_{th}$), $\hat{\mu}_0$ can be estimated from the mean of the reconstructed period:

$$\hat{\mu}_0 = \frac{1}{N} \sum_{b=0}^{N-1} \hat{\mu}_b \quad (5.14)$$

Applying the same method to cases where $\mu_0\tau_m < V_{th}$ is not as reliable (see Section 5.5).




FIGURE 5.2: μ estimation procedure. (a): Input signal with noise, i.e., $\mu(t) + \sigma(t)W(t)$. (b): Output spike train of the neuron. (c): Estimated μ values resulting from the estimator given by eqn. 5.11. (d): The same estimated values from panel C, aligned to a single period length, using eqn. 5.12. (e): The averages of the estimates from panel D (dots connected with solid line) that fall within the same bin, as described in eqn. 5.13. In all panels, the dashed grey line represents the original, noiseless input signal $\mu(t)$.

5.4 Estimation of σ

While our focus is on input synchrony estimation, which involves the changes in the input population's firing rate $\mu(t)$, the same procedure used in the estimation of the μ parameters can be applied to the corresponding noise σ parameters. We therefore use a similar procedure to estimate the parameters of $\sigma(t)$ (σ_0 and σ_a) based on the σ^2 estimation method in Bibbona et al. [149]:

$$\hat{\sigma}_i^2 = \frac{2}{K_i - 1} \sum_{k=1}^{K_i} \frac{[V_k^i - \hat{\mu}_i \tau_m + (\hat{\mu}_i \tau_m - V_{k-1}^i) e^{-h/\tau_m}]^2}{\tau_m (1 - e^{-2h/\tau_m})} \quad (5.15)$$

where the symbols have the same meaning as the ones of eqn. 5.11.

eqn. 5.11 (which estimates individual μ values) is replaced by eqn. 5.15 (which estimates individual σ^2 values) and the rest of the procedure (i.e., eqn. 5.12 eqn. 5.13) is followed again, with the square root of the σ^2 estimates replacing their μ counterparts.

5.4.1 Estimation of σ_p and σ_0

For obtaining $\hat{\sigma}_p$ and $\hat{\sigma}_0$ from the reconstructed trajectory of σ we follow the same procedure as for $\hat{\mu}$. The maximum value corresponds to estimates of the peak σ_p , while the average provides estimates for the baseline σ_0 .

5.5 Results

Simulations were run using the input parameter values and value ranges shown in Table 5.1. Parameter combinations were chosen such that data from all valid configurations were generated. Discarding simulations which did not produce enough output spikes (i.e., when the output rate is less than 1 Hz) resulted in data from 14048 simulations, which were run for 5 s of simulated time each.

Parameter values covered the supra- and sub-threshold ranges, for both the baseline of the signal μ_0 and the peak μ_p . As expected, cases where the peak of the signal is sub-threshold (i.e., $\mu_p\tau_m < V_{th}$) were much less likely to produce spikes, since threshold crossings in such cases are dependent on the amplitude of the noise. Only 564 of the simulations have a sub-threshold μ_p .

For each estimated parameter, we calculate the relative estimation error as $|x - \hat{x}|/x$ and report the mean relative error ε_x , averaged over all simulations which share the same parameter value. Where the correct value is zero ($x = 0$), the error is taken to be the value of \hat{x} .

Average estimation errors are summarised in Table 5.2, separated into sub-threshold signal peak ($\mu_p\tau_m < V_{th}$), supra-threshold peak with sub-threshold baseline ($\mu_0\tau_m < V_{th} \leq \mu_p\tau_m$) and supra-threshold baseline ($\mu_0\tau_m \geq V_{th}$).

Parameter	ε_x for $\mu_p\tau_m < V_{th}$	ε_x for $\mu_0\tau_m < V_{th} \leq \mu_p\tau_m$	ε_x for $\mu_0\tau_m \geq V_{th}$	ε_x (overall)
μ_p	0.33	0.10	0.05	0.08
μ_0	0.35	0.43	0.07	0.20
μ_a	0.89	0.24	0.19	0.28
σ_p	0.19	0.09	0.05	0.07
σ_0	0.27	0.18	0.08	0.12
σ_a	0.29	0.25	0.14	0.18

TABLE 5.2: Parameter estimation errors (relative) for sub-threshold signal peak, supra-threshold peak with sub-threshold baseline and supra-threshold baseline. The number of simulations for each range is 564 (4%) for $\mu_p\tau_m < V_{th}$, 4844 (34%) for $\mu_0\tau_m < V_{th} \leq \mu_p\tau_m$ and 8640 (62%) for $\mu_0\tau_m \geq V_{th}$ (total simulations: 14048).




FIGURE 5.3: Histograms of (a) relative and (b) absolute estimation errors from ~ 3500 realisations of the frequency estimator for each value of $f \in \{5, 10, 15, 20\}$ (in Hz). The total number of simulations was 14048. Note that due to the very low number of cases with nonzero estimation error, the vertical axis has a logarithmic scale. The μ and σ parameters (baselines and amplitudes) varied across their entire ranges, found in Table 5.1. It should be noted that the number of misestimated instances is only 66 and the average relative estimation error is $\varepsilon_f < 0.01$.

5.5.1 Frequency estimation results

Figure 5.3 shows the results of the frequency estimation. Figure 5.3a shows a histogram of the relative frequency estimation errors ($\varepsilon_f^r = |f - \hat{f}|/f$) and Figure 5.3b shows a histogram of the absolute estimation errors ($\varepsilon_f^a = |f - \hat{f}|$). The histograms are shown using a logarithmic scale due to the very low number of nonzero estimation errors.

The average relative estimation error is $\varepsilon_f < 0.01$. Of the 14048 samples used, only 66 had a non-zero estimation error. The high amount of perfect estimates is a result of the rounding that occurs during the calculation of the PSD, as mentioned in Section 5.2, which removes any




FIGURE 5.4: The triangles represent μ_0 and μ_a values for simulations where the frequency estimation failed. The frequency for simulations with values represented by the points was estimated perfectly.

small deviations from the true value that may have occurred if a smaller bin size had been used. The small number of misestimations deviate substantially from the correct values.

It is important to understand when misestimations occur and to be able to predict whether it is possible to estimate the underlying frequency with no prior knowledge of the input. From our investigation of the simulated data, it became apparent that the frequency estimation failed when the amplitude of the input sine wave μ_a was very small compared to the baseline μ_0 (see Figure 5.4).

However, since prior knowledge of the input parameters is not assumed, the prediction of whether the frequency estimation will succeed or not, should rely on the effects that the combination of μ parameters have on the firing of the neuron. The high- μ_0 /low- μ_a combination does not allow for the underlying frequency to be estimated because in such cases, the fluctuations of the input sine wave are very small (low amplitude). The firing becomes very regular and does not reflect the subtle variations in the input firing rates. We can use the firing variability, commonly measured using the coefficient of variation (CV) of the ISIs, to test whether the estimation of the frequency of the input fluctuations (if any) has a high probability of failing. Figure 5.5 (top) shows the frequency estimation error of each simulation ε_f as a function of its CV. Figure 5.5 (bottom) shows the probability of misestimated frequencies conditioned




FIGURE 5.5: Top: CV of the firing ISIs vs the frequency estimation error. Bottom: Normalised (conditional) CV distribution of simulations with non-zero ε_f . High relative estimation errors occur only at low CV values (regular firing) and with higher probability in the range [0.1–0.3], which are caused by the input sine wave having very subtle fluctuations around its baseline value.

on the CV. The figures clearly show that not only is the frequency estimation error larger for CVs between 0.1 and 0.3, but misestimations are also more probable in that same range.

5.5.2 $\hat{\mu}$ estimation results

Using the estimated frequencies \hat{f} , we began the reconstruction of the input signal $\hat{\mu}_i$, as described in Section 5.3. Figure 5.6 shows the result of applying eqn. 5.11 to a single simulation's generated data (see caption for parameter values). This corresponds to Figure 5.2c of the illustrated methodology.

The data shown in Figure 5.6 has been truncated to 1 second for clarity. Figure 5.7 shows the application of eqn. 5.12 to the $\hat{\mu}_i$ data shown in Figure 5.6, which corresponds to Figure 5.2d.

The data is then binned and averaged (eqn. 5.13) to produce Figure 5.8, which corresponds to Figure 5.2e. We used $N = 10$ bins for all data in this study. The number was chosen such that the bins were large enough to contain enough spikes for the averaging process, but small enough to avoid containing both the high and low parts of the sine wave. In addition, it is best




FIGURE 5.6: $\hat{\mu}_i$ for a single simulation with parameters $\mu_0 = 1.6$ mV/ms, $\mu_a = 0.4$ mV/ms, $\sigma_0 = 0.4$ mV/ $\sqrt{\text{ms}}$, $\sigma_a = 0.4$ mV/ $\sqrt{\text{ms}}$ and $f = 5$ Hz. The dots connected with the solid line represent the estimated trajectory of the input signal, while the dashed grey line is the actual input sine wave. The results in this figure correspond to the step in the estimation procedure shown in Figure 5.2c.

if the number of bins is chosen such that $N/2$ is odd, so that there is a bin centred on each peak (positive and negative) of the sine wave. This ensures that all spikes around the peak of the sine wave fall within one bin, and slightly increases the estimation accuracy for $\hat{\mu}_p$. From the binned data, we can extract the estimates for μ_p ($\max(\hat{\mu}_b)$, Figure 5.9a), μ_0 (eqn. 5.14, Figure 5.9b), and subsequently estimate $\hat{\mu}_a = \hat{\mu}_p - \hat{\mu}_0$ (Figure 5.9c).

The average relative error between the actual and estimated values for the peak of the input signal is $\varepsilon_{\mu_p} = 0.08$. The average relative estimation error for the baseline value is $\varepsilon_{\mu_0} = 0.2$. It should be noted however that this includes supra- and sub-threshold values of μ_0 . Taking each case individually, the estimator is very reliable for supra-threshold values $\varepsilon_{\mu_0 \tau_m \geq V_{th}} = 0.07$ but as expected, it has a very high error for sub-threshold values $\varepsilon_{\mu_0 \tau_m < V_{th}} = 0.43$. Since the estimation of the amplitude relies on all other estimated parameters, each estimated value is affected by all previous errors. Its average relative error is $\varepsilon_{\mu_a} = 0.28$ (see Table 5.2 for full results).




FIGURE 5.7: $\hat{\mu}_i$ grouped into a single period length based on \hat{f} . This is the same data used for Figure 5.6. The results in this figure correspond to the step in the estimation procedure shown in Figure 5.2d.

5.5.3 $\hat{\sigma}$ estimation results

Figures 5.10a–5.10c show the estimates $\hat{\sigma}_p$, $\hat{\sigma}_0$ and $\hat{\sigma}_a$ for the parameter ranges shown in Table 5.1. The overall, average estimation errors are $\varepsilon_{\sigma_p} = 0.07$, $\varepsilon_{\sigma_0} = 0.12$ and $\varepsilon_{\sigma_a} = 0.18$. As expected, the estimates were more reliable when the input signal $\mu(t)$ was supra-threshold at the peak (i.e., $\mu_p \tau_m \geq V_{th}$) and especially when the baseline of the signal was also supra-threshold (i.e., $\mu_0 \tau_m \geq V_{th}$). Table 5.2 shows the full set of estimation errors.

5.5.4 Frequency misestimation

Since all parameter estimates rely on the frequency, in order to align estimated values to the length of a single period, we also studied the effect that frequency misestimations have on the subsequent estimates. Figure 5.11 shows ε_{μ_p} , the average estimation error for μ_p , given a frequency misestimation α_f for all samples. μ_p estimates were calculated using incorrect frequencies calculated as $(1 + \alpha_f)f$.




FIGURE 5.8: $\hat{\mu}_b$: the binned average of the estimates seen in Figure 5.7 for $N=10$ bins, i.e., the width of each bin is 20 ms. The results in this figure correspond to the step in the estimation procedure shown in Figure 5.2e.

As Figure 5.11 shows, ε_{μ_p} is very sensitive to errors in the estimation of the frequency. This signifies the importance of accurately identifying the intervals of synchronous activity, as frequency misestimations disrupt the averaging of estimates that occur within the same intervals of the period.

5.6 Discussion

Our methods rely on estimating the frequency of the periodic synchronised activity from the power spectrum of the fired spikes, which allows us to align estimates based on the length of the period. As mentioned in Section 5.1, the methods presented in this chapter rely on existing estimation methods for stationary inputs. The existing methods rely on the membrane potential trace during an inter-spike interval [149]. Using these estimators, we initially obtain a noisy discrete reconstruction of the supra-threshold input signal. We adapt the estimators to time-varying (sinusoidal) inputs by binning the response spikes and aligning bins based on the input period. This may also extend to sub-threshold ranges in cases where noise-induced




FIGURE 5.9: Estimated vs actual μ parameters. (a): μ_p estimates. $\varepsilon_{\mu_p} = 0.08$. (b): μ_0 estimates. $\varepsilon_{\mu_0} = 0.20$. (c): μ_a estimates. $\varepsilon_{\mu_a} = 0.28$.

firing occurs, with less accuracy. From the reconstructed input we subsequently obtain accurate estimates of the input signal and noise by averaging individual estimates that fall within corresponding intervals of each period. The use of estimators for stationary inputs causes the method to assume that the input within each bin is stationary, which provides a good approximation as long as the bin width is small enough.

Even though the binning and alignment of input estimates is similar to the methodology presented in [151], in their work they assumed the sine wave frequency (f) was known and their model only used stationary noise (σ). These authors binned the spike intervals of the neuron, with each bin representing a discrete phase shift of the input signal, and represented all points within each bin by the midpoint of their respective bin. One further fundamental difference of our work, however, is that our binning procedure is performed after the estimation of the input signal parameters and the binning is used to align and obtain averages of the estimates. The estimation procedure, in our case, is performed on the original spike time intervals.

The accuracy of our method is based on several underlying assumptions regarding the model




FIGURE 5.10: Estimated vs actual σ parameters. (a): σ_p estimates. $\varepsilon_{\sigma_p} = 0.07$. (b): σ_0 estimates. $\varepsilon_{\sigma_0} = 0.12$. (c): σ_a estimates. $\varepsilon_{\sigma_a} = 0.18$.

and its input. We simulate periodic synchronous input volleys using a sinusoidal signal. Our method therefore assumes that the input signal can be accurately described by a sine wave and also that this sine wave represents a sudden increase in synchronous excitatory input firing rates. The requirement that the oscillations are caused by an increase in excitation can be tested by deducing the excitatory and inhibitory input rates from the estimated μ and σ values [34, 166]. More generally, it is important to note that synchronous inputs are not always accurately represented by sinusoidal oscillations. The sine wave is an appropriate model for periodic synchronous volleys of excitatory inputs, with relatively large variance between spike times within the volley. When spike times are highly correlated however [see for example Figure 1c, in 33], the membrane potential responds with near-instantaneous jumps. Additionally, the model's explicit threshold function assumes that the threshold is constant, which is a common assumption when working with models of reduced complexity, like the LIF. The method we used to estimate the input parameters assumes a constant input signal across an entire spike train. By using this method to estimate the parameters for individual inter-spike intervals, we implicitly assume that the input during each interval can be approximated by a constant




FIGURE 5.11: ϵ_{μ_p} vs α_f . The average estimation error for the peak of the input, assuming the frequency has been misestimated. α_f is a fixed error ratio applied to all samples in the data.

current with added constant noise.

Our frequency estimation method requires that the amplitude of the input spike rates μ_a is large compared to the baseline input rate μ_0 , so that the observable fluctuations in the firing frequency reflect the input sine wave. This requirement can be tested by measuring the CV of the neuron's firing. Lower CV values increase the probability of the estimation failing, since regular firing provides no information about input rate changes.

It is interesting to note that the sinusoidal nature of our inputs causes stochastic resonance to occur in our model. Stochastic resonance in neurons occurs when a sub-threshold input signal ($\mu(t)\tau_m < V_{th}$) causes firing due to the fluctuations caused by the added noise ($\sigma(t)$) (i.e., noise-induced firing), when in fact firing would not occur had the signal been noiseless [167, 168]. From our results (see Table 5.2) we can see that firing does indeed occur when the signal is sub-threshold, but the estimation accuracy in such cases is low.

Alternative input parameter estimation methods, such as the Bayesian method developed by

Kobayashi et al. [34, 169] and the state-space method by Kim and Shinomoto [170] do not require estimating the frequency of the input sine wave in the case of sinusoidal inputs, but they do assume that the input variability is low. When estimating inputs that vary significantly, the estimation error increases with the variability (e.g., amplitude of the sine wave). In our case, the estimation error is lower for larger amplitudes and in fact, more variable input causes higher output variability, which increases the accuracy of the frequency estimation (see Figure 5.5). In the course of exploring the model and methods, it was not possible to find a way to estimate the input signal with the same level of accuracy without a frequency estimation first. We may therefore consider our method as an alternative to all other methods and complementary to the Bayesian method [34] which estimates the input parameters accurately when no assumptions about the input's variability can be made, or when the input is known to be smooth. Our method however provides a good input estimation accuracy when the input is known to be sinusoidal with large amplitude and high variability and when the aforementioned assumptions hold.

The dependence of the method's accuracy on the underlying input sine wave frequency may hint at an underlying phenomenon. It may suggest that information in the brain is transmitted via synchronous regular signals. This is likely related to the idea that information is encoded in spike times relative to the oscillatory signal [49]. In other words, the inability to accurately reconstruct the input signal without knowing the frequency of the oscillations, may suggest that the information in the signal itself depends on this frequency.

This method may be considered an extension of the NPSS as well as an attempt to address some of its shortcomings. The NPSS was limited by the fact that the effects of inhibition were not accounted for. The current measure overcomes this limitation by treating input synchrony as a brief increase in the firing rates of the summed input population. Inhibition is implicit in the sinusoidal input, which models the pre-synaptic population activity as a single input signal. However, the input estimator described in this chapter is not a direct estimator of the operational mode, as it estimates actual input synchrony, not utilised input synchrony. The degree of utilised input synchrony can be inferred from our results by observing the firing times with respect to the period of sinusoidal input signal.

By estimating the baseline and peak values of the input rates (μ_0 and μ_p) we effectively obtain a measure of the synchronous modulations in pre-synaptic spike rates. Our previous work also required the use of a temporal window that directly preceded a spike, within which all input

spikes were regarded as coincident. This required determining (or making an assumption for) the timescale on which a neuron processes information, as any temporal information was lost within the window length. The method presented here requires only one parameter selection, the number of averaging bins N , the effect of which is negligible on the results and has no effect on their interpretation, as long as the bin size is small enough to not contain both the high and the low parts of the sine wave. As mentioned above, it is good to select N such that $N/2$ is odd, in order to have a bin centred on the peak of the sine wave. Given this requirement, we suggest that at least $N = 6$ should be used.

In the case of non-periodic inputs, our estimates would introduce large amounts of error, since the methodology assumes periodic input when grouping spikes based on their relative position in the sine wave period. In such a case, the frequency estimator could be replaced by a peak detector that could detect peaks in the firing rate of the neuron. Using such a method, the spike intervals could then be grouped based on their relative timing within each inter-peak interval, rescaled to a single, common length. This would again assume that the input during each interval is sinusoidal, but with varying frequency. If, on the other hand, this assumption does not hold, the individual estimates of μ_i could be grouped based on the instantaneous firing rate, such that estimated values that fall within periods with similar firing rates are grouped together and averaged.

Although our focus is on estimating the degree of synchrony at the input, which is represented by the magnitude of change of $\mu(t)$, we can use the same method to estimate the noise $\sigma(t)$. In our model, we use noise to represent variability in firing rates of the input population. However, since noise in neural systems can induce synchronisation [171], the ability to estimate the noise driving a neuron, or population, may be useful in studying this phenomenon.

The methods and results presented in this chapter were published in *Brain Research* in 2013 [37].

Chapter 6

Can a biophysical model of a pyramidal neuron learn time delays between the spikes from different input neurons?

In this chapter, we investigate the ability of a pyramidal neuron to learn to distinguish, with high precision, correlated input spikes of varying delays. Our goal is to discover a learning mechanism that allows the cell to reinforce pairs of synapses that activate with very specific time delays between them (on the millisecond scale), while not reinforcing other synapses on the same dendritic tree. A side effect of our analysis is the measurement of the precision with which such a mechanism can function and an investigation into the properties that define or affect this precision.

The detection of sequences of sensory inputs with specific short time delays (e.g., velocity sensitive motion detection or decoding of the firing of Geniculate lagged cells, see [172]) is a function of biological systems. Sequence detectors are usually modelled as coincidence detectors that exploit appropriate delays of asynchronous individual input to cause a coincidence after the arrival of the last input of the sequence (see for example [173]). Given the adaptability of neural systems, the question arises as to whether learning mechanisms exist that develop appropriate coincidence detectors and then stabilize them during use.

The widely used Spike Timing Dependent Plasticity (STDP) [174–178] learning rule normally requires the post-synaptic neuron to fire a spike and will reinforce all synapses with inputs arriving shortly before that spike. Synapses on distant dendrites whose earlier inputs also contribute to the spike undergo a much weaker reinforcement than proximal dendrites and end up disappearing when resource limitations are considered in the model, as proposed by Letzkus et al. [179]. Branco et al. [173] have shown that, on the contrary, synapses at various distances from the soma stay strong and contribute to sequence-specific neuronal responses. They did that by activating a succession of synapses by optical uncaging and noted that if the uncaging sequence moves from distal to proximal synapses, the soma showed a higher increase in potential than if the sequence moved away from the soma. Given the results by Branco et al. [173], it should be possible to reinforce synapses at any distance.

In this chapter, we are interested in reinforcing pairs of synapses that are separated by a propagation time delay corresponding to the arrival time difference of spikes from two different input neurons. We initially examined whether a detector based on dendritic propagation delays in a biophysical model of a pyramidal neuron [179] can be developed in a bottom-up, unsupervised fashion, i.e., without the soma firing a prior spike to trigger learning on pre-synaptic inputs, following a hypothesis formulated by Bugmann and Christodoulou [180]. A bottom-up approach is in the spirit of experiments conducted by Marom and Shahaf [181] showing learning without supervisory spiking by the target. The examined mechanism is based on non-linear summation of synaptic EPSPs and their effects, as described for example in Denham and Denham [182], followed by the backpropagation of the summed EPSP to the dendrites, triggering a learning mechanism at the originating synapse. Simulating this initial approach revealed that the learning mechanism appears to be insufficiently sensitive to differences in time delays. This led to the development of a second approach using a backpropagating AP.

In the second approach, a background input current is added to the neuron (at the somatic compartment), to allow the coincidence of pairs of small EPSPs to generate a spike that can then activate learning mechanisms when backpropagating. That background current can be seen as a “learning-enable” signal that is activated when the organism decides that there is a need to learn the current input situation. These processes are designed to allow learning of weights in conditions where they are initially too weak to induce output spikes.

A key element of both approaches is the assumption that inputs from each pre-synaptic neuron initially target several pre-existing synapses at various positions on the dendrite. These

synapses have a probabilistic behaviour and will activate at most one at a time, thus probing various dendritic propagation times [180]. The learning rule should then select synaptic pairs separated by the appropriate distance and reinforce them.

This approach differs from the supervised approach used by van Leeuwen [183] who assumes synaptic relocation along the dendritic tree, or the model by Hüning et al. [184] that assumes delay modification. The principle of selection of existing synapses is also used by Gerstner et al. [185], where time-differences between pre- and post-synaptic spikes determine weight changes, or the work by Eurich et al. [186] who use a Hebbian learning rule depending on correlations between pre- and post-synaptic activity within a certain time window. Senn et al. [187] also proposed the use of stochastic synapses, for adapting synaptic delays. Note that the problem treated here is different from that of detecting temporal patterns in a single input spike train, like in Hunzinger et al. [188], or global oscillations in multiple spike trains like in Kerr et al. [189]. In the context of dendritic delays selection, in this chapter we examine the capability of a pyramidal neuron to provide a learning signal selective enough to certain input time differences. In this chapter, we use the term *learning* to refer to synaptic plasticity and in particular, synaptic plasticity that reinforces a specific neural behaviour. Here, the neural behaviour we aim to reinforce is the identification of precise input delays between synaptic pairs. Moreover, the use of the term *learning* should not be confused with learning as a higher cognitive function, but only as a mechanism for making a single cell sensitive to particular input signals.

6.1 Methods

6.1.1 Overview

Figure 6.1 shows a simplified sketch of our model's architecture. Four synapses attach to a neuron's dendrite at increasing distances from the soma. The synapse that is closest to the soma, synapse B, originates from pre-synaptic neuron B. The rest of the synapses, A1, A2 and A3, originate from pre-synaptic neuron A.

Pre-synaptic neurons A and B fire the same spike trains with a fixed time delay, Δt . In other words, whenever neuron A fires at a time t , B fires at $t + \Delta t$. In this chapter, we aim at reinforcing synapse A2. We consider two scenarios (Figure 6.1):




FIGURE 6.1: Schematic of our model's architecture consisting of a simple neuron with 4 synapses. B is a proximal synapse, while synapses A_1 – 3 are at increasing distances from the soma. All A -synapses originate from the same pre-synaptic neuron (neuron A) and B originates from a different one (neuron B). See the text for an explanation of the two back-propagation diagrams.

- (i) An EPSP from A_2 arrives at B after Δt time, thus coinciding with the time the EPSP at B is created. The coinciding EPSPs are amplified, creating an increase in post-synaptic potential at B , which travels back to the A -synapses.
- (ii) The EPSPs from A_2 and B coincide at their arrival at the soma and trigger a somatic spike, creating a back-propagating AP which travels back to the synapses.

In both cases, the back-propagating potential is expected to cause weight changes in the active synapses (i.e., the synapses that have recently been active).

The purpose of both scenarios is to make the post-synaptic neuron sensitive to the firing delay between pre-synaptic neurons A and B , by reinforcing only synapse B and the corresponding A -synapse whose distance from B is such that the EPSP from A coincides with the EPSP from B , at location B . In other words, if c is the propagation speed and Δt is the firing delay between pre-synaptic neurons A and B , the learning mechanism should reinforce an A -synapse that is

at a distance $c\Delta t$ from synapse B. In all our scenarios, the A-synapse that is located at the ideal distance from B will be labelled A2. Our methods require that synapses are stochastic with a low probability of release [187, 190, 191], since synapses A1, A2 and A3 all originate from the same pre-synaptic neuron, but should receive individual reinforcement. By setting the release probability sufficiently low, we can consider that the probability of having two or more A-synapses active at the same time is negligible.

The main difference between the two approaches is the lack of somatic spiking in the first approach. The first scenario relies on the amplification and backpropagation of a potential, caused by the coinciding EPSPs at the dendritic location of synapse B. Plasticity, in this scenario, would occur as a result of the changes caused along the dendrite by the backpropagating amplified EPSP, in the absence of somatic spiking. The second scenario follows a more traditional approach to learning, where the coinciding EPSPs trigger a somatic AP that is able to cause synaptic changes based on a STDP-type learning rule.

6.1.2 Model

For our simulated experiments, we used the NEURON simulation environment [192] using a reconstructed layer 5 pyramidal neuron model, originally built by Stuart and Spruston [193]. This model was modified by Letzkus et al. [179] to account for active properties, by including voltage-gated ion channels at the following densities (in $\text{pS } \mu\text{m}^{-2}$):

- Soma: $g_{Na} = 3000$, $g_{Kv} = 30$, $g_{Ka} = 0.06$, $g_{Kca} = 2.5$, $g_{Km} = 2.2$, $g_{CaT} = 0.0003$.
- Dendrites: $g_{Na} = 40$, $g_{Kv} = 30$, $g_{Ka} = 0.03$, $g_{Kca} = 2.5$, $g_{Km} = 0.05$, $g_{CaT} = 0.0003$.
- Distal dendrites ($> 600 \mu\text{m}$ from the soma): $g_{Na} = 40$, $g_{Kv} = 30$, $g_{Ka} = 0.03$, $g_{Kca} = 2.5$, $g_{Km} = 0.05$, $g_{CaT} = 0.001$, $g_{Ca} = 1.25$ (slow high-voltage activated calcium).
- Axon: $g_{Na} = 30\,000$, $g_{Kv} = 400$.

This model, in principle, is able to generate dendritic calcium spikes, however these are not produced in the scenarios simulated in this chapter. The specific membrane resistance was $15 \text{ k}\Omega \text{ cm}^2$, the membrane capacitance was $1 \mu\text{F cm}^{-2}$ (giving a 15 ms time constant), the axial resistance was $125 \Omega \text{ cm}$. Somatic resting potential, in the absence of any background current, was -79 mV and the AP threshold was -63 mV . We further modified the resulting model by Letzkus with the addition of synaptic locations and their weights.




FIGURE 6.2: Synaptic locations on a morphological reconstructed neuron model. The grey location, indicated by S is the location of the soma.

Synapses were placed along the main branch of the dendrite at increasing distances, as shown in Figure 6.2. Excitatory post-synaptic currents (EPSCs), produced due to AMPA (α -amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid) receptors, were modelled as double exponentials with a rise time $\tau_{rise} = 0.2$ ms, a decay time $\tau_{decay} = 2$ ms and a maximum amplitude (the *weight* of the synapse). Each synapse also contained NMDA (N-Methyl-D-Aspartate) receptors, with a NMDA-AMPA ratio of 0.2, as in Letzkus et al. [179]. The NMDA component was simulated with a kinetic model, developed and described in Kampa et al. [194], and was part of the model taken from Letzkus et al. [179]. Potentials were recorded at the dendritic compartment of each synapse as well as at the soma.

6.1.3 Synaptic weights and scaling

As mentioned in the beginning of this chapter, our second approach relies on a constant background current that raises the resting potential of the neuron, such that two coinciding EPSPs would trigger an AP at the soma, but a single EPSP should not. The current is applied directly to the somatic compartment of the model. These requirements define a range of background current amplitudes (0.22 to 0.24 nA) and we use this range to explore the behaviour of the neuron. For each value of the current amplitude, a different set of synaptic weights is required.

The effect of the current amplitude on the results is discussed in Section 6.2.2.4. The figures shown in this section illustrate the behaviour for an example case of 0.222 nA constant background current, which raises the effective resting potential to -72.5 mV (6.5 mV higher than the default rest).

Synaptic weights were scaled such that an EPSP from any of the synapses had the same peak depolarisation level at the soma. The reason for the uniform scaling of synaptic weights was to make the effective difference at the soma between specific pairs of EPSPs depend only on the propagation delay of the potential and the input spike time difference. This reflects the results of Häusser [195] showing that the amplitude of the EPSP arriving at the soma is independent of the distance of the originating synapse.

The value range for the weights was assumed to be continuous (double precision floating point values, in nA) and no bounds were set on the value of the weight for each synapse. However, this does not mean that the weights were practically unbounded. If a synapse was strong enough to trigger a somatic spike on its own, i.e., without the need for a second, coincident EPSP, the weight would be considered too high. The values of the weights were therefore limited with respect to the peak depolarisation an EPSP from a synapse can cause at the soma. This was a limiting factor during synaptic scaling and background current amplitude calibration, but was not a relevant consideration during our investigation into the theoretical reinforcement of synapses due to synaptic plasticity. While investigating the strength of reinforcement caused by various plasticity rules, we assumed a continuous, unbounded range for the synaptic weights.

Figure 6.3 shows the weight of each synapse after the calibration was complete. The effects of the background current on the scaled weights can be seen in figure Figure 6.4. Figure 6.4a shows the potential recorded at the soma for four EPSPs, evoked separately at each of the synaptic locations with a constant background current applied (top) or not applied (bottom). Figure 6.4b shows the peak depolarisation caused at the soma by an EPSP from each synapse, with (top line) and without (bottom line) the background current applied. These values are calculated by subtracting the effective resting potential, i.e., the membrane potential immediately before the EPSP is triggered, from the peak potential reached during the EPSP. Figure 6.4c shows the time when the peak potential was reached, both with (top line) and without (bottom line) the background current applied. The results show that adding a background current has three effects on the EPSPs:







FIGURE 6.3: Example of synaptic weights after scaling. Weights (excitatory post-synaptic peak currents) were scaled to produce a depolarisation of 1 mV at the soma when a constant background current is applied with an amplitude of 0.222 nA. The symbols used to represent each point will be used in certain figures in this chapter to reference each synaptic location: B ●, A1 ◆, A2 ■, A3 ▲.



(a)



(b)



(c)

FIGURE 6.4: (a) Depolarisation at soma for EPSPs evoked at each synaptic location with a constant background current of 0.222 nA (top) and without (bottom). (b) Peak potential at the soma with background current (blue, solid line) and without (red, dashed line). The peak is measured as the difference between the maximum potential reached when the EPSP is evoked and the *effective* resting potential (i.e., the resting potential immediately before the EPSP is triggered). (c) Time to peak (peak time minus onset of EPSP) at the soma with background current (blue, solid line) and without (red, dashed line).

- It shifts the voltage upwards, as expected, but also increases the depolarisation caused by an EPSP. This effect becomes a little stronger with greater distance from the soma (Figure 6.4b).
- It increases the time required for the EPSP to reach its peak at the soma (Figure 6.4c).
- It also affects the shape of the EPSPs arriving at the soma, with the expected widening with distance largely being eliminated (in Figure 6.4a, top curve, the width at 95 % of the height is for B: 7.1 ms, A1: 7.1 ms, A2: 7.1 ms and A3: 7.3 ms). Note that, without background current, EPSP peaks are narrower and show the usual distance dependence (in Figure 6.4a, lower curve, the width at 95 % of the height is for B: 2.4 ms, A1: 2.8 ms, A2: 3.2 ms and A3: 3.6 ms).

These effects have also been observed in *in vitro* studies on cortical and hippocampal pyramidal cells [196–200]. Fricker and Miles [198] suggest that when the membrane potential is depolarised, the activation of inward currents tends to increase the amplitude of EPSPs and prolongs their decay.

6.1.4 Simulation procedure

Only two synapses were considered in each simulation. The assumption in our model is that synapses are probabilistic with low activation probability and we examine the cases where at most only one of the A-synapses is active at a time, along with B. Therefore, activation of synapses was induced in pairs of one of the A synapses and the B synapse, with a delay Δt , in order to examine the AP firing behaviour for separate instances of stimulus arrival. The simulation was run for 1 s before evoking any EPSPs at the synapses, to allow the potential to stabilise from initial conditions. This was especially required in the case where background current was applied. The background current is always applied directly to the somatic compartment of the neuron.




FIGURE 6.5: Integral of the NMDA conductance during the activation of a pair of synapses, A_i, B . The conductance was recorded at the location of the A-synapse (A_i) in each case. The symbols at the peak of each curve represent the three synaptic locations, A1, A2 and A3 respectively. No background current was added in this scenario.

6.2 Results

6.2.1 Approach 1: Backpropagating coincident EPSPs

In this section we report on a simulation of the forward propagation of an EPSP from A2, its coincidence with an EPSP generated by an input at B and then the backpropagation of the resulting coincident EPSP to A2. In order to evaluate the expected degree of reinforcement of synapse A_i , we calculate the integral of the NMDA conductance at that synaptic location, as in Letzkus et al. [179]. Figure 6.5 shows the integral of the NMDA conductance (normalised to the maximum) across a range of input delays (Δt) after a pair of input spikes activated a pair of synapses. For each simulation, one A-synapse was activated followed by an activation of synapse B after a delay of Δt . The order of synapse activation was reversed (B A) for negative Δt . The NMDA conductance was recorded at the location of the A-synapse being activated and the integral over the entire simulation time was calculated (see Figure 6.5).

As Figure 6.5 shows, the integrated NMDA conductance for each synapse peaks at different delays, corresponding to the dendritic propagation times from each A-synapse to B. The maximum conductance, used to normalise the conductance change, was taken to be the integral of

TABLE 6.1: Normalised integral of the NMDA conductance for specific input delays Δt . Each value is the delay which maximises the integral at one of the three locations, 0.8 ms for A1, 1.6 ms for A2, and 2.6 ms for A3.

	Δt		
	0.8 ms	1.6 ms	2.6 ms
A1	1.0000	0.9997	0.9997
A2	0.9997	1.0000	0.9999
A3	0.9997	0.9996	1.0000




FIGURE 6.6: NMDA conductance learning rule. The figure shows the degree of synaptic plasticity as a function of the NMDA conductance integral. Image taken from Letzkus et al. [179].

the conductance at the synapse when the two synapses are activated with a delay that matches the propagation delay of the EPSP between the synaptic locations. This is what we refer to as the ideal time delay, for this scenario. However, the difference in conductance integrals between synapses, for a given time delay, are very small. For instance, when the input delay is such that it maximises the integral at A2, i.e., $\Delta t = 1.6$ ms, spike pairs (A1 B) and (A3 B) have normalised NMDA conductance integrals of 0.9997 and 0.9996 respectively (see Table 6.1). A learning rule which is based on the integral of the NMDA conductance, as in Letzkus et al. [179], would apply reinforcements with minute differences between synapses, for delays -1 to 3 ms (Figure 6.6). This can be seen in Figure 6.5, where all curves are above 0.995 within the aforementioned range of delays. Ideally, in order for the neuron to learn

to respond to a specific spike delay which corresponds to the propagation delay between A2 and B ($\Delta t = 1.6$ ms), the reinforcement at that delay should be significantly different for each synapse, such that A2 receives a significantly stronger reinforcement when activated compared to A1 or A3. Our results show that this is not the case. Backpropagating coincident EPSPs are therefore unlikely to provide sufficiently strong differentiating signals.

Given this result, we have investigated a different approach using backpropagating APs instead of EPSPs. This is the topic of the next section.


6.2.2 Approach 2: Backpropagating action potential

Here we examine the idea that, by adding a background input current to the neuron, the two small coincident EPSPs become able to cause a spike that then triggers learning mechanisms when backpropagating to the recently active synapses. As learning by backpropagating APs is well documented, the main issue dealt with in this section is the appropriate generation of that AP. Again, we focus on learning at synapse A2, ignoring B.


Figure 6.7 shows two ways in which AP generation is affected by the application of a constant background current to the soma. Figure 6.7a shows how the width of the *firing domain* grows with increased background current. Each pair of curves show the minimum and maximum input delays (the delay between activation of an A-synapse and synapse B) that trigger a somatic spike across a range of background current amplitudes, for each synaptic pair (A_i B). Background currents below 0.22 nA do not raise the effective resting potential high enough to enable a somatic spike, even when EPSPs coincide perfectly at the soma. Background currents above 0.227 nA, which would drive the spike domain outside the range -20 to 20 ms, were not considered. Figure 6.7b shows the effect of increasing the background current on the somatic spike delay (the time of the somatic spike minus the time of the second EPSP's onset). For each synaptic pair (A_i B), the optimal input delay was used (see Section 6.2.2.3 for calculations and Table 6.2 for the values).

6.2.2.1 Relative timing conditions for action potential generation

To measure the selectivity of the coincidence detection mechanisms, we ran simulations where we varied the time delay between spike arrival time at an A-synapse and synapse B. In the example shown in Figure 6.8 (background current with amplitude 0.222 nA), the neuron fires



(a)



(b)

FIGURE 6.7: (a) Effect of the background current amplitude on the range of input delays that cause a response. Input delays outside the range -20 to 20 ms were not considered. (b) Effect of the background current amplitude on the somatic spike time for $(A_i B)$ spike pairs. For each pair, the optimal input delay was used (A1: 1.8 ms, A2: 3.8 ms, A3: 6.5 ms). The dashed vertical line in both figures indicates the background current amplitude of 0.222 nA used as an example throughout this chapter. The number of samples on both figures have been decimated for clarity.




FIGURE 6.8: Spike occurrence for the indicated combinations of input delay (Δt) and synapse location for an example case with background current amplitude of 0.222 nA. The bigger dots denote cases where a spike was fired, while the smaller ones show cases where no somatic spike occurred. Negative delays denote cases where the B synapse triggered an EPSP before the A-synapse. The leftmost column of dots represents simulations where both EPSPs were triggered at the B synapse. The dashed grey lines correspond to the maximum and minimum delays that cause a somatic AP according to the model described in Section 6.2.2.3. The diamond shaped points on each column represent the optimal delay between each synapse and synapse B for maximum depolarisation, which is the difference in peak times at the soma between an A-synapse and synapse B ($d_{A_i B}$, third column of Table 6.2).

for a range of time delays, and cannot discriminate, in a binary way (which is by firing or not firing) time differences shorter than around ± 7 ms. The domain of firing rises towards more positive delays for more distant synapses, showing the effect of an increasing dendritic propagation time. The rise and width of the firing domain are explained in Section 6.2.2.3.

6.2.2.2 Action potential timing

Figure 6.9a shows the time of AP production relative to the input time. The somatic spike delay (d_S) was calculated as the somatic spike time (t_S) minus the initiation time of the second input spike, $d_S = t_S - \max(t_A, t_B)$, where t_A and t_B are the times of EPSP initiation from

pre-synaptic neurons A and B respectively. The amplitude of the background current in all cases was 0.222 nA.

On the right-hand side of the horizontal axis (which shows the input delay), positive delays indicate that synapse B receives the last input. On the left-hand side, synapse A receives the last input. The four curves show a lateral shift by a time slightly shorter than the difference between the time-to-peak (T_p) delays of each individual EPSP (Figure 6.4). Table 6.2 in the next section shows the time-to-peak (T_p) values. The lateral shift shows clearly that the dendritic propagation time affects the response to inputs with different time difference. However, it also shows a very flat minimum that provides little differentiation between time differences close to the difference in T_p between synapses. All AP production times are at least 22 ms after the arrival of the last of the two inputs, due to the integration of long lasting NMDA currents needed to cause firing. They take around 2 to 3 ms to propagate back to an A-synapse, giving a round-trip time of 24 ms or more.

6.2.2.3 Modelling the selectivity

Figure 6.8 shows that an AP is generated for a range of time differences of input spikes. To explain the upper and lower boundaries of the domains of firing, we have formulated a simple model. The EPSPs arriving at the soma from B and from A_i ($i = 1, 2, 3$) cause a maximal potential increase when their peaks are coincident. This occurs when the delay ($d_{A_i B}$) between input spikes matches the difference in arrival times (differences in time to peak) at the soma:

$$d_{A_i B} = T_{p_{A_i}} - T_{p_B} \quad (6.1)$$

where $T_{p_{A_i}}$ and T_{p_B} are the times to peak for respectively inputs from A_i and B. When both inputs arrive at B, the delay is zero ($d_{BB} = 0$).

Now, let us assume that an exact coincidence is not needed and that the peaks are efficient over a given time window. A simple way to estimate this is to measure the duration T_w of the width of the EPSP peak at, say, 95 % of its maximum height. The results are shown in Table 6.2. All the peaks have a width of about 7 ms. So, the inputs can be desynchronized by ± 7 ms and still produce APs. More precisely, the maximum delay Δt_{\max} and minimum delays Δt_{\min} are given by:




FIGURE 6.9: (a) AP production time after the last input spike for the indicated synapses and for different delays between spike input at B and A. Positive delays correspond to the B input occurring after the A input. The range of delays correspond to the ones in Figure 6.8. (b) Example voltage traces for a single simulation where an AP was generated. The case shown corresponds to the point in subfigure (a) marked by the black asterisk (synapse A2, $\Delta t = 5$ ms). The top panel shows the voltage at all three locations, A2, B and soma. The bottom panel shows the depolarisation at dendritic locations B and A2 in an expanded ordinate, for clarity. In subfigure (b), the time of the input at synapse B corresponds to time 0. Both subfigures are sample cases with background current set at 0.222 nA.

TABLE 6.2: Model Data and Results. The results Δt_{\max} and Δt_{\min} are shown on Figure 6.8. T_p is the time-to-peak, i.e., the time when the EPSP reaches maximum depolarisation at the soma. d_{A_iB} is the difference between the two peak times from A_i and B respectively. This value is 0 when both EPSPs are triggered at B . T_w is the *peak duration*, which is measured as the duration where an EPSP is above 95 % of its maximum height. Δt_{\max} and Δt_{\min} are the calculated maximum and minimum delays that cause a somatic AP for an EPSP coming from a given synapse, while coupled with an EPSP from B . All values on the table are in milliseconds.

Synapse	T_p	d_{A_iB}	Delay from Figure 6.9a	T_w	Δt_{\max}	Δt_{\min}
B	6.9	0	0	7.1	7.3	-7.3
A1	8.7	1.8	1.12	7.1	8.48	-5.92
A2	10.8	3.9	2.8	7.1	10	-4.4
A3	13.4	6.5	4.8	7.3	12.1	-2.5

$$\Delta t_{\max} = d_{A_iB} + 0.5T_{w_A} + 0.5T_{w_B} \quad (6.2)$$

$$\Delta t_{\min} = d_{A_iB} - 0.5T_{w_A} - 0.5T_{w_B} \quad (6.3)$$

It turns out that this simple model explains the behaviour of the domain of firing times in Figure 6.8 (dashed lines). To obtain the correct slope of the curves, we had to replace the delay calculated from eqn. 6.1 with the actual best delay given by data in Figure 6.9a (see Table 6.2).

6.2.2.4 Effect of background current on selectivity

In order to reinforce only one of the A-synapses (A2), the input pair (A2 B) should cause a somatic spike, while the other pairs (A1 B) and (A3 B) should not. This requires the maximum input delay Δt_{\max} for A2 to be smaller than the Δt_{\min} for A3, and for the minimum delay Δt_{\min} for A2 to be larger than the Δt_{\max} for A1. This defines a very narrow domain of operation due to the shallow slopes of the max and min model curves in Figure 6.8. Given the difference between d_{A3B} and d_{A1B} (taken from Figure 6.9a and Table 6.2), the firing domain at A2 cannot be wider than ± 0.92 ms.

We can show that such a selectivity is difficult to achieve when membrane noise is taken into account. Experimental data in low noise conditions show fluctuations of the membrane potential reflecting excitatory inputs by fast (15 to 30 ms) potential rise, followed by a slower

(200 to 300 ms) decay to the average potential (estimated from Figure 2 in 201). We can therefore assume, as a first approximation, that the potential stays constant long enough for the duration of one ($A_i B$) input sequence and the summation of the two EPSCs at the soma. According to Destexhe and Paré [201], the smallest observed voltage fluctuations have a standard deviation of around 0.4 mV. Their data suggest a Beta distribution of potentials (Figure 6.10a).

$$B(V) = C \left(\frac{V}{2.5} \right)^{\alpha-1} \left(1 - \frac{V}{2.5} \right)^{\beta-1} \quad (6.4)$$

where C is a normalisation constant ensuring that the integral of $B(V)$ is equal to 1 between $V = 0$ and $V = 2.5$ mV.

We conducted a coarse fitting of the data and found a satisfactory match with values of $\alpha = 2$ and $\beta = 7$. With a maximum fluctuation of 2.5 mV, as seen in Figure 2b in Destexhe and Paré [201], this distribution yields a standard deviation of 0.4 mV and a peak at 0.5 mV, both matching experimental values. In order to relate data on potential fluctuations to results of our simulations, we calculated the equivalent background current I for each potential using the locally (around 0.22 nA) linear relation found in our simulations:

$$V(I) = -72.66 \text{ [mV]} + 55 \text{ [mV/nA]} I \quad (6.5)$$

For each current we then calculate the time window for firing using eqn. 6.6 below, based on Figure 6.7a for inputs at the A2 synapse. Figure 6.7a shows that the half range Δt of delays between an input at B and an input at A2 for which the neuron fires increases approximately as:

$$\Delta t(I) = \frac{15.5 \text{ [ms]}}{0.006 \text{ [nA]}} (I - 0.22 \text{ [nA]}) \quad (6.6)$$

This is valid for currents above 0.22 nA and no firing takes place below that current. The value of the average half time window for currents between a base current I_0 and the max current $I_0 + \Delta I_{2.5}$ corresponding to a potential increase of $\Delta I = 2.5$ mV is calculated through:

$$\overline{\Delta t} = \frac{\int_{V_0}^{V_0+2.5 \text{ mV}} \Delta t(I(V)) B(V) S(I(V)) dV}{\int_{V_0}^{V_0+2.5 \text{ mV}} B(V) S(I(V)) dV} \quad (6.7)$$

where $S(I)$ is either 1 or 0, reflecting whether the neuron can fire a somatic spike at the voltage V or not. The integral of the distribution alone in the denominator of eqn. 6.7 is normally equal

to 1, but if parts of the distribution fall below the voltage corresponding to 0.22 nA, there cannot be firing in that range and the integral in the denominator is needed for normalisation. The denominator represents the fraction of cases where the voltage is above the minimum needed for firing. From eqn. 6.5 we find that a background current increase of 0.045 nA is needed to increase the somatic membrane voltage by 2.5 mV. Therefore, we considered the following two cases:

- (i) The background current fluctuates between 0.22 nA and 0.265 nA.
- (ii) The background current fluctuates between 0.20 nA and 0.245 nA.

The results are as follows: For case (i), on average, inputs to A and B need to be no more than ± 30 ms away from the ideal time difference (the propagation time difference between the A_i and B synapse, d_{A_iB}). For case (ii), on average, inputs need to be only within ± 10 ms, but there is only a 7% probability that the background voltage is high enough for inputs to generate a spike. In the first case the potential (input current) is always sufficient to cause a spike if the time interval between input spikes is within the firing domain (input current > 0.22 nA). The potential is often larger and allows for a large variation in timing. In the second case, the base membrane potential is often (93% of the time) below the minimum potential value (produced by a current of 0.22 nA), and even inputs with the perfect time interval will not always be able to cause firing. In the few cases where the current is above 0.22 nA, the average time window is still quite high, with ± 10 ms. The more accurate the tuning of the somatic spike, the less frequent that spike is. When we reduce I_0 until the ideal window of ± 0.92 ms is achieved ($I_0 = 0.177253$ nA), the neuron is responsive in 2×10^{-8} of cases, which is very close to never. The image that emerges is that generating somatic spikes as a learning signal is almost impossible.

6.2.2.5 Effect of background current on frequency of learning feedback

We have also explored whether there is enough difference in somatic firing probabilities to provide a learning signal that differentiates A2 from A1 and A3. Ideally, the firing probability of the neuron when (A2 B) are activated should be significantly higher than when (A1 B) and (A3 B) are activated with the same input delay Δt , for a given level of baseline background current. Figure 6.10b shows the probability of firing of all three pairs for input delay $\Delta t = 2$ ms across the full range of baseline input currents I_0 that do not cause firing with probability 1.




FIGURE 6.10: (a) Distribution of membrane potential amplitudes above *effective rest*, i.e., the minimum (baseline) resting potential of the membrane when the neuron is injected with a fluctuating current. (b) Firing probability of the neuron as a function of the baseline injected background current, I_0 . The probabilities are calculated using eqn. 6.7. The smaller figures on the right hand side show the lower and upper parts of the probability curves at different scale to emphasize the small difference between the two curves. In order to calculate the firing probability of (A1 B) and (A3 B), we set the minimum current required for firing to 0.2205 nA, instead of 0.22 nA, which is the minimum current required for firing for those synaptic pairs when the input delay is $\Delta t = 2$ ms (seen in Figure 6.7a).

Since the noisy current causes fluctuations of the membrane voltage up to 2.5 mV (which corresponds to current fluctuations up to 0.045 nA), the lowest baseline input current I_0 that has a non-zero probability of causing a spike, when (A2 B) are activated with an input delay of $\Delta t = 2$ ms, is 0.175 nA. The maximum I_0 considered was 0.22 nA, since above this value the probability of (A2 B) firing is 1 (for $\Delta t = 2$ ms). The input delay $\Delta t = 2$ ms was chosen as the ideal delay of (A2 B) in terms of firing probability, since this is the only delay which causes firing for that synapse pair when the background current is 0.22 nA. In other words, it is the leftmost point of the A2 curve on Figure 6.7a (0.22 mA, 2 ms). The firing probabilities for (A1 B) and (A3 B) are the same, since at $\Delta t = 2$ ms, both synapse pairs require a minimum background current of 0.2205 nA to trigger a somatic spike.

The curves in Figure 6.10b show that the difference in firing probabilities is very small for the full range of background current amplitudes. Is this difference too small to cause significantly higher reinforcement of A2 than A1 or A3? To address this, we have examined two extreme cases:

- (i) The background current fluctuates between 0.195 nA and 0.24 nA.
- (ii) The background current fluctuates between 0.220 nA and 0.256 nA.

For case (i) the probability of firing for (A2 B) is 1.8 % and for (A1 B) and (A3 B) 1.5 %. This means that (A2 B) triggers 20 % more spikes than the other two pairs. However, the response frequency for all pairs is very low. In other words, although the relative difference in probabilities is not negligible, (A2 B) would only trigger 3 more spikes than (A1 B) or (A3 B) for every 1000 input pairs.

In case (ii), the inverse problem occurs. The probability of firing for (A2 B) is 100 % and for (A1 B) and (A3 B) 99.6 %. In this case, the response frequency is high for all pairs, but the difference in firing probability is too small for A2 to gain a significantly higher reinforcement frequency. Therefore, it is unlikely that the firing probability would affect the frequency of reinforcement in a way which would differentiate A2 from the other two synapses.

6.3 Discussion

The aim of the work of this chapter was to examine whether it is possible to train a neuron to detect input spikes from different sources with specific delays. All mechanisms considered attempted to exploit dendritic propagation times between pairs of synapses. This chapter has focused on the possibility to generate appropriate training signals to reinforce pairs of synapses with internal propagation time matching the external spike arrival time difference.

In the first approach, using backpropagating EPSPs, we found that even in the best case, where the peak of an EPSP from A2 coincided perfectly with the peak of an EPSP at B, the difference in integrated NMDA conductance was too small to be of practical use. Thus, the use of a non-spiking approach as suggested in Bugmann and Christodoulou [180] is not effective in a pyramidal neuron.

In the second approach, we generated a somatic AP in response to the coincidence of EPSPs from synapses A and synapse B. Unfortunately, our work has revealed two problems: (1) In the case of a noiseless constant background current at its minimum of 0.22 nA, the time window for firing appears to be smaller than the desired ± 0.92 ms, but considering minimal noise as observed by Destexhe and Paré [201] reveals that the practical time window is closer to ± 10 ms. Triggering a spike only when the input time delay matches the propagation time delay of a

specific synaptic pair is therefore impossible. (2) We then examined whether the probabilities of producing an AP were sufficiently different between synapse pairs (A1 B), (A2 B) and (A3 B). We found that either (i) there is a difference but with a very low probability of an AP being produced, or (ii) many APs are produced but their firing probability does not differentiate between synaptic pairs.

We also observed that there is a variation of the firing time of the APs, depending on the difference between ideal input time difference and the actual one. In Figure 6.9a, this difference is at most around 10% for large input time differences, but around values relevant to distinguishing between A1, A2 and A3, the curves show little sensitivity to be exploited, for instance, in a STDP-type learning rule. Furthermore, in the presence of noise, it is unlikely that these somatic AP delays remain reliable indicators of input time delays.

Thus, there appears to be no usable signal to reinforce one A-synapse against all the others. Although pyramidal neurons in the brain vary in electrical properties, e.g., with membrane time constants ranging from 9 to 60 ms (15 ms in the model used in this chapter), they share sufficient common structural properties [202] for these results to be of a general validity.

The selectivity of the neuron can be measured by the height of the row of dots in Figure 6.8. Our example for a background current of 0.222 nA shows that APs are generated for delays between the two input spikes varying typically by ± 7 ms for any pair of synapses. Our computational model proposes that the selectivity is related to the width of the EPSP arriving at the soma (where width refers to some effective top part of the EPSP profile, estimated at around 95% of the height in our example). Even synapses attached directly to the soma show a width of around ± 7 ms. One may ask if this is also the time window for coincidence detection. Without background current, EPSPs arriving at the soma have much narrower peaks, around 3 ms vs 7 ms in our example (Table 6.2). However, our simulations show that, when the potential rises towards the firing threshold, EPSPs become wider. This raises the question of whether a number of superposed narrow EPSPs actually stay narrow. This may be addressed in further studies.

The methods and results presented in this chapter have been recently submitted as a paper to *BioSystems* which has been accepted subject to minor revisions [38].

Chapter 7

General Discussion & Conclusions

The research presented in this thesis addressed multiple aspects of the problem of identifying the temporal precision of the neural code. This problem arises as a generalisation of the more traditional distinction between temporal integration and coincidence detection, which have been regarded as the two distinct neural operational modes [70, 76]. The generalisation used throughout our research regards the two modes as opposing extremes of a continuum: temporal integration is characterised by low temporal precision and wide (long window) spike rate averaging, while coincidence detection is characterised by high precision and narrow (short window) spike binning (Section 2.2) [31, 51, 74, 75, 111]. More importantly, the operational mode has direct implications for neural encoding mechanisms. Temporal integrators integrate all incoming activity, causing a neuron to fire responses based on the average rate of input spikes, a behaviour which characterises a rate code. Conversely, coincidence detectors respond to precise spike timing coincidences and synchronous events, which characterises a temporal code [75, 76, 111]. This generalisation of the two operational modes extends to the coding schemes, which yields the concept of a continuum between rate and temporal coding.

The general contribution of the thesis is the development of tools and methods for determining the operational mode of a neuron, through observations of its behaviour. While the results do not provide definitive answers to the question of which operational mode neurons employ, our methods provide general conclusions regarding the ways in which the operational mode is defined by various properties of the neuron and its behaviour.

More specifically, our contributions towards solving the problem of understanding the neural code, involved identifying where a neuron's behaviour lies on the continuum of operational

modes. We also studied the ways in which the neural and input parameters affect the relationship between input synchrony and operational mode, to further understand how a neuron's mode of operation is defined. A more general approach was developed and presented for inferring the parameters that define the entire signal driving a neuron, under the assumption that it is sinusoidal, representing periodic bursts of correlated or synchronous activity. Finally, we investigated ways in which a biophysical neuron model can learn (through synaptic plasticity) to distinguish input spike pairs of fixed delay and the temporal precision of the distinguishing mechanisms. The following list summarises the specific contributions that result from this work:

- We developed the *Normalised Pre-Spike membrane potential Slope* (NPSS, Chapter 3), which is a method for determining the degree of input synchrony that was responsible for the firing of a neuron and is a measure of the operational mode of a neuron. It depends on knowledge of a neuron's intrinsic properties (membrane leak time constant, resting potential, firing threshold) and membrane voltage traces to determine, using the slope of depolarisation prior to each spike, the strength of the input that was driving the neuron at the time of firing [36]. Using this method we determined that:
 - A leaky integrate-and-fire neuron with partial reset firing highly irregular spikes at high rates is primarily a temporal integrator with short leak time constant and low threshold.
 - At very high firing rates, the definitions of two operational modes converge and the two modes become indistinguishable.
- We defined the distinction between *actual* input synchrony and *utilised* input synchrony and determined the way in which neural and input properties affect the operational mode of a neuron (Chapter 4). The former—actual input synchrony—can be measured using spike train distance metrics, while the latter—utilised input synchrony—is analogous to the operational mode and is therefore measured by the aforementioned method we developed, the NPSS. The study of the relationship between the two types of synchrony revealed that:
 - Input synchrony directly determines the operational mode when a synchronous input volley depolarises the membrane enough to fire a spike from rest.
 - High synchrony brings a neuron's operational mode to pure coincidence detection only when (i) the synchronous volley causes a high enough depolarisation to cause

a response ($\Delta_v \geq V_{th}$) and (ii) spike volleys are highly synchronous (i.e., no jitter, $\sigma_{in} = 0$ ms). These conditions are both necessary and sufficient for a neuron to operate in a pure coincidence detection mode, as measured by the NPSS.

- The relationship between input synchrony and operational mode is not affected significantly by the asymptotic, threshold-free membrane potential.
- A neuron can operate in a mode between temporal integration and coincidence detection when (i) it is driven by high degrees of (actual) input synchrony and (ii) a single volley does not depolarise the membrane enough to fire from rest, but multiple volleys arrive at high enough frequency to trigger an output spike. In such cases, we define the behaviour of the neuron as “integration of coincidences”, since a high degree of coincidences occur at the input of the neuron, but multiple highly synchronous volleys are integrated before the neuron fires.
- We developed a method for inferring the input parameters and by extension the degree of input synchrony of a neuron driven by a periodic, sinusoidal signal (Chapter 5). The method determines the frequency of the sine wave using the variation in spike frequency of the neuron across time and by using previously developed input reconstruction methods for stationary signals, infers with high accuracy, the parameters of the sinusoidal input [37].
- We developed a method for analysing the ability of a detailed reconstructed pyramidal neuron to distinguish, or learn to distinguish, between pairs of input spikes separated by temporal delays of varying width (Chapter 6). Our investigation determined that:
 - The firing domain (the range of delays between a pair of spikes that triggers a somatic response) is defined in part by the amplitude of a background current which affects the resting potential of the neuron. More specifically, the higher the amplitude, the wider the firing domain of the neuron.
 - For a neuron to learn to distinguish between input spike pairings of precise delay (on the order of several milliseconds), a highly sensitive learning rule is required.
 - A pyramidal neuron with a compact dendritic tree or short propagation times (fast propagation speed), cannot learn to distinguish input time delays that vary by a few milliseconds, as long as it fires at a reasonable rate (i.e., the firing probability is not very low).

7.1 NPSS: A measure of the operational mode

The NPSS, measures the degree of input synchrony that was responsible for firing response spikes in a LIF model neuron. It estimates this quantity by normalising the slope of the membrane potential prior to firing between two theoretical bounds, each of which represents either the maximum or minimum obtainable voltage change in a small period. The maximum (upper) and minimum (lower) slope bounds represent the voltage change that occurs from completely synchronous and completely random (dispersed) input spikes causing the neuron to fire, respectively.

Our method is based on the assumption that a neuron's operational mode is solely determined by the degree of synchrony responsible for triggering spikes. This is a safe assumption, since in the presence of intrinsic neural properties which may determine the neuron's operational mode, the degree of input synchrony which causes a spike would reflect the neuron's *preference* for input spikes of that form. For example, a neuron with very short time constant ($\tau_m \ll 10$ ms), or low threshold ($N_{th} \ll 5$), would operate primarily (or even exclusively) as a coincidence detector and the neuron's sensitivity to coincidences (or insensitivity to uncorrelated input spikes) would be reflected in the dominance of coincident input events in the short time before each output spike.

We explicitly defined the bounds of the slope of the membrane potential for models with arbitrary reset potentials (i.e., models where V_{reset} is not necessarily equal to V_{rest}). Although the bounds are general, we were interested in the case of partial reset in particular ($V_{reset} > V_{rest}$), in order to study the case of high firing variability at high rates, which a partial reset model is able to accurately reproduce. Bugmann et al. [35] attempted to determine the causes of firing for this case using reverse correlations of the firing and the input spike trains, but concluded that the method was not a reliable indicator of the relative contribution of each mechanism (temporal integration and coincidence detection) to the firing of spikes. Using our method, we were able to determine that the partial reset model, when driven by random, high-rate inputs which cause it to fire highly irregular spike trains at high rates, is predominantly a temporal integrator of incoming EPSPs. In accordance with the analysis by Softky and Koch [73], the partial reset model operates as an equivalent temporal integrator with very low threshold and very short membrane leak time constant.

Our slope bound definitions also had an interesting implication. As the firing rate of a neuron increases, which causes firing ISIs to become shorter, the two bounds converge to the same value. Both bounds depend on a small time window preceding each output spike—called the coincidence window—which defines a period within which all input spikes are considered synchronised. When the firing ISI is equal to the length of the coincidence window, the two bounds become equal. The implication of this behaviour is that if the firing rate is high enough, then the two operational modes become indistinguishable. This can also be expressed intuitively. If a firing ISI is equal to or less than the defined coincidence window w , then all input spikes that triggered the output spike arrived within w . Therefore, there is no distinction between describing the neuron's behaviour as *integrating* EPSPs for a period w or responding to a *coincidence* of EPSPs in a period w . Conversely, when an ISI is much longer than w , the difference between temporal integration and coincidence detection is defined by the degree of depolarisation that occurred prior to the start of the coincidence window, with respect to the length of the ISI. For small enough coincidence window, such as in our case where $w = 2$ ms, this issue only appears at unrealistically high firing rates (above 400 Hz).

Given the existence of neurons with a very high sensitivity to highly precise coincidences, as mentioned in Section 3.7, it would be beneficial to investigate the effects of such small values of w on the results and consider the implications of such high precision. Based on the way in which the NPSS measures the operational mode, a very small value of w would result in small values for \bar{M} , meaning that the NPSS would measure a dominance of temporal integration for even very low jitter or very high synchrony. This is by design, since as mentioned already, the value of w should reflect the maximum time between two spikes that may still be regarded as coincident. If spikes are only regarded as coincident when they are within 0.1 ms of each other, then the definition of *coincidence detection* becomes much stricter and it is less likely to occur in the presence of noise.

7.2 Input synchrony and the operational mode

As mentioned in the previous section, our measure assumes that the operational mode is determined by input synchrony. However, there is an important distinction between the degree of synchrony between the neurons that are driving a single, post-synaptic cell and the degree of synchrony between the spikes that are causing the post-synaptic neuron to fire spikes. We label these two features *actual* and *utilised* input synchrony respectively. The latter of the two,

utilised input synchrony, is what determines the operational mode of the neuron: high utilised synchrony defines coincidence detection while low utilised synchrony defines temporal integration. We used a spike time distance metric to measure the actual input synchrony and studied the relationship between the two features. Our goal was to determine how the neuron's intrinsic parameters, as well as the features of the input spike trains that are not relevant to synchrony, shape the relationship between input synchrony and operational mode.

Our analysis revealed that input synchrony directly defines the operational mode when a single synchronous volley of spikes from a pre-synaptic population is able to trigger a spike directly from rest. This is caused by a large number of synchronised input spike trains or large synaptic weights at the inputs (or both), such that when the pre-synaptic population fires in synchrony, a response spike is almost surely triggered. This case is characterised by the NPSS measuring 1 when the spike time distance measures 0 and conversely, the NPSS measures 0 when the spike time distance is near the maximum value (~ 0.5). The relationship of the two measures between these two extremes varies based on the method used to generate the synchronous events. Our results and analysis show that high synchrony brings a neuron's operational mode to pure coincidence detection only when (i) the synchronous volley causes a high enough depolarisation to cause a response ($\Delta_v \geq V_{th}$) and (ii) spike volleys are highly synchronous (i.e., there is no jitter, $\sigma_{in} = 0$ ms). These conditions are both necessary and sufficient for a neuron to operate in a pure coincidence detection mode, as measured by the NPSS.

When the peak depolarisation of synchronous spike volleys is not high enough to trigger a spike, the relationship between the two measures follows the same trend as in the case described above, but the degree of utilised synchrony is generally lower. In the extreme case of absolutely synchronised inputs, the operational mode lies near the midpoint between temporal integration and coincidence detection. This value reflects the integration of multiple synchronous volleys, which are required to trigger a spike.

We also examined whether the asymptotic threshold-free potential affects the relationship between the two measures, but found no apparent effects. This indicates that the frequency of the input does not define in any way the operational mode, as long as the input frequency is high enough to cause spiking.

7.3 Input synchrony of the Ornstein-Uhlenbeck LIF

In order to extend the NPSS and address some of its shortcomings, we studied the estimation of input synchrony for a stochastic, diffusion LIF, which is based on the Ornstein-Uhlenbeck process. This allowed us to develop methods for estimating the degree of input synchrony and by extension the operational mode for a similar neuron under similar conditions to the LIF, but with different parameters and processes describing the input. We first derived the correspondence between the parameters of the two models (LIF and OU), in order to show that our choice of inputs is compatible with the pulse-packet-like input description we used in Chapters 3 and 4. Our methods focused on estimating input parameters that allow us to reconstruct the input signal driving the neuron. From the input signal or its parameters, we can determine the actual input synchrony and subsequently infer the degree of utilised input synchrony by observing the relationship between the neuron's response and input.

The input signal was described by a sine wave with added noise. This caused behaviour which is qualitatively similar to a neuron driven by periodic volleys of synchronous spikes. Our methods adapted existing estimators which were developed to estimate input parameters for stationary inputs signals. The main contribution of our adaptation was the estimation of the frequency of the underlying sine wave that describes the input. Using the frequency, we were able to obtain accurate estimates of the input parameters by aligning initial estimates of the input signal and averaging points which fall within the same segment of the sinusoidal period.

The sinusoidal input signal was chosen as a way to emulate the pulse-packet formalisation used in the previous chapters, Chapters 3 and 4, using the parameters of the OU model. Section 5.1.1 describes how each input parameter of the OU model with sinusoidal inputs translates to the LIF with periodic pulse packets and random background inputs. However, our choice of input signal also has implications relating to synchronous oscillations. In particular, oscillations are considered to be instrumental for the synchronisation of neural firing [203]. Additionally, oscillations have been found to promote reliable propagation of synchronous signals along neural pathways [204]. Our study of synchrony as a determinant of the operational mode is therefore reinforced by the addition of oscillatory signal analysis, since the two phenomena are closely related. While the sine wave frequencies used in our work (5 to 20 Hz) fall within the range of physiological oscillations (theta waves: 4 to 8 Hz, alpha waves: 8 to 13 Hz, and beta waves: 13 to 30 Hz), we did not aim to simulate a particular frequency band. Our results

and methods are general enough that they can be used to analyse neural oscillations of any frequency.

Since we used estimators for stationary signals, our initial estimates assume that the input signal during an ISI is stationary, which does not hold. However, this inaccuracy is averaged out by our binning and averaging procedure.

Although our methods are tailored to estimate sinusoidal input signals with fixed frequency, a similar technique can be used for varying frequencies or even random arrivals of synchronous input volleys. In either case, a peak detector can be used to determine the times of the peaks of the input signal and subsequently align estimates based on their relation to the peak.

Compared to the NPSS, the methods described in Chapter 5 do not require defining a coincidence window, which affects both the results and the interpretation of the temporal precision of the neuron. Additionally, the input parameter estimator is not affected by the existence of inhibitory inputs, as opposed to the NPSS which assumes purely excitatory input spikes.

7.4 Learning temporal delays between input spike pairs

Coincidence detection in neurons allows for temporally precise signal propagation. Coinciding EPSPs are more likely to trigger a response than temporally dispersed input spikes, which implies that a neuron is able to signal the existence of precisely timed firing in its pre-synaptic population. In Chapter 6 we investigated possible mechanisms for a neuron to learn to distinguish between input spike pairs with a specific fixed delay, without learning spike pairs with slightly different delays. Our results suggest that it is not possible for a pyramidal neuron to learn to distinguish time delays of the order of dendritic propagation time between synapses. Our study of a biophysical model of a pyramidal neuron considered multiple learning mechanisms and various features of the behaviour of the neuron which could, in principle, give rise to signals that can distinguish between different input time delays.

Firstly, a learning mechanism which relies on the changes in NMDA conductance at the synaptic site, in the presence of coincident inputs and in the absence of somatic spiking, was shown to be insufficient for producing a learning signal that is significantly higher for the coincident synaptic pair, when compared to other synaptic pairs.

Secondly, methods relying on a somatic spike being generated also failed to provide a mechanism for learning temporal delays. In the presence of somatic spiking, we examined:

- (i) the range of input time delays, at each synaptic pair, which produce a somatic response spike (we refer to this as the *firing domain*);
- (ii) the timing of the AP, relative to the input, for the full range of input time delays;
- (iii) the effect of background current noise on the width of the *firing domain* of each synaptic pair;
- (iv) the effect of background current noise on the firing frequency (or firing probability) of the neuron for inputs from different synaptic pairs.

None of these features provided a strong differentiating signal for different input synapses. More specifically:

- (i) The *firing domain* is not sufficiently distinct between synapses (i.e., the domains always overlap).
- (ii) APs are generated with delays that do not vary significantly when triggered by different synapses, in the range of input delays considered by our study.
- (iii) In the presence of noise, the average background current can be reduced to the point where only current peaks cause firing. In this case, the firing domain is very narrow and differentiated, but the probability of firing becomes negligible.

A central issue underlying our results is the width of the firing domain for all synapse pairs. González et al. [205] showed that supra-linear responses can be obtained for delays up to 13 ms between the activation of subsequent synapses. Similarly, we have observed that somatic responses can occur for input spikes with arrival delays beyond 20 ms. The goal of our work in Chapter 6 revolves around finding a case where spikes arriving at synapses (A2 B) within Δt of each other cause a somatic response, but spikes arriving at (A1 B) or (A3 B) within the same Δt of each other do not. Our investigation found that this was not possible, for a very wide range of values for Δt , since, in accordance with the findings in González et al. [205], delays up to 13 ms between subsequent synapses can cause a somatic response. In our case, it was found that delays of this length (~ 10 ms) either cause a somatic response or do not,

regardless of which pair of synapses was activated. A neuron with different morphology, or different biophysical properties, which has significantly slower propagation speeds across its dendritic tree, would likely be able to distinguish input time delays at much higher precision.

Several changes to the simulation setup can be made that may give rise to learning signals that differentiate between the synaptic pairs. The NMDA/AMPA ratio used in our study was 0.2, as it was in the original model [179]. This low value could be a possible explanation for the lack of sufficient difference in the NMDA activation integrals between different synaptic pairs. Higher values, such as the 1.1 ratio found in neurons in layer 5 of the pre-frontal cortex [206] would normally increase the NMDA activation integrals significantly. This would likely cause an amplification in the NMDA peaks, which in turn would affect the normalised integrals studied in the first scenario of Chapter 6. This change would most probably make the peaks in Figure 6.5 easier to differentiate, since the peaks would be higher and more pronounced. If the differentiation is high enough, i.e., the peaks are high and narrow, and the range of delays that maximise the integral for each synaptic pair is unique, then the first scenario would be successful in creating a differentiating learning signal for a specific pair of input synapses at a specific time delay and our experiments would give positive results.

The change in NMDA/AMPA ratio would also affect the values of the *normalised* NMDA conductance integrals, which we use to calculate the change in synapse weight based on a plasticity rule, described in Section 6.2.1. In our analysis, we assumed that the maximum NMDA conductance is the conductance change caused by coinciding EPSPs, which occurs when the input time delay is equal to the dendritic propagation delay. However, higher conductance changes would be possible with higher NMDA concentrations and even in our current model, the assumed maximum is not necessarily the true maximum conductance. Since our model has active conductances in the dendritic compartments, dendritic spikes are possible, which would cause a non-linear increase in NMDA conductances. The large conductance integral that would result from a dendritic spike would be a better candidate for the normalisation maximum which would affect our analysis of the results of applying the plasticity mechanism. However, in the current state of the model, this would not likely cause a differentiating signal. We expect that, even though the increase in synaptic conductance in the presence of a dendritic spike is non-linear, the relative changes in maximum conductance across all synapses would be comparable. The resulting changes in synaptic strength would be similar for all synapses at similar delays, which would result in the same lack of differentiability between synaptic pairs. The major difference would be the synaptic weight change, as a higher maximum may cause

the normalised NMDA integral to always fall within the negative part (depression) of the plasticity rule. Regardless, the effect of dendritic spikes and the change in maximum conductance integral should be investigated along with the change in NMDA/AMPA ratio.

The choice of synaptic *pairs* could also be a limiting factor. Synaptic *groups* could be studied in the same way in order to enhance input cooperativity. In this case, instead of triggering EPSPs on two synapses, where the firing times between the two are determined by a fixed delay, each synapse could be replaced by a group of synapses, and the firing between groups would occur with the same fixed delays. Furthermore, this change would also remove the assumption that multiple synapses originate from the same pre-synaptic neuron, an assumption with little experimental evidence [207–209]. It should be noted, however, that this assumption can be relaxed if we simply assume that each A-synapse originates from a different neuron, but all pre-synaptic neurons are correlated. In the case of synaptic groups, each synapse would originate from a different neuron and each synaptic group would connect a presynaptic population, which fires correlated spikes, to the postsynaptic neuron. We expect that this change would cause a significant amplification in the potential at the synaptic location where multiple coincident EPSPs are triggered. This is due to the supra-linear depolarisation, which is observed in cases where multiple synapses are active on the same dendritic compartment or spine [210] and should amplify the difference in peak depolarisation between different pairs of synaptic groups. In turn, it would cause learning signals that are significantly different for each synaptic group. This should provide a differentiating signal, regardless of the scenario examined (with or without somatic spiking). The inclusion of multiple synapses at each synaptic location would have several implications for the experiments and the model. A number of new variables would arise from this change, such as the number of synapses per location, the density of the synapses, i.e., the proximity of each synapse to its neighbour, and the distribution of initial synaptic weights. Given this change, the calibration of synaptic weights could be made only for cases where there is no background current, which would better conform with the experimental data [195], since the background current may not be required given the amplification that would be caused by the multiple synapses. These variables would have to be studied within the ranges found in the appropriate experimental literature and adjusted to accommodate the goals of our scenarios.

All synapses in Chapter 6 were placed on the main branch (the apical stem) of the dendritic tree (see Figure 6.2). Therefore, an EPSP triggered at any of the synapses caused a depolarisation at the other synaptic sites as well. Placing synapses on different branches of the dendritic

tree should limit the effects of EPSPs between synapses. This would be particularly useful when multiple synapses at each location are included, as mentioned above, since the amplified depolarisation caused by multiple coinciding synapses would need to be isolated from other synaptic locations to avoid reinforcing all synaptic groups. If multiple synapses are kept at the original locations, then the coincident firing at one synaptic location would cause a large depolarisation at all other synaptic locations as well, effectively making the production of a differentiating signal harder, or even impossible.

An interesting issue to consider is the distribution and nature of the membrane potential noise, used in Sections 6.2.2.4 and 6.2.2.5. Our calculations of the effect of noise on the behaviour of the model assumed a voltage fluctuation distribution which was obtained from experimental recordings [201]. It is unclear whether the noise recorded in these experiments represents somatic or synaptic noise. We believe it is more important to incorporate membrane potential fluctuations that fit experimentally recorded data, with no assumption as to the source of the fluctuations. Noise of a different form or source would potentially follow a different distribution and may allow for a greater difference between firing domains and firing probabilities. The effect this change would have on our methods would solely depend on the distribution of fluctuations. Different sources and intensities of noise fluctuations could be studied, alongside the changes mentioned above, to examine whether differentiation becomes feasible.

Given the biophysical neuron model used in this study, as long as it fires at a reasonable rate (i.e., the firing probability is not extremely low), the neuron shows little sensitivity to fluctuations in input time delays on the order of the effective width of the EPSP.

7.5 Simulation, theory and experiment

7.5.1 Simulations

Various choices regarding the simulation of neurons affect the accuracy of the results. Here we discuss the relevant simulation strategies of this thesis, which includes our choice of simulation environments, differential equation solvers, and simulation time steps.

The time step of a simulation or data can have a significant impact on methods and results [211]. Generally, the smaller the time step the more accurate the simulation, at the cost of longer run

times and higher memory usage. The choice of differential equation solver also provides similar trade offs between accuracy and run time [211, 212].

In our simulations of LIF-type models, we maintained a simulation time step of 0.1 ms. All such simulations were run in the Brian simulator using exact integration [213]. For the work in Chapter 3 the simulations for the publication [36] were originally generated using a custom-built simulator which used the fourth order Runge-Kutta method [36] and were reproduced in Brian for this thesis. The simulations for Chapter 6 were run in NEURON [192] using an integrator with adaptive time step called CVODE.

The exact integration method implies there is no error (i.e., deviation from an exact solution) with respect to the evolution of the membrane potential, regardless of time step. However, given that the simulator is time-driven (as opposed to event-driven), the simulation time step affects the temporal resolution of spike generation. More specifically, spikes can only occur at multiples of the time step. Our choice of 0.1 ms for the time step was chosen as a good balance between the aforementioned accuracy and simulation run time and was kept throughout the entire research project, both for simplicity and for compatibility between results. The CVODE integrator used in NEURON allows the simulator to achieve very high accuracy by shrinking the time step when necessary, while speeding up computations by growing the time step when the model's variables remain in steady states, or vary slowly.

7.5.2 Time scales: Theory and experiment

As with simulations, the temporal resolution of data from experiments is also a factor that can affect results [214]. More importantly, when developing analysis methods, it should not be assumed that the data discretization step is arbitrarily small. This is especially important when working with simple models, such as the LIF, that lend themselves to analytical treatment. Direct analysis of the behaviour of the OU, for example, is not uncommon and in fact, it is regarded as one of the model's advantages [158–162]. Similar analysis on digitized experimental recordings may be infeasible, if the developed methods rely on data (voltage traces or spike trains) of arbitrary precision.

In our work, the methods did not generally make any assumptions about the time scales of the data. However, since all our LIF-based simulations used the same time step, as mentioned above, this topic was not investigated during the research. It would be interesting to investigate

the effect of different time scales on our analysis and whether adjustments need to be made to our methods.

7.5.3 **Biological relevance of results and conclusions**

In addition to the above discussion on the relevance of time scales in both theoretical and experimental work, as well as the accuracy of simulations, we may examine the relevance of our work to biology. The degree of biological relevance of our results depends on the corresponding biological relevance of our models, the choice of parameter values and the assumptions made by our methods.

The LIF models used in Chapters 3 and 4 are relatively simple and use various artificial mechanisms (firing, reset, refractoriness) to simulate the behaviour of real neurons. The relevance of our methods is limited by the ability of such models to reproduce realistic membrane potential trajectories and firing statistics. However, they do provide a very good baseline for developing methods which can then be extended to more realistic models and data. Parameter values were chosen within the range of biologically plausible values and where appropriate, were adjusted to produce realistic firing statistics. For example, the LIFwPR model was chosen for its ability to reproduce the highly variable high rate firing observed in cortical neurons by Softky and Koch [73].

However, in more recent years, it has been observed that high firing variability at high rates may not represent the usual behaviour of cortical neurons [215–219]. More specifically, physiological firing rates are usually much slower, in the range of 1 to 5 Hz, or at most on the order of tens of Hz. Fortunately, this does not affect our methodology. The highly variable high rate firing served as our initial motivation for discovering the operational mode of neurons exhibiting this behaviour, however the methods we developed are generally applicable, irrespective of the underlying firing rate of the cell. In fact, lower firing rates make the NPSS in particular more reliable, by widening the continuum between the two extreme operational modes.

More biologically relevant results could be obtained by directly applying our methods to experimentally recorded data. Furthermore, spike trains from neurons known to employ a particular encoding scheme, such as rate coders [39, 40], could be used to adjust the firing statistics of our models, or even calibrate the models to reproduce the spike times precisely. Fitting of models to experimental data could also extend to intracellular membrane potentials, which would

increase the biological relevance of all aspects of our model (integration of potentials as well as spiking). This would provide partial experimental validation for our methods as well as our results.

7.6 General contributions

The general contributions of this thesis can be examined from different aspects which reflect the multidisciplinary nature of the work. As mentioned in Chapter 1 and Section 1.1, computational neuroscience relies on contributions from a number of fields of study. Conversely, contributions to computational neuroscience often feed back into the contributing fields. In this section, we discuss how our work contributes to computational neuroscience as a general field, the closely related field of neuroinformatics, and also specific sub-fields of computer science.

Neuroinformatics involves research which intersects the fields of neuroscience and computer science. They generally involve:

- the development and application of computational methods for the study of brain;
- the application of advanced IT (information technology) methods to deal with the large quantities and a highly complex neuroscientific data;
- the exploitation of our insights into the principles underlying brain function to develop new IT technologies.

Our work falls under the field of computational neuroscience, which is represented by the first category in the list above. We have developed methods for studying single neuron function in order to gain insight into the ways in which information is encoded in the brain. Our methods and results could produce future contributions towards the third category as well, since any advancements in neural coding can lead to better computational methods, within the realm of computational intelligence, machine learning or artificial intelligence.

7.6.1 Contributions computational neuroscience

The general goal of the work presented in this thesis is to gain a better understanding of the neural code. In other words, the aim of this work is to improve our insight into how neurons

encode and decode information. The specific goal of identifying the operational mode of neurons and measuring the temporal precision of their integration and encoding, was specifically aimed at addressing this general goal.

The problem of the neural code is a fundamental and integral issue of computational neuroscience and neuroscience in general. Our work contributes towards solving it through the development and analysis of methods that help us understand the temporal precision of signal encoding and propagation, by analysing neural behaviour, i.e., the relationship of neural responses to their input, and framing the discussion around the problem of identifying the operational mode. As described in Section 2.1, the identification of operational modes relates directly to the issue of rate codes vs temporal codes. Therefore, by developing methods of studying neural operational modes, we are enhancing the body of available tools for understanding neural coding and directly contribute towards understanding the function and information processing mechanisms of the brain, at the level of signal processing in single neurons.

7.6.2 Contributions to computer science

Neuroscience related contributions to computer science usually revolve around the fields of artificial intelligence, computational intelligence, and machine learning. While the terms artificial intelligence and computational intelligence are often used interchangeably, they differ slightly in meaning. Artificial intelligence is considered a more general term which implies any intelligence that is not biological. Computational intelligence, on the other hand, suggests a kind of intelligence that can be expressed through computational methods (e.g., artificial neural networks, genetic/evolutionary algorithms, fuzzy logic).

The most prolific and obvious contribution of computational neuroscience to computer science is the study of artificial neural networks. While the earliest forms of artificial neural networks were based on very simplified abstractions of how neurons were thought to operate [220], newer incarnations have been drawing more and more inspiration from our understanding of biological neural networks and learning mechanisms [221–226].

In this context, understanding how neurons encode and decode information in the brain helps researchers develop new neural network architectures and learning rules, in order to better replicate biological functions. Such advancements could lead to computational intelligence

techniques that can better emulate human behaviour and performance. Our work on operational modes and temporal precision could be used to understand the effects of different modes (on the temporal integration – coincidence detection continuum) on machine learning algorithms. A common issue in the training and usage of such algorithms is the trade-off between generalisation and accuracy in, for example, classification problems: higher accuracy makes the classifier more sensitive to differences between samples, while lower accuracy enhances generalisation at the cost of higher classification errors. In these cases, generalisation is desirable, as long as misclassification is low.

Understanding the role of operational modes with respect to the sensitivity of a neuron to distinguishing signals (stimuli) in realistic neural networks, could benefit the design of artificial neural networks. Networks could be designed to adjust their sensitivity to differences between signals, by shifting their mode of operation along the continuum. Such an endeavour would benefit directly from the work in this thesis, as it would require understanding how the operational mode is defined and measured in simple spiking neuron models (Chapters 3 and 5), in order to design neurons that can operate in a given mode as well as adjust it. It would also benefit from understanding the role that synchrony plays in a neuron's mode of operation and the ways in which this relationship can be enhanced, weakened, or generally affected by other properties of the neuron model (Chapter 4), such that they can be used as parameters for defining the mode of a node in the network. A neuron's sensitivity to coincidences and its potential ability to learn to distinguish between highly precise input signals (Chapter 6) could aid in the design of the input signal encoding. It should, for instance, take into account the smallest temporal shift between two signals that a neuron can detect to encode inputs into signals with differences large enough for the neuron to distinguish and small enough to allow for a wide input domain. In addition, hybrid networks can be designed, where neurons are separated into groups whose nodes share the same properties, but differ between subnetworks, such as in Poirazi et al. [222, 224]. In this case, the difference between subnetworks would be the use of different operational modes which may allow the network to take advantage of the strengths of both temporal integration and coincidence detection.

The study of all of the above issues would use our work as a foundation but would require further research to assess their applicability and benefit.

7.7 Dissemination of PhD work

The research carried out during this PhD has been widely disseminated, both in publications in refereed archival journals and by presentations in various conferences/workshops/meetings.

The publications in refereed archival journals are as follows:

- Achilleas Koutsou, Chris Christodoulou, Guido Bugmann, and Jacob Kanev. Distinguishing the causes of firing with the membrane potential slope. *Neural Computation*, 24:2318–2345, 2012.
- Achilleas Koutsou, Jacob Kanev, and Chris Christodoulou. Measuring input synchrony in the Ornstein-Uhlenbeck neuronal model through input parameter estimation. *Brain Research*, 1536:97–106, 2013.
- Achilleas Koutsou, Guido Bugmann, Chris Christodoulou. On learning temporal correlations between input neurons' spikes in a biophysical model of a pyramidal neuron. *BioSystems*, accepted subject to minor revisions, May 2015.
- Achilleas Koutsou, Jacob Kanev, Maria Economidou, Chris Christodoulou. Integrator or coincidence detector – What shapes the relation of stimulus synchrony and the operational mode of a neuron?. Submitted to *Mathematical Biosciences and Engineering*, under review, March 2015.

In addition to the above publications, which were a direct result of this PhD work, we have been working in parallel on other very relevant topics related to neural coding, which resulted in the following submission:

- Jacob Kanev, Achilleas Koutsou, Chris Christodoulou, and Klaus Obermayer. Integrator or Coincidence Detector – a novel Measure based on the discrete Reverse Correlation to determine a Neuron's operational Mode. Submitted to *Journal of Neurophysiology*, under review, June 2015.

Part of this work has also been orally presented by the candidate as an invited lecture entitled *Distinguishing the causes of firing with the membrane potential slope*, at the Department of Computational Neuroscience, Institute of Physiology of the Academy of Sciences, Prague, Czech Republic.

The presentations of this work by the candidate in conferences/workshops/meetings are as follows:

- Achilleas Koutsou, Chris Christodoulou, Jacob Kanev, and Guido Bugmann. Quantification of the contribution of temporal integration and coincidence detection to the irregularity of cortical neurons at high rates. Poster, *STM2010 – Spike Train Analysis Workshop*, Plymouth, UK, June 2010.
- Achilleas Koutsou, Chris Christodoulou, Jacob Kanev, and Guido Bugmann. Causes of firing in cortical neurons revisited: Temporal integration vs coincidence detection. Poster, *AREADNE 2012: Research in Encoding And Decoding of Neural Ensembles*, Santorini, Greece, June 2010.
- Achilleas Koutsou and Chris Christodoulou. Measuring single neuron operational modes using a metric based on the membrane potential slope. Oral presentation, *3rd Cyprus Workshop on Signal Processing and Informatics*, Nicosia, Cyprus, July 2010.
- Achilleas Koutsou, Chris Christodoulou, Guido Bugmann, and Jacob Kanev. Distinguishing the causes of firing with the membrane potential slope. *9th International Workshop on Neural Coding*, Limassol, Cyprus, October/November, 2010.
- Achilleas Koutsou, Chris Christodoulou, Guido Bugmann, and Jacob Kanev. Determining pre-synaptic synchrony and neural operational modes using the membrane potential slope. Poster, *FENS (Federation of European Neuroscience Societies) – IBRO (International Brain Research Organization) – Hertie Winter School*, Obergurgl, Austria, January, 2012.
- Achilleas Koutsou, Petr Lánský, Jacob Kanev, and Chris Christodoulou. Input synchrony estimation in the Ornstein-Uhlenbeck model through the slope of depolarisation at threshold crossing. Oral presentation, *10th International Workshop on Neural Coding*, Prague, Czech Republic, September, 2012.

- Achilleas Koutsou, Chris Christodoulou, Guido Bugmann, and Jacob Kanev. Understanding the Neural Code through Exploration of the Causes of Firing. *Research Work of Postgraduate Students*, Faculty of Pure and Applied Sciences, University of Cyprus, Nicosia, Cyprus, November, 2012.
- Achilleas Koutsou, Jacob Kanev, Maria Economidou, and Chris Christodoulou. Comparison of synchrony measures and implications for inter-network neural connectivity. *11th International Workshop on Neural Coding*, Versailles, France, October 2014.
- Achilleas Koutsou, Guido Bugmann, and Chris Christodoulou. Learning temporal correlations in input spike trains. *11th International Workshop on Neural Coding*, Versailles, France, October 2014.

Work on other relevant topics, which were not a direct result of this PhD work, was also presented by the candidate in the following workshop:

- Jacob Kanev, Achilleas Koutsou, and Chris Christodoulou. Can discrete Response-Stimulus Correlation distinguish Integration from Coincidence Detection? *10th International Workshop on Neural Coding*, Prague, Czech Republic, September, 2012.

Chapter 8

Future work

The primary general objective of this thesis was the study of the distinction between operational modes and ways to identify and measure where the behaviour of a neuron lies on the continuum between *temporal integration* and *coincidence detection*. This involved studying synchrony and its effect on the integration and firing precision of neurons, within the context of understanding the neural code and the time-scales of neural processing. This was achieved through the development of methods of identifying and measuring response relevant input synchrony from single cell observations. Further insight into the identification and measurement of synchrony and the effects it has on the operational mode of neurons can be achieved by extending our existing work and generalising our methods to multi-neuron behaviour, network behaviour and more realistic models.

Extending the NPSS

Alternative models

The NPSS can be adapted to study the behaviour of alternative models, such as the MAT model mentioned in Chapter 3. Although the two models are equivalent under certain conditions, the threshold decay parameter of the MAT model allows for different dynamics when set to values other than the membrane leak time constant. In their paper, Kobayashi et al. [80] showed how the MAT model's parameters can be adjusted to accurately reproduce various spiking characteristics, based on experimentally recorded data. The slope bounds of the NPSS could

be adapted to apply to the membrane potential dynamics of the MAT model, namely the behaviour of the threshold and the lack of a post-spike reset of the membrane potential, in order to determine the operational mode responsible for the firing characteristics that it can reproduce. Furthermore, we could adapt the NPSS to analyse the behaviour of biophysical models. For this, we would be unable to derive slope bounds analytically, as we did for the LIF model. Rather, we would derive the upper and lower slope bounds by observing the model's behaviour under conditions of very high or no synchrony, respectively. The method would therefore be more empirically based, but it would not rely on any single neuron model. Instead, it would allow us to analyse the behaviour of any model, given that the model can be simulated with inputs at both extremes of the continuum and with a variety of parameter configurations. This would make our results and conclusions more compatible with biologically relevant conditions and behaviour.

Inhibitory inputs

One of the drawbacks of the NPSS is its inability to analyse the behaviour of a neuron in the presence of inhibitory inputs. As a first extension of the work presented in this thesis, it would be interesting and useful to adapt the NPSS to measure the operational mode of neurons when driven by varying degrees of inhibition, as well as the excitation that causes firing. The method can be modified to account for the higher slopes which are caused by the membrane potential being driven below rest. Also, we expect that the inclusion of inhibition will change the relationship between input synchrony and operational mode, as the amplitude of synchronous spike volleys required to cause output spikes will need to be higher when the membrane is hyperpolarised. Generally, this extension will require redefining the slope bounds to account for the effects of an inhibitory drive, which may require making assumptions about the nature of inhibition, e.g., whether it is constant and how strong it is in relation to the excitation.

Experimentally recorded data

Another extension to the NPSS would be to adapt it to analyse recordings made from real neurons. In this case, the estimation of the bounds of the slope would be based on empirical observations of the range of the pre-spike slope, under different conditions. This is necessary since theoretical estimations of the maximum and minimum rate of change of the membrane potential would be inaccurate for a real neuron. A preliminary study on this extension has

already been performed, using intracellular data recorded *in vivo* while the cell was being stimulated by a fluctuating current [162, 227]. In this preliminary work, ISIs were binned based on their length and the bounds of the slopes were chosen for each bin separately, by observing the highest and lowest recorded slopes for each binned ISI range. An important difficulty in the adaptation of our method was the proper selection of the length of the coincidence window. Since spike initiation in real neurons is not as rapid as in LIF models, the 2 ms window used in our study is too short to capture a significant part of the stimulus that drives the neuron to fire. Since the length of the coincidence window affects both the results and the underlying assumption of what we consider coincident, this issue requires more in depth study.

Extending the input parameter estimator

The input parameter estimators described in Chapter 5 can be extended to situations where the frequency of the underlying sine wave of the input signal varies with time. Additionally, the methods can be tested on more realistic, biophysical models, that are driven by spike trains with oscillating firing frequencies, in order to determine the accuracy of our methods in more realistic circumstances.

Further investigation into learning temporal delays between input spike pairs

A better understanding of the operational mode in real neurons could be achieved by extending the work of Chapter 6. In particular, we could further investigate how the biophysical properties of different types of neurons affect the precision of coincidence detection and determine the necessary conditions for highly precise input delay learning and detection.

The operational mode of networks and single neurons within them

A more general continuation of the project might involve simulating neural networks, which exhibit synchronous activity by calibrating network parameters appropriately (e.g., external

inputs, amount and strength of inhibitory connections) within the ranges of biologically plausible values. This would allow us to probe individual neurons within the network as well as observe the behaviour of populations of neurons, in order to ascertain how our existing results hold when the simulated cell is embedded within a neural population. By observing single neuron activity in networks using the NPSS, we might be able to determine the predictability of synchronous activity within the network. This is closely related to the concept of synfire chains [70, 74, 123, 130], mentioned in Section 2.3, but could also be defined in more general terms. For instance, at the single neuron level, the observed synchrony is expected to be defined as a function of (i) the overall synchrony in the network and (ii) the connectivity density of the network [50]. In other words, synchronous activity in any network is expected to be sustained and propagated to further layers only if the connectivity is strong enough to allow it. The level of synchrony that can be sustained is expected to rely on the amount of connections from one layer to the next and the strength of that connectivity. Additionally, a network study might investigate the precision with which synchrony is propagated through the network, which aspects facilitate this propagation (e.g., connectivity patterns, variance between synaptic delays, excitation-inhibition distribution) and which inhibit it.

References

- [1] Patricia S. Churchland, Christof Koch, and Terrence J. Sejnowski. What is computational neuroscience? In E. L. Schwartz, editor, *Computational Neuroscience*, pages 46–55. MIT Press, Cambridge, MA, 1993.
- [2] Dale Purves, George J. Augustine, David Fitzpatrick, William C. Hall, Anthony-Samuel LaMantia, James O. McNamara, and S. Mark Williams. *Neuroscience*. Sinauer Associates, Sunderland, MA, 3rd edition, 2004.
- [3] Martin Rosvall and Carl T. Bergstrom. Mapping change in large networks. *PLoS ONE*, 5(1):e8694, 2010.
- [4] Peter Dayan and Larry F. Abbott. *Theoretical Neuroscience: Computational and Mathematical Modelling of Neural Systems*. MIT Press, Cambridge, MA, 2001.
- [5] Terrence J. Sejnowski, Christof Koch, and Patricia S. Churchland. Computational neuroscience. *Science*, 241:1299–1306, 1988.
- [6] Bruno A. Olshausen. Book Review: Theoretical Neuroscience. *Journal of Cognitive Neuroscience*, 15(1):154–155, 2003.
- [7] Alan L. Hodgkin and Andrew Huxley. A Quantitative Description of Membrane Current and its Application to Conduction and Excitation in Nerves. *Journal of Physiology*, 117(4):500–544, 1952.
- [8] Louis Lapicque. Recherches quantitatives sur l’excitation électrique des nerfs traitée comme une polarisation. *Journal de Physiologie et de Pathologie Générale*, 9:620–635, 1907.
- [9] Richard B. Stein. A theoretical analysis of neuronal variability. *Biophysical Journal*, 5(2):173–194, 1965.

- [10] Richard B. Stein. Some models of neuronal variability. *Biophysical Journal*, 7(1):37–68, 1967.
- [11] Bruce W. Knight. Dynamics of encoding in a population of neurons. *Journal of General Physiology*, 59:734–766, 1972.
- [12] Larry F. Abbott. Lapicque’s introduction of the integrate-and-fire model neuron (1907). *Brain Research Bulletin*, 50(5):303–304, 1999.
- [13] Nicolas Brunel and Mark C. W. van Rossum. Lapicque’s 1907 paper: From frogs to integrate-and-fire. *Biological Cybernetics*, 97(5):337–339, 2007.
- [14] Hugo de Garis, Chen Shuo, Ben Goertzel, and Lian Ruiting. A world survey of artificial brain projects Part I: Large-scale brain simulations. *Neurocomputing*, 74:3–29, 2010.
- [15] Fred Rieke, David K. Warland, Rob R. de Ruyter van Steveninck, and William Bialek. *Spikes: Exploring the neural code*. MIT Press, Cambridge, MA, 1997.
- [16] Wulfram Gerstner and Werner M. Kistler. *Spiking Neuron Models*. Cambridge University Press, Cambridge, 2002.
- [17] Henry C. Tuckwell. *Introduction to Theoretical Neurobiology: Volume 1, Linear Cable Theory and Dendritic Structure*. Cambridge University Press, Cambridge, 1988.
- [18] Henry C. Tuckwell. *Introduction to Theoretical Neurobiology: Volume 2, Nonlinear and Stochastic Theories*. Cambridge University Press, Cambridge, 1988.
- [19] Eric Schwartz. Introduction. In Eric L. Schwartz, editor, *Computational Neuroscience*, pages ix–xii. MIT Press, Cambridge, MA, 1993.
- [20] Michael Häusser. The Hodgkin-Huxley theory of the action potential. *Nature Neuroscience*, 3:1165, 2000.
- [21] David H. Hubel and Torsten N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *Journal of Physiology*, 160(1):106–154, 1962.
- [22] Jane Cronin. *Mathematical aspects of Hodgkin-Huxley neural theory*. Cambridge University Press, Cambridge, NY, 1987.
- [23] John Rinzel and G. Bard Ermentrout. Analysis of neural excitability and oscillations. In Christof Koch and Idan Segev, editors, *Methods in Neuronal Modeling*, pages 251–292. MIT Press, Cambridge, MA, 1998.

- [24] Eugene M. Izhikevich. *Dynamical Systems in Neuroscience: The Geometry of Excitability and Bursting*. MIT Press, Cambridge, MA, 2007.
- [25] Claude E. Shannon. A Mathematical Theory of Communication. *The Bell System Technical Journal*, 27:379–423, 1948.
- [26] William Bialek and Fred Rieke. Reliability and information transmission in spiking neurons. *Trends in Neurosciences*, 15(11):428–434, 1992.
- [27] Alexander Borst and Frédéric E. Theunissen. Information theory and neural coding. *Nature Neuroscience*, 2(11):947–957, 1999.
- [28] Donald H. Perkel, George L. Gerstein, and George P. Moore. Neuronal spike trains and stochastic point processes I. The single spike train. *Biophysical Journal*, 7(4):391–418, 1967.
- [29] Donald H. Perkel, George L. Gerstein, and George P. Moore. Neuronal spike trains and stochastic point processes II. Simultaneous spike trains. *Biophysical Journal*, 7(4):419–440, 1967.
- [30] Jean-Pierre Rospars and Petr Lánský. Stochastic model neuron without resetting of dendritic potential: application to the olfactory system. *Biological Cybernetics*, 69(4):283–294, 1993.
- [31] Michael A. Kisley and George L. Gerstein. The continuum of operating modes for a passive model neuron. *Neural Computation*, 11(5):1139–1154, 1999.
- [32] Sonja Goedeke and Markus Diesmann. The mechanism of synchronization in feed-forward neuronal networks. *New Journal of Physics*, 10(1):015007, 2008.
- [33] Michael R. DeWeese and Anthony M. Zador. Non-Gaussian membrane potential dynamics imply sparse, synchronous activity in auditory cortex. *Journal of Neuroscience*, 26(47):12206–12218, 2006.
- [34] Ryota Kobayashi, Shigeru Shinomoto, and Petr Lánský. Estimation of time-dependent input from neuronal membrane potential. *Neural Computation*, 23(12):3070–3093, 2011.
- [35] Guido Bugmann, Chris Christodoulou, and John G. Taylor. Role of temporal integration and fluctuation detection in the highly irregular firing of a leaky integrator neuron model with partial reset. *Neural Computation*, 9(5):985–1000, 1997.

- [36] Achilleas Koutsou, Chris Christodoulou, Guido Bugmann, and Jacob Kanev. Distinguishing the causes of firing with the membrane potential slope. *Neural Computation*, 24: 2318–2345, 2012.
- [37] Achilleas Koutsou, Jacob Kanev, and Chris Christodoulou. Measuring input synchrony in the Ornstein-Uhlenbeck neuronal model through input parameter estimation. *Brain Research*, 1536:97–106, 2013.
- [38] Achilleas Koutsou, Guido Bugmann, and Chris Christodoulou. On learning temporal correlations between input neurons’ spikes in a biophysical model of a pyramidal neuron. *BioSystems*, accepted, May 2015.
- [39] Edgar Douglas Adrian. The impulses produced by sensory nerve endings. *Journal of Physiology*, 61(1):4–72, 1926.
- [40] Edgar Douglas Adrian and Detlev W. Bronk. The discharge of impulses in motor nerve fibres. *Journal of Physiology*, 67(2):118–151, 1929.
- [41] Andrew B. Schwartz, Ronald E. Kettner, and Apostolos P. Georgopoulos. Primate motor cortex and free arm movements to visual targets in three-dimensional space. I. Relations between single cell discharge and direction of movement. *Journal of Neuroscience*, 8(8): 2913–2927, 1988.
- [42] William Bialek, Fred Rieke, Rob R. de Ruyter van Steveninck, and David K. Warland. Reading a neural code. *Science*, 252(5014):1854–1857, 1991.
- [43] Michael N. Shadlen and William T. Newsome. Noise, neural codes and cortical organization. *Current Opinion in Neurobiology*, 4(4):569–579, 1994.
- [44] Michael London, Arnd Roth, Lisa Beeren, Michael Häusser, and Peter E. Latham. Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature*, 466(7302):123–127, 2010.
- [45] Wyeth Bair and Christof Koch. Temporal Precision of Spike Trains in Extrastriate Cortex of the Behaving Macaque Monkey. *Neural Computation*, 8:1185–1202, 1996.
- [46] Rob R. de Ruyter van Steveninck, Geoffrey D. Lewen, Steven P. Strong, Roland Koberle, and William Bialek. Reproducibility and variability in neural spike trains. *Science*, 275 (5307):1805–1808, 1997.

- [47] Takeaki Shimokawa and Shigeru Shinomoto. Estimating Instantaneous Irregularity of Neuronal Firing. *Neural Computation*, 21(7):1931–1951, 2009.
- [48] Gary R. Holt, William R. Softky, Christof Koch, and Rodney J. Douglas. Comparison of discharge variability in vitro and in vivo in cat visual cortex neurons. *Journal of Neurophysiology*, 75(5):1806–1814, 1996.
- [49] Katrina MacLeod, Alex Bäcker, and Gilles Laurent. Who reads temporal information contained across synchronized and oscillatory spike trains? *Nature*, 395(6703):693–698, 1998.
- [50] W. Martin Usrey and R. Clay Reid. Synchronous activity in the visual system. *Annual Review of Physiology*, 61(1):435–456, 1999.
- [51] Jonathan D. Victor. Temporal aspects of neural coding in the retina and lateral geniculate. *Network: Computation in Neural Systems*, 10:R1–R66, 1999.
- [52] Adam L. Jacobs, Gene Fridman, Robert M. Douglas, Nazia M. Alam, Peter E. Latham, Glen T. Prusky, and Sheila Nirenberg. Ruling out and ruling in neural codes. *Proceedings of the National Academy of Sciences*, 106(14):5936–5941, 2009.
- [53] William R. Softky. Simple codes versus efficient codes. *Current Opinion in Neurobiology*, 5(2):239–247, 1995.
- [54] Michael W. Oram, Dengke Xiao, Barbara Drietschel, and Kevin R. Payne. The temporal resolution of neural codes: Does response latency have a unique role? *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 357:987–1001, 2002.
- [55] Wulfram Gerstner, Raphael Ritz, and J. Leo van Hemmen. Why spikes? Hebbian learning and retrieval of time-resolved excitation patterns. *Biological Cybernetics*, 69(5-6):503–515, 1993.
- [56] Apostolos P. Georgopoulos, Andrew B. Schwartz, and Ronald E. Kettner. Neuronal population coding of movement direction. *Science*, 233(4771):1416–1419, 1986.
- [57] Apostolos P. Georgopoulos, Ronald E. Kettner, and Andrew B. Schwartz. Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population. *Journal of Neuroscience*, 8(8):2928–2937, 1988.

- [58] Sheila Nirenberg and Peter E. Latham. Decoding neuronal spike trains: How important are correlations? *Proceedings of the National Academy of Sciences*, 100(12):7348–7353, 2003.
- [59] Don H. Johnson and Will Ray. Optimal stimulus coding by neural populations using rate codes. *Journal of Computational Neuroscience*, 16(2):129–138, 2004.
- [60] Wolf Singer and Charles M. Gray. Visual feature integration and the temporal correlation hypothesis. *Annual Review of Neuroscience*, 18(1):555–586, 1995.
- [61] Christoph von der Malsburg. Binding in models of perception and brain function. *Current Opinion in Neurobiology*, 5(4):520–526, 1995.
- [62] Zachary F. Mainen and Terrence J. Sejnowski. Reliability of spike timing in neocortical neurons. *Science*, 268(5216):1503–1506, 1995.
- [63] Gustavo Deco and Bernd Schürmann. The coding of information by spiking neurons: an analytical study. *Network: Computation in Neural Systems*, 9(3):303–317, 1998.
- [64] Gustavo Deco and Bernd Schürmann. Spatiotemporal coding in the cortex: information flow-based learning in spiking neural networks. *Neural Computation*, 11(4):919–934, 1999.
- [65] Sheila Nirenberg, Steve M. Carcieri, Adam L. Jacobs, and Peter E. Latham. Retinal ganglion cells act largely as independent encoders. *Nature*, 411(6838):698–701, 2001.
- [66] Naama Brenner, Steven P. Strong, Roland Koberle, William Bialek, and Rob R. de Ruyter van Steveninck. Synergy in a neural code. *Neural Computation*, 12(7):1531–1552, 2000.
- [67] Don H. Johnson. Four top reasons mutual information does not quantify neural information processing. In *Proceedings of the Annual Computational Neuroscience Meeting*, Chicago, IL, 2002.
- [68] Christopher J. Rozell and Don H. Johnson. Examining methods for estimating mutual information in spiking neural systems. *Neurocomputing*, 65-66:429–434, 2004.
- [69] Alexander G. Dimitrov, Aurel A. Lazar, and Jonathan D. Victor. Information theory in neuroscience. *Journal of Computational Neuroscience*, 30:1–5, 2011.
- [70] Moshe Abeles. Role of the cortical neuron: integrator or coincidence detector? *Israel Journal of Medical Sciences*, 18(1):83–92, 1982.

- [71] Michael N. Shadlen and William T. Newsome. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *Journal of Neuroscience*, 18(10):3870–3896, 1998.
- [72] William R. Softky and Christof Koch. Cortical cells should fire regularly, but do not. *Neural Computation*, 4(5):643–646, 1992.
- [73] William R. Softky and Christof Koch. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *Journal of Neuroscience*, 13(1):334–350, 1993.
- [74] Ad Aertsen, Markus Diesmann, and Mark-Oliver Gewaltig. Propagation of synchronous spiking activity in feedforward neural networks. *Journal of Physiology Paris*, 90(3-4):243–247, 1996.
- [75] Michelle Rudolph and Alain Destexhe. Tuning neocortical pyramidal neurons between integrators and coincidence detectors. *Journal of Computational Neuroscience*, 14(3):239–251, 2003.
- [76] Peter König, Andreas K. Engel, and Wolf Singer. Integrator or coincidence detector? The role of the cortical neuron revisited. *Trends in Neurosciences*, 19(4):130–137, 1996.
- [77] Chris Christodoulou and Guido Bugmann. Coefficient of variation (CV) vs mean interspike interval (ISI) curves: What do they tell us about the brain? *Neurocomputing*, 38-40:1141–1149, 2001.
- [78] Todd W. Troyer and Kenneth D. Miller. Physiological gain leads to high ISI variability in a simple model of a cortical regular spiking cell. *Neural Computation*, 9(5):971–983, 1997.
- [79] Shigeru Shinomoto and Shinsuke Koyama. A solution to the controversy between rate and temporal coding. *Statistics in Medicine*, 26(21):4032–4038, 2007.
- [80] Ryota Kobayashi, Yasuhiro Tsubo, and Shigeru Shinomoto. Made-to-order spiking neuron model equipped with a multi-timescale adaptive threshold. *Frontiers in Computational Neuroscience*, 3(9), 2009.
- [81] Anthony J. Bell, Zachary F. Mainen, Misha Tsodyks, and Terrence J. Sejnowski. ‘Balancing’ of conductances may explain irregular cortical spiking. Technical report, Institute of Neural Computation. UCSD, Sad Diego, CA, 1995.

- [82] Jianfeng Feng and David Brown. Impact of temporal variation and the balance between excitation and inhibition on the output of the perfect integrate-and-fire model. *Biological Cybernetics*, 78(5):369–376, 1998.
- [83] Marius Usher, Martin Stemmler, Christof Koch, and Zeev Olami. Network amplification of local fluctuations causes high spike rate variability, fractal firing patterns and oscillatory local field potentials. *Neural Computation*, 6(5):795–836, 1994.
- [84] Misha Tsodyks and Terrence J. Sejnowski. Rapid state switching in balanced cortical network models. *Network: Computation in Neural Systems*, 6(2):111–124, 1995.
- [85] Carl van Vreeswijk and Haim Sompolinsky. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274(5293):1724, 1996.
- [86] Carl van Vreeswijk and Haim Sompolinsky. Chaotic balanced state in a model of cortical circuits. *Neural Computation*, 10(6):1321–1371, 1998.
- [87] Juan K. Lin, Klaus Pawelzik, Udo Ernst, and Terrence J. Sejnowski. Irregular synchronous activity in stochastically-coupled networks of integrate-and-fire neurons. *Network: Computation in Neural Systems*, 9(3):333–344, 1998.
- [88] Charles F. Stevens and Anthony M. Zador. Input synchrony and the irregular firing of cortical neurons. *Nature Neuroscience*, 1(3):210–217, 1998.
- [89] Yutaka Sakai, Shintaro Funahashi, and Shigeru Shinomoto. Temporally correlated inputs to leaky integrate-and-fire models can reproduce spiking statistics of cortical neurons. *Neural Networks*, 12(7-8):1181–1190, 1999.
- [90] Adrián Ponce-Alvarez, Bjørg Elisabeth Kilavik, and Alexa Riehle. Comparison of local measures of spike time irregularity and relating variability to firing rate in motor cortical neurons. *Journal of Computational Neuroscience*, 29(1-2):351–365, 2010.
- [91] Martin P. Nawrot, Clemens Boucsein, Victor Rodriguez Molina, Alexa Riehle, Ad Aertsen, and Stefan Rotter. Measurement of variability dynamics in cortical spike trains. *Journal of Neuroscience Methods*, 169(2):374–390, 2008.
- [92] Shigeru Shinomoto, Keisetsu Shima, and Jun Tanji. Differences in spiking patterns among cortical neurons. *Neural Computation*, 15(12):2823–2842, 2003.

- [93] Ronnie M. Davies, George L. Gerstein, and Stuart N. Baker. Measurement of time-dependent changes in the irregularity of neural spiking. *Journal of Neurophysiology*, 96(2):906–918, 2006.
- [94] Keiji Miura, Masato Okada, and Shun-ichi Amari. Estimating spiking irregularities under changing environments. *Neural Computation*, 18(10):2359–2386, 2006.
- [95] Jianfeng Feng and David Brown. Coefficient of variation of interspike intervals greater than 0.5. How and when? *Biological Cybernetics*, 80(5):291–297, 1999.
- [96] Jianfeng Feng and David Brown. Impact of correlated inputs on the output of the integrate-and-fire model. *Neural Computation*, 12(3):671–692, 2000.
- [97] Chris Christodoulou and Guido Bugmann. Near Poisson-type firing produced by concurrent excitation and inhibition. *Biosystems*, 58(1-3):41–48, 2000.
- [98] Moshe Gur, Alexander Beylin, and D. Max Snodderly. Response variability of neurons in primary visual cortex (V1) of alert monkeys. *Journal of Neuroscience*, 17(8):2914–2920, 1997.
- [99] Ethan D. Gershon, Matthew C. Wiener, Peter E. Latham, and Barry J. Richmond. Coding strategies in monkey V1 and inferior temporal cortices. *Journal of Neurophysiology*, 79(3):1135–1144, 1998.
- [100] Prakash Kara, Pamela Reinagel, and R. Clay Reid. Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron*, 27(3):635–646, 2000.
- [101] James R. Müller, Andrew B. Metha, John Krauskopf, and Peter Lennie. Information conveyed by onset transients in responses of striate cortical neurons. *Journal of Neuroscience*, 21(17):6978–6990, 2001.
- [102] Asohan Amarasingham, Ting-Li Chen, Stuart Geman, Matthew T. Harrison, and David L. Sheinberg. Spike count reliability and the Poisson hypothesis. *Journal of Neuroscience*, 26(3):801–809, 2006.
- [103] Markus Meister and Michael J. Berry. The neural code of the retina. *Neuron*, 22:435–450, 1999.
- [104] Phillip J. Best, Aaron M. White, and Ali Minai. Spatial processing in the brain: The activity of hippocampal place cells. *Annual Review of Neuroscience*, 24(1):459–486, 2001.

- [105] Ranulfo Romo and Emilio Salinas. Flutter discrimination: Neural codes, perception, memory and decision making. *Nature Reviews Neuroscience*, 4(3):203–218, 2003.
- [106] Rachel I. Wilson and Zachary F. Mainen. Early events in olfactory processing. *Neuroscience*, 29(1):163, 2006.
- [107] Christoph E. Schreiner and Jeffrey A. Winer. Auditory cortex mapmaking: Principles, projections, and plasticity. *Neuron*, 56(2):356–365, 2007.
- [108] Greg D. Field and E. J. Chichilnisky. Information processing in the primate retina: Circuitry and coding. *Annual Review of Neuroscience*, 30(1):1–30, 2007.
- [109] Nicholas J. Priebe and David Ferster. Inhibition, spike threshold, and stimulus selectivity in primary visual cortex. *Neuron*, 57(4):482–497, 2008.
- [110] Jonathon Shlens, Fred Rieke, and E. J. Chichilnisky. Synchronized firing in the retina. *Current Opinion in Neurobiology*, 18(4):396–402, 2008.
- [111] R. Christopher DeCharms and Anthony M. Zador. Neural representation and the cortical code. *Annual Review of Neuroscience*, 23(1):613–647, 2000.
- [112] Jonathan D. Victor and Keith P. Purpura. Nature and precision of temporal coding in visual cortex: A metric-space analysis. *Journal of Neurophysiology*, 76(2):1310–1326, 1996.
- [113] Mark C. W. van Rossum. A novel spike distance. *Neural Computation*, 13(4):751–764, 2001.
- [114] Jonathan D. Victor. Spike train metrics. *Current Opinion in Neurobiology*, 15(5):585–592, 2005.
- [115] Conor Houghton and Kamal Sen. A new multineuron spike train metric. *Neural Computation*, 20(6):1495–1511, 2008.
- [116] Conor Houghton. Studying spike trains using a van Rossum metric with a synapse-like filter. *Journal of Computational Neuroscience*, 26(1):149–155, 2009.
- [117] Wulfram Gerstner, Richard Kempter, J. Leo van Hemmen, and Hermann Wagner. A neuronal learning rule for sub-millisecond temporal coding. *Nature*, 383(6595):76–78, 1996.

- [118] Donata Oertel. The role of timing in the brain stem auditory nuclei of vertebrates. *Annual review of physiology*, 61:497–519, 1999.
- [119] Donata Oertel, Ramazan Bal, Stephanie M. Gardner, Philip H. Smith, and Philip X. Joris. Detection of synchrony in the activity of auditory nerve fibers by octopus cells of the mammalian cochlear nucleus. *Proceedings of the National Academy of Sciences*, 97(22):11773–11779, 2000.
- [120] Micheal C. Reed, Jacob J. Blum, and Colleen C. Mitchell. Precision of neural timing: Effects of convergence and time-windowing. *Journal of Computational Neuroscience*, 13(1):35–47, 2002.
- [121] Petr Marsalek and Petr Lánský. Proposed mechanisms for coincidence detection in the auditory brainstem. *Biological Cybernetics*, 92(6):445–451, 2005.
- [122] Colleen C. Mitchell. Precision of neural timing: The small ϵ limit. *Journal of Mathematical Analysis and Applications*, 309(2):567–582, 2005.
- [123] Markus Diesmann, Mark-Oliver Gewaltig, and Ad Aertsen. Stable propagation of synchronous spiking in cortical neural networks. *Nature*, 402(6761):529–533, 1999.
- [124] Markus Diesmann, Mark-Oliver Gewaltig, and Ad Aertsen. Characterization of synfire activity by propagating pulse packets. In J. Bower, editor, *Computational Neuroscience: Trends in Research*, pages 59–64. Academic Press, San Diego, CA, 1996.
- [125] Mark-Oliver Gewaltig, Markus Diesmann, and Ad Aertsen. Propagation of cortical synfire activity: Survival probability in single trials and stability in the mean. *Neural Networks*, 14(6-7):657–673, 2001.
- [126] Anthony N. Burkitt and G.M. Clark. Analysis of integrate-and-fire neurons: synchronization of synaptic input and spike output. *Neural Computation*, 11(4):871–901, 1999.
- [127] Emilio Salinas and Terrence J. Sejnowski. Correlated neuronal activity and the flow of neural information. *Nature Reviews Neuroscience*, 2(8):539–550, 2001.
- [128] Yoshiyuki Asai and Alessandro E. P. Villa. Integration and transmission of distributed deterministic neural activity in feed-forward networks. *Brain Research*, 1434:17–33, 2012.
- [129] Romain Brette. Computing with neural synchrony. *PLoS Computational Biology*, 8(6):e1002561, 2012.

- [130] Sven Schrader, Sonja Grün, Markus Diesmann, and George L. Gerstein. Detecting synfire chain activity using massively parallel spike train recording. *Journal of Neurophysiology*, 100(4):2165–2176, 2008.
- [131] Michael Denker, Bernd Wiebelt, Denny Fliegner, Markus Diesmann, and Abigail Morrison. Practically trivial parallel data processing in a neuroscience laboratory. In Sonja Grün and Stefan Rotter, editors, *Analysis of Parallel Spike Trains*, pages 413–436. Springer US, Boston, MA, 2010.
- [132] Sonja Grün, Markus Diesmann, and Ad Aertsen. Unitary events in multiple single-neuron spiking activity: I. Detection and significance. *Neural Computation*, 14(1):43–80, 2002.
- [133] Sonja Grün. Data-driven significance estimation for precise spike correlation. *Journal of Neurophysiology*, 101:1126–1140, 2009.
- [134] Sonja Grün, Markus Diesmann, and Ad Aertsen. Unitary event analysis. In S Grün and S. Rotter, editors, *Analysis of Parallel Spike Trains*, volume 7 of *Springer Series in Computational Neuroscience*, pages 191–220. Springer, New York, NY, 2010.
- [135] Benjamin Staude, Stefan Rotter, and Sonja Grün. CUBIC: Cumulant based inference of higher-order correlations in massively parallel spike trains. *Journal of Computational Neuroscience*, 29(1-2):327–350, 2010.
- [136] Moshe Abeles. *Corticonics: Neural circuits of the cerebral cortex*. Cambridge University Press, Cambridge, 1991.
- [137] Moshe Abeles and I. Gat. Detecting precise firing sequences in experimental data. *Journal of Neuroscience Methods*, 107(1-2):141–154, 2001.
- [138] Jonathan D. Victor and Keith P. Purpura. Metric-space analysis of spike trains: theory, algorithms and application. *Network: Computation in Neural Systems*, 8(2):127–164, 1997.
- [139] Thomas Kreuz, Julie S. Haas, Alice Morelli, Henry D. I. Abarbanel, and Antonio Politi. Measuring spike train synchrony. *Journal of Neuroscience Methods*, 165(1):151–161, 2007.
- [140] Thomas Kreuz, Daniel Chicharro, Ralph G. Andrzejak, Julie S. Haas, and Henry D. I. Abarbanel. Measuring multiple spike train synchrony. *Journal of Neuroscience Methods*, 183(2):287–99, 2009.

- [141] Jonathan D. Victor and Keith P. Purpura. Spike Metrics. In Sonja Grün and Stefan Rotter, editors, *Analysis of Parallel Spike Trains*, pages 129–156. Springer US, Boston, MA, 2010.
- [142] Thomas Kreuz, Daniel Chicharro, Martin Greschner, and Ralph G. Andrzejak. Time-resolved and time-scale adaptive measures of spike train synchrony. *Journal of Neuroscience Methods*, 195:92–106, 2011.
- [143] Thomas Kreuz, Daniel Chicharro, Conor Houghton, Ralph G. Andrzejak, and Florian Mormann. Monitoring spike train synchrony. *Journal of Neurophysiology*, 109(5):1457–1472, 2013.
- [144] Cătălin V. Rusu and Răzvan V. Florian. A New Class of Metrics for Spike Trains. *Neural Computation*, 26:306–348, 2014.
- [145] Dimitrie Pompeiu. Sur la continuité des fonctions de variables complexes. *Annales de la Faculté des Sciences de Toulouse*, 2/7(3):265–315, 1905.
- [146] Felix Hausdorff. *Grundzüge der Mengenlehre*. Leipzig: Verlag von Veit & Comp., 1914.
- [147] P. Lansky. Inference for the diffusion models of neuronal activity. *Mathematical Biosciences*, 67(2):247–260, 1983.
- [148] Susanne Ditlevsen and Petr Lánský. Estimation of the input parameters in the Ornstein-Uhlenbeck neuronal model. *Physical Review E*, 71(1):011907(1)–011907(9), 2005.
- [149] Enrico Bibbona, Petr Lánský, Laura Sacerdote, and Roberta Sirovich. Errors in estimation of the input signal for integrate-and-fire neuronal models. *Physical Review E*, 78:011918, 2008.
- [150] George Uhlenbeck and Leonard S. Ornstein. On the theory of the Brownian motion. *Physical Review*, 36(5):823–841, 1930.
- [151] Alexandre Iolov, Susanne Ditlevsen, and André Longtin. Fokker-Planck and Fortet equation-based parameter estimation for a leaky integrate-and-fire model with sinusoidal and stochastic forcing. *Journal of Mathematical Neuroscience*, 4(4), 2014.
- [152] Chris Christodoulou, Trevor G. Clarkson, Guido Bugmann, and John G. Taylor. Analysis of fluctuation-induced firing in the presence of inhibition. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, volume III, pages 115–120, Como, Italy, 2000. IEEE Computer Society Press.

- [153] Alexandre Kuhn, Stefan Rotter, and Ad Aertsen. Correlated input spike trains and their effects on the response of the leaky integrate-and-fire neuron. *Neurocomputing*, 44:121–126, 2002.
- [154] Emilio Salinas and Terrence J. Sejnowski. Integrate-and-fire neurons driven by correlated stochastic input. *Neural Computation*, 14(9):2111–2155, 2002.
- [155] Attila Losonczy and Jeffrey C. Magee. Integrative Properties of Radial Oblique Dendrites in Hippocampal CA1 Pyramidal Neurons. *Neuron*, 50(2):291–307, 2006.
- [156] Michael Denker, Sébastien Roux, Henrik Lindén, Markus Diesmann, Alexa Riehle, and Sonja Grün. The local field potential reflects surplus spike synchrony. *Cerebral Cortex*, 21:2681–2696, 2011.
- [157] Thomas Kreuz. Measures of spike train synchrony. *Scholarpedia*, 6(10):11934, 2011.
- [158] Luigi M. Ricciardi and Laura Sacerdote. The Ornstein-Uhlenbeck process as a model for neuronal activity. *Biological Cybernetics*, 35(1):1–9, 1979.
- [159] Luigi M. Ricciardi and Shunsuke Sato. First-passage-time density and moments of the Ornstein-Uhlenbeck process. *Journal of Applied Probability*, 25:43–57, 1988.
- [160] Petr Lánský and Jean-Pierre Rospars. Ornstein-Uhlenbeck model neuron revisited. *Biological Cybernetics*, 72(5):397–406, 1995.
- [161] Laura Sacerdote and Petr Lánský. Interspike interval statistics in the Ornstein-Uhlenbeck neuronal model with signal-dependent noise. *Biosystems*, 67(1-3):213–219, 2002.
- [162] Petr Lánský, Pavel Sanda, and Jufang He. The parameters of the stochastic leaky integrate-and-fire neuronal model. *Journal of Computational Neuroscience*, 21(2):211–223, 2006.
- [163] Nicolas Fourcaud and Nicolas Brunel. Dynamics of the firing probability of noisy integrate-and-fire neurons. *Neural Computation*, 14(9):2057–110, 2002.
- [164] Benjamin Lindner, Jordi Garcia-Ojavlo, Alexander Neiman, and Lutz Schimansky-Geier. Effects of noise in excitable systems. *Physics Reports*, 392(6):321–424, 2004.
- [165] Donald E. Knuth. Two notes on notation. *The American Mathematical Monthly*, 99(5):403–422, 1992.

- [166] Petr Lánský and Laura Sacerdote. The Ornstein – Uhlenbeck neuronal model with signal-dependent noise. *Physics Letters A*, 285:132–140, 2001.
- [167] Luca Gammaitoni, Peter Hänggi, Peter Jung, and Fabio Marchesoni. Stochastic resonance. *Reviews of Modern Physics*, 70(1):223–287, 1998.
- [168] Gregor Wenning and Klaus Obermayer. Activity Driven Adaptive Stochastic Resonance. In *Physical Review Letters*, volume 90, page 120602, 2003.
- [169] Ryota Kobayashi, Tsubo Yasuhiro, Petr Lánský, and Shigeru Shinomoto. Estimating time-varying input signals and ion channel states from a single voltage trace of a neuron. In J. Shawe-Taylor, R. S. Zemel, P. Bartlett, F. C. N. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, pages 217–225. Curran Associates, Red Hook, NY, 2011.
- [170] Hideaki Kim and Shigeru Shinomoto. Estimating nonstationary input signals from a single neuronal spike train. *Physical Review E*, 86(5):051903, 2012.
- [171] G. Bard Ermentrout, Roberto F. Galán, and Nathaniel N. Urban. Reliability, synchrony and noise. *Trends in Neurosciences*, 31(8):428–434, 2008.
- [172] Alan B. Saul. Lagged cells in alert monkey lateral geniculate nucleus. *Visual Neuroscience*, 25(5-6):647–59, 2008.
- [173] Tiago Branco, Beverley A. Clark, and Michael Häusser. Dendritic discrimination of temporal input sequences in cortical neurons. *Science*, 329(5999):1671–1675, 2010.
- [174] Henry Markram, Joachim Lübke, Michael Frotscher, and Bert Sakmann. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, 275(5297):213, 1997.
- [175] Guo-Qiang Bi and Mu-Ming Poo. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of Neuroscience*, 18(24):10464–10472, 1998.
- [176] Li I. Zhang, Huizhong W. Tao, Christine E. Holt, William A. Harris, and Mu-Ming Poo. A critical window for cooperation and competition among developing retinotectal synapses. *Nature*, 395:37–44, 1998.

- [177] Robert C. Froemke and Yang Dan. Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature*, 416:433–438, 2002.
- [178] Yang Dan and Mu-Ming Poo. Spike timing-dependent plasticity of neural circuits. *Neuron*, 44(1):23–30, 2004.
- [179] Johannes J. Letzkus, Björn M. Kampa, and Greg J. Stuart. Learning rules for spike timing-dependent plasticity depend on dendritic synapse location. *Journal of Neuroscience*, 26(41):10420–10429, 2006.
- [180] Guido Bugmann and Chris Christodoulou. Learning temporal correlation between input neurons by using dendritic propagation delays and stochastic synapses. In *Proceedings of the International Workshop on Neural Coding*, pages 131–132, Plymouth, UK, 2001.
- [181] Shimon Marom and Goded Shahaf. Development, learning and memory in large random networks of cortical neurons: lessons beyond anatomy. *Quarterly Reviews of Biophysics*, 35(1):63–87, 2002.
- [182] Michael J. Denham and Susan L. Denham. A synaptic learning rule based on the temporal coincidence of pre- and postsynaptic activity. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN'01)*, volume 1 Jul 2001, pages 1–6, Washington D.C., USA, 2001.
- [183] Matthijs van Leeuwen. Spike timing dependent structural plasticity in a single model neuron. Technical Report INF/SCR-04-15, Institute for Information and Computing Sciences, Utrecht University, 2004.
- [184] H. Hünig, H. Glünder, and Günther Palm. Synaptic delay learning in pulse-coupled neurons. *Neural Computation*, 10:555–565, 1998.
- [185] Wulfram Gerstner, Richard Kempter, J. Leo Van Hemmen, and Hermann Wagner. Hebbian Learning of Pulse Timing in the Barn Owl Auditory System. In W. Maass and C. M. Bishop, editors, *Pulsed Neural Nets*, chapter 14, pages 353–377. MIT Press, Cambridge, MA, 1998.
- [186] Christian W. Eurich, Klaus Pawelzik, Udo Ernst, Andreas Thiel, Jack D. Cowan, and John G. Milton. Delay adaptation in the nervous system. *Neurocomputing*, 32-33:741–748, 2000.

- [187] Walter Senn, Martin Schneider, and Berthold Ruf. Activity-Dependent Selection of Axonal and Dendritic Delays Or, Why Synaptic Transmission Should Be Unreliable. *Neural Computation*, 14(3):583–619, 2002.
- [188] Jason F. Hunzinger, Victor H. Chan, and Robert C. Froemke. Learning complex temporal patterns with resource-dependent spike timing-dependent plasticity. *Journal of Neurophysiology*, 108(2):551–566, 2012.
- [189] Robert R. Kerr, Anthony N. Burkitt, Doreen A. Thomas, Matthieu Gilson, and David B. Grayden. Delay selection by spike-timing-dependent plasticity in recurrent networks of spiking neurons receiving oscillatory inputs. *PLoS Computational Biology*, 9(2):e1002897, 2013.
- [190] R. Y. Pun, E. A. Neale, P. B. Guthrie, and P. G. Nelson. Active and inactive central synapses in cell culture. *Journal of Neurophysiology*, 56(5):1242–1256, 1986.
- [191] S Redman. Quantal analysis of synaptic potentials in neurons of the central nervous system. *Physiological Reviews*, 70(1):165–98, 1990.
- [192] Michael Hines and Ted Carnevale. The NEURON simulation environment. *Neural Computation*, 9(6):1179–1209, 1997.
- [193] Greg J. Stuart and Nelson Spruston. Determinants of voltage attenuation in neocortical pyramidal neuron dendrites. *Journal of Neuroscience*, 18(10):3501–3510, 1998.
- [194] Björn M. Kampa, John Clements, Peter Jonas, and Greg J. Stuart. Kinetics of Mg²⁺ unblock of NMDA receptors: implications for spike-timing dependent synaptic plasticity. *Journal of Physiology*, 556(Pt 2):337–345, 2004.
- [195] Michael Häusser. Synaptic function: Dendritic democracy. *Current Biology*, 11(1):R10–R12, 2001.
- [196] Rudolph A. Deisz, Gilles Fortin, and Walter Zieglgansberger. Voltage dependence of excitatory postsynaptic potentials of rat neocortical neurons. *Journal of Neurophysiology*, 65(2):371–382, 1991.
- [197] Greg Stuart and Bert Sakmann. Amplification of EPSPs by axosomatic sodium channels in neocortical pyramidal neurons. *Neuron*, 15(5):1065–1076, 1995.

- [198] Desdemona Fricker and Richard Miles. EPSP amplification and the precision of spike timing in hippocampal neurons. *Neuron*, 28(2):559–569, 2000.
- [199] Guillermo González-Burgos and German Barrionuevo. Voltage-gated sodium channels shape subthreshold EPSPs in layer 5 pyramidal neurons from rat prefrontal cortex. *Journal of Neurophysiology*, 86(4):1671–1684, 2001.
- [200] Veronika Zsiros and Shaul Hestrin. Background synaptic conductance and precision of EPSP-spike coupling at pyramidal cells. *Journal of Neurophysiology*, 93(6):3248–3256, 2005.
- [201] Alain Destexhe and Denis Paré. Impact of network activity on the integrative properties of neocortical pyramidal neurons in vivo. *Journal of Neurophysiology*, 81:1531–1547, 1999.
- [202] Nelson Spruston. Pyramidal neurons: dendritic structure and synaptic integration. *Nature Reviews Neuroscience*, 9:206–221, 2008.
- [203] Peter J. Uhlhaas, Gordon Pipa, Bruss Lima, Lucia Melloni, Sergio Neuenschwander, Danko Nikolić, and Wolf Singer. Neural synchrony in cortical networks: history, concept and current status. *Frontiers in Integrative Neuroscience*, 3, 17, 2009.
- [204] Sven Jahnke, Raoul-Martin Memmesheimer, and Marc Timme. Oscillation-induced signal transmission and gating in neural circuits. *PLoS Computational Biology*, 10(12):e1003940, 2014.
- [205] José Francisco Gómez González, Bartlett W. Mel, and Panayiota Poirazi. Distinguishing linear vs. non-linear integration in CA1 radial oblique dendrites: It’s about time. *Frontiers in Computational Neuroscience*, 5:1–12, 2011.
- [206] Kyriaki Sidiropoulou and Panayiota Poirazi. Predictive features of persistent activity emergence in regular spiking and intrinsic bursting model neurons. *PLoS Computational Biology*, 8(4):e1002489, 2012.
- [207] Nicolas Toni, Pierre Alain Buchs, Irina Nikonenko, C. R. Bron, and Dominique Muller. LTP promotes formation of multiple spine synapses between a single axon terminal and a dendrite. *Nature*, 402(6760):421–5, 1999.

- [208] Kara D. Federmeier, Jeffrey A. Kleim, and William T. Greenough. Learning-induced multiple synapse formation in rat cerebellar cortex. *Neuroscience Letters*, 332(3):180–184, 2002.
- [209] N. I. Medvedev, Glenn Dallérac, V. I. Popov, J. J. Rodriguez Arellano, H. A. Davies, I. V. Kraev, V. Doyère, and M. G. Stewart. Multiple spine boutons are formed after long-lasting LTP in the awake rat. *Brain Structure & Function*, 219(1):407–14, 2014.
- [210] Mark T. Harnett, Judit K. Makara, Nelson Spruston, William L. Kath, and Jeffrey C. Magee. Synaptic amplification by dendritic spines enhances input cooperativity. *Nature*, 491(7425):599–602, 2012.
- [211] Stephan Henker, Johannes Partzsch, and René Schüffny. Accuracy evaluation of numerical methods used in state-of-the-art simulators for spiking neural networks. *Journal of Computational Neuroscience*, 32(2):309–326, 2011.
- [212] Stefan Rotter and Markus Diesmann. Exact digital simulation of time-invariant linear systems with applications to neuronal modeling. *Biological Cybernetics*, 81(5-6):381–402, 1999.
- [213] Dan Goodman and Romain Brette. Brian: A simulator for spiking neural networks in Python. *Frontiers in Neuroinformatics*, 2(5):1–10, 2008.
- [214] Valeria Del Prete, Laura Martignon, and Alessandro E P Villa. Detection of syntopies between multiple spike trains using a coarse-grain binarization of spike count distributions. *Network: Computation in Neural Systems*, 15(1):13–28, 2004.
- [215] Roland Baddeley, Larry F. Abbott, Michael C. A. Booth, Frank Sengpiel, Tobe Freeman, Edward A. Wakeman, and Edmund T. Rolls. Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proceedings of the Royal Society of London B*, 264(1389):1775–1783, 1997.
- [216] David Attwell and Simon B. Laughlin. An energy budget for signaling in the grey matter of the brain. *Journal of Cerebral Blood Flow and Metabolism*, 21(10):1133–1145, 2001.
- [217] Bruno A. Olshausen and David J. Field. Sparse coding of sensory inputs. *Current Opinion in Neurobiology*, 14(4):481–487, 2004.
- [218] Bruno A. Olshausen and David J. Field. How close are we to understanding V1? *Neural Computation*, 17(8):1665–1699, 2005.

- [219] Ben D B Willmore, James A Mazer, and Jack L Gallant. Sparse coding in striate and extrastriate visual cortex. *Journal of Neurophysiology*, 105(6):2907–2919, 2011.
- [220] Warren S. McCulloch and Walter H. Pitts. A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biology*, 5(4):115–133, 1943.
- [221] Wolfgang Maass. Networks of spiking neurons: The third generation of neural network models. *Neural Networks*, 10(9):1659–1671, 1997.
- [222] Panayiota Poirazi, Costas Neocleous, Costantinos S. Pattichis, and Christos N. Schizas. A biologically inspired neural network composed of dissimilar single neuron models. In *Proceedings of the 23rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society, October 2001, Istanbul Turkey*, volume 1, pages 676–679, 2001.
- [223] Wolfgang Maass. Computing with spikes. *Special Issue on Foundations of Information Processing of TELEMATIK*, 8(1):32–36, 2002.
- [224] Panayiota Poirazi, Costas Neocleous, Costantinos S. Pattichis, and Christos N. Schizas. Classification capacity of a modular neural network implementing neurally inspired architecture and training rules. *IEEE Transactions on Neural Networks*, 15(3):597–612, 2004.
- [225] Wolfgang Maass and Henry Markram. On the computational power of circuits of spiking neurons. *Journal of Computer and System Sciences*, 69(4):593–616, 2004.
- [226] Juergen Schmidhuber. Deep Learning in Neural Networks: An Overview. Technical Report Technical Report IDSIA-03-14, The Swiss AI Lab IDSIA Istituto DalleMolle di Studi sull’Intelligenza Artificiale University of Lugano & SUPSI, Manno-Lugano, Switzerland, 2014.
- [227] Jufang He. Slow Oscillation in Non-Lemniscal Auditory Thalamus. *Journal of Neuroscience*, 23(23):8281–8290, 2003.