

Redundant Input Cancellation by a Bursting Neural Network

by

Kieran Bol

Thesis submitted to the
Faculty of Graduate and Postdoctoral Studies
In partial fulfillment of the requirements
For the M.Sc. degree in
Physics

Department of Physics
Faculty of Science
University of Ottawa

© Kieran Bol, Ottawa, Canada, 2011

Abstract

Although there is a general understanding of what methods the brain uses to process sensory information, little is known about how the brain implements these techniques using a neural network. One of the most powerful and important applications is solving the sensory “cocktail party problem:” to adaptively suppress extraneous signals in an environment. Theoretical studies suggest that the solution to the problem involves an adaptive filter, which learns to remove the redundant noise. However, neural learning is also in its infancy, and there are still many questions about the stability and application of synaptic learning rules in neural computation.

In their natural surroundings, weakly electric fish solve a simplified version of the cocktail party problem: they can selectively remove redundant signals from their sensory input. Cancellation of predictable signals is accomplished in *A. Leptorhynchus* using a neural network that putatively operates as an adaptive filter. In this thesis, the implementation of an adaptive filter in the brain of *A. Leptorhynchus* was found to require a cerebellar architecture that could supply independent frequency channels of delayed feedback and multiple burst learning rules that could shape this feedback. This extends the concept of spike-timing dependent plasticity to include bursts as events for learning and unifies two ideas about the function of the cerebellum that were previously separate: the cerebellum as an adaptive filter and as a generator of precise temporal inputs. Further, the success of the neural network in this manuscript shows that the cerebellum can stably learn a set of synaptic weights even with unsupervised learning. In addition, the output parameters of the model during local and global stimulation paradigms were accurately predicted by theoretical analysis.

Acknowledgements

I would like to thank Prof. André Longtin for supervising and assisting me throughout this undertaking. He was always encouraging, flexible and there was no question he was fully committed to helping me succeed. I would also like to thank NSERC for their financial support and my colleagues in the Physics department for their tips, hints, tricks, and their company. Finally, I would like to thank Prof. Len Maler for very helpful discussions and Dr. Gary Marsat, who, among other things, performed the experimental data collection this investigation is based on.

Dedication

For my family, the best support team a person could ask for, and for my friends, who make the journey worthwhile.

Contents

1	Introduction	1
1.1	Sensory Processing and Adaptive Filtering	1
1.2	Neurons	2
1.3	Neural Encoding	4
1.4	Neural Learning and Information Storage	6
1.5	Spike-Timing Dependent Plasticity	8
1.6	Supervised Learning	11
1.7	Cerebellum	11
1.8	Weakly Electric Fish	13
1.9	Organization of Thesis	16
2	Biological Background and Experimental Data	17
2.1	Weakly Electric Fish	17
2.2	Experimental Data	22
2.2.1	Neural Spiking Statistics	22
2.2.2	Surgery and <i>In Vivo</i> Recordings	25
2.2.3	Local and global stimulation data	26
3	Modelling Methods	30
3.1	Neural Modeling	30
3.2	Modelling the ELL	33
3.2.1	Superficial pyramidal neurons	33
3.2.2	Electroreceptors	33
3.2.3	Depolarizing after-potential	34
3.2.4	Granule cells	36
3.2.5	Pyramidal cell-parallel fibre synaptic plasticity	37
3.2.6	Potentiating rule	39

3.2.7	Disynaptic inhibition	39
3.2.8	Summary	40
3.3	Computer Simulations	41
3.3.1	Simulating the noise	41
3.3.2	Simulating the weights	42
4	Parameter Fitting	45
4.1	Local Stimulation	45
4.2	Global Stimulation	48
4.3	Granule Cell Activity	52
5	Frequency-tuned cerebellar channels and burst-induced LTD lead to the cancellation of redundant sensory inputs	57
5.1	Contribution Statement and acknowledgements	57
5.2	Abstract	58
5.3	Introduction	58
5.4	Materials and Methods	60
5.4.1	LIF model	60
5.4.2	Modeling the bursting dynamics	61
5.4.3	Local stimuli inputs (feedforward)	62
5.4.4	Parallel fibre inputs (feedback)	62
5.4.5	Burst definition and learning rule	63
5.4.6	Quantitative analysis	63
5.4.7	<i>In vivo</i> electrophysiology	64
5.5	Results	65
5.5.1	Model replicates experimentally observed responses in the absence of feedback	65
5.5.2	Incorporating a feedback that cancels low frequency coding	67
5.5.3	Independent parallel fibre frequency channels revealed <i>in vivo</i>	72
5.6	Discussion	75
5.7	Figures and Legends	79
6	Analytics	80
6.1	Local Stimulation	80
6.1.1	Behaviour	80
6.1.2	Theoretical Analysis	82

6.2	Global Stimulation	90
6.2.1	Qualitative investigation of \bar{w}	90
6.2.2	Qualitative investigation of the weight distribution	94
6.2.3	Theoretical Analysis	97
6.2.4	Power laws	104
6.3	Contrast Invariance	106
7	Discussion	113
7.1	Both 2- and 4-spike burst rules are required for optimal cancellation . . .	113
7.2	Biological Realism	115
7.3	Model Predictions	116
7.4	Theoretical Analysis	117
7.5	Comparison to other Systems	120
8	Conclusion	123
8.1	Summary of Results	123
8.2	Future Directions	124
	Nomenclature	126

List of Tables

2.1	Phase shifts between data and stimulus	27
3.1	DAP parameter values	36
5.1	Parameter values for the optimal model	60
5.2	DAP parameter values	62
5.3	$\kappa(f)$ for different AM frequencies	62

List of Figures

1.1	Schematic of a neuron's anatomy	4
1.2	Schematic of a neuron's membrane potential	5
1.3	A plot of electroreceptor firing rate as the stimulus intensity is varied . .	6
1.4	Examples of different experimental STDP rules	9
1.5	Example of the ambiguity of using STDP rules in an active network . . .	10
1.6	Schematic of the mammalian cerebellum	12
1.7	Overview of weakly electric fish	15
2.1	Schematic of local and global stimulation of <i>Apteronotus leptorhynchus</i> .	18
2.2	Superficial cell activity during local and global stimulation	19
2.3	Schematic of ELL circuitry in <i>Apteronotus leptorhynchus</i>	20
2.4	An outline of the variety of anatomical delays in the feedback pathways .	22
2.5	<i>In vitro</i> single spike and burst-induced synaptic plasticity	23
2.6	Example of a post-stimulus time histogram (PSTH)	24
2.7	Temporal organization of spikes and definition of a burst	25
2.8	<i>In vivo</i> experimental data of the superficial cell response to local stimulation	28
2.9	<i>In vivo</i> experimental data of the superficial cell response to global stimulation	29
3.1	Schematic of the modelled feedback pathway	38
3.2	Model Summary	40
3.3	Effect of changing the integration time-step in the model when noise is rectified	42
3.4	Weight evolution during global stimulation	43
4.1	Parameter fitting to mimic experimental data under local stimulation . .	47
4.2	Optimal model fit of experimental data under local stimulation	49
4.3	<i>In vitro</i> burst-induced depression at the superficial cell-parallel fibre synapse	50

4.4	Parameter fitting to match experimental data under global stimulation when both burst learning rules are in effect	51
4.5	Optimal global stimulation fits when only the 2-spike burst rule is active, only the 4-spike burst rule is active, and when both are active	53
4.6	Model behaviour during global stimulation when Λ decreases at high AM frequencies	54
4.7	Model behaviour in the global η model	55
4.8	Values of Λ and η at each AM frequency during global stimulation to optimally fit experimental data	56
5.1	Properties of the system to be replicated by the model	66
5.2	Model replication of local stimulation and global data	68
5.3	Burst-induced depression measured <i>in vitro</i> with model fit	70
5.4	Comparison of model and data during global stimulation with unconstrained granule cell bursting	71
5.5	Decreasing η to mimic decreased granule cell bursting at high AM frequencies	72
5.6	Comparison of model and data during global stimulation	73
5.7	Experimental evidence of frequency channels	76
6.1	General behaviour of the model under local stimulation	81
6.2	R-I curve of stochastic LIF model and the first passage time approximation	83
6.3	Model behaviour with low-pass filtered noise and analytical fit with FPT equation	84
6.4	Model behaviour with low-pass filtered noise and analytical fit with low-pass filtered noise correction	85
6.5	Model firing rates when input is rectified and analytical approximation with and without the rectification correction	87
6.6	Analytical fit of model behaviour during local stimulation without a DAP	88
6.7	Analytical fit of model behaviour during local stimulation with a DAP .	90
6.8	Typical simulation of the average weight value during a single run of global simulation	92
6.9	Weight evolution during global stimulation in the final global model . . .	95
6.10	Weight evolution during global stimulation when $\kappa(f)$, η_2 and η_4 are constant	96
6.11	R-I curve of the model with different values of inhibition	99
6.12	Br-R curve of the model during global stimulation	100
6.13	Schematic of the dynamics of the average weight value at equilibrium . .	100

6.14	A plot of \bar{w}^{bin} versus the equilibrium burst rates for constant inputs . . .	102
6.15	Analytical fit of the model during global stimulation	103
6.16	Analytical fit of the model's weight distribution during global stimulation	104
6.17	Equilibrium mean weight value, averaged over an AM cycle, plotted against the burst rate during global stimulation	105
6.18	The power law dependence between the AM frequency and the average firing rate and average weight	106
6.19	Evidence of power laws holding over variations in the AM frequency . . .	107
6.20	Average weight distributions learned at different contrasts under global stimulation	108
6.21	Investigation of contrast invariance in the model when the weights are learned at 75% contrast	109
6.22	Investigation of contrast invariance in the model when the weights are learned at 25% contrast	110
6.23	Investigation of contrast invariance in the model when the weights are learned at 25% contrast and the relative change in contrast is taken into account	111
6.24	<i>In vivo</i> data of the AM contrast as a function of the distance between two fish	112

Chapter 1

Introduction

The complex nature of the brain has made neuroscience one of the least explored scientific disciplines. Although the macroscopic functions of cerebral regions have been well studied, it is still unclear how the brain accomplishes these tasks using neurons. This is made all the more challenging by the brain's intricate and highly-connected structure, where specific inputs and outputs to an area are difficult to identify and even more difficult to apply meaning to, functionally. The study of sensory processing centres can simplify these issues, since at least the inputs to such regions are known and controllable. Furthermore, recent advances in experimental technology have clarified some of the essential microscopic processes that operate within the brains of certain model organisms. Investigating sensory processing in these creatures presents a unique opportunity to better understand how the brain performs computations using neural networks.

1.1 Sensory Processing and Adaptive Filtering

The brain has few tasks more important than sensory processing. An organism must be able to quickly scan a noisy environment and detect weak, but potentially important, signals with high precision. Fortunately, the brain is remarkably good at this feat. One of the best examples is in the auditory system, the so-called “cocktail party problem:”

One of our most important faculties is our ability to listen to, and follow, one speaker in the presence of others. This is such a common experience that we may take it for granted; we may call it “the cocktail party problem.” No machine has been constructed to do just this, to filter out one conversation from a number jumbled together. [1]

This is a powerful form of tunable noise suppression, and, although it is not fully understood, this skill appears to be based on non-linear computations in the sensory processing regions of the brain [2]. Unfortunately, auditory noise cancellation is difficult to investigate because of the complex statistical structure that distinguishes each voice in an environment.

Nonetheless, previous analyses have concluded that auditory and other similar sensory signals are likely processed in the brain using an adaptive filter, which preferentially removes redundant background input in order to detect weak, novel stimuli [3, 4]. The general structure of a filter involves a decomposition of the input signal into its basic elements by a filter bank. These elements are then recombined with different weightings to create the output. Crucially, however, an adaptive filter will adjust its own parameters according to a learning algorithm to optimize performance.

In the case of cancelling predictable signals for sensory processing, the optimal filter bank would separate the input into independent frequency bandwidths. If identified prior to processing, the redundant signal could be removed by adding the predictable signal in anti-phase to the original input, as is done *e.g.* in noise-cancelling headphones. Adaptive filters are very powerful processing tools and can be applied to a range of biological problems, such as motor control, sensory processing, and state estimation [5]. In many organisms, the neural connectivity of the region known to perform this task has been well characterized, and often has a cerebellar architecture [4, 5]. The question, then, is how a biophysically based neural network can operate as such a filter. This is an open problem in bioacoustics and in auditory neuroscience. A specific partial solution to this problem is the focus of this thesis.

1.2 Neurons

The fundamental unit in the brain is the neuron (Figure 1.1). An archetypal neuron is composed of highly branched appendages, called dendrites; the cell body, called the soma; and a long, thin projection called an axon (for more information, see [6, 7]). Neurons, like every cell, have an exterior membrane that is impermeable to ions. Due to this membrane, there is a voltage difference between the extracellular fluid and the cell's interior, which is regulated by the opening and closing of ion channels in the membrane. The critical feature of neurons, however, is the regulation of the membrane potential for computation and signalling. This is accomplished using unique, non-linearly activating ion channels in the soma and axon of neurons that only open once the membrane potential

in the vicinity of these channels is above a certain threshold. At rest, a neuron's interior is typically at a potential 70 mV lower than its exterior, and fluctuations of the membrane potential will decay exponentially back to this resting potential given sufficient time [6].

If enough current is injected into the neuron, however, the membrane potential will rise above threshold and these voltage-gated channels will start to open. As they open, it will depolarize (*i.e.* increase) the membrane potential more, causing more channels to open in a chain reaction. This will produce an action potential: a sudden and strong depolarization of the membrane. This depolarization usually starts in the soma but will, by diffusion of the ions, depolarize the initial segment of the axon, which also has these unique channels. Their activation will cause a rapid depolarization in this region and depolarize the next axonal segment. Sequential activation of these channels allows the action potential to quickly propagate down the axon of the neuron (see *e.g.* [8] for a further explanation). The speed of action potential propagation can be significantly increased if the axon is sheathed in a lipid-based insulator, called myelin.

During the action potential, slower but more powerful hyperpolarizing voltage-gated channels also open, which eventually dominate the depolarizing channels and cause the membrane potential to fall back below threshold to its resting potential. The voltage-gated channels then close and reset in an interval called the refractory period, during which the neuron cannot create another action potential. Afterwards, the neuron is ready to fire an action potential again.

In vivo, subthreshold current is injected into a neuron from ion channels activated by synapses [6]. A synapse is a one-way connection between two neurons, and the neuron receiving the input is called the post-synaptic neuron while the neuron supplying the input is the pre-synaptic neuron. If an action potential occurs in the pre-synaptic neuron, the synapse will be activated and transmit chemical messengers across the synaptic cleft, which open nearby ion channels in the post-synaptic neuron. The classical synapse is axo-dendritic, which means that the synapse is between the pre-synaptic neuron's axon and attaches to the dendrite of the post-synaptic neuron. Synapses can either depolarize the post-synaptic membrane potential (excitatory synapses) or hyperpolarize the membrane potential (inhibitory synapses), depending on the ion channels they activate. The ion currents from each activated synapse, called post-synaptic potentials, flow through the dendrites and are summed in the soma, the cell body of the neuron. If enough excitatory synapses are active, the neuron will fire an action potential (Figure 1.2).




Figure 1.1: Schematic of a neuron’s anatomy. Inputs from other neurons synapse at the neuron’s dendrites. When a synapse is activated by a pre-synaptic action potential, the pre-synaptic neuron releases chemicals across the synaptic cleft that open ion channels in the post-synaptic neuron that will change the post-synaptic neuron’s membrane potential. These inputs are summed in the soma and, if beyond a certain threshold, will cause the neuron to fire an action potential that will electronically propagate down the neuron’s axon. The axon ends in synapses that in turn supply inputs to other neurons. From [6].

1.3 Neural Encoding

In general, action potentials are stereotyped: they have the same temporal shape and voltage amplitude regardless of the neuron or the size of depolarization as long as it is above threshold [6]. The effect of an action potential on a synapse, on the other hand, is not stereotyped, and the strength of a synapse (*i.e.* the size of the PSP generated in the post-synaptic neuron by the synapse’s activation) can take different values. Linked together, neurons will send and receive input via action potentials and the effect of those inputs on each neuron is weighted by its synaptic strengths. This parallels an electronic network with neurons acting as independent processors and sequences of action potentials like digital signals.

Cerebral neurons receive information about the external environment via sensory neurons. Each sensory neuron responds to a single stimulus modality (*e.g.* cutaneous,




Figure 1.2: Schematic of a neuron’s membrane potential when receiving synaptic inputs as measured in the soma. The resting membrane potential (RMP) of this hypothetical neuron is -70 mV. An inhibitory post-synaptic potential (IPSP) will cause the membrane potential to hyperpolarize, or decrease. An excitatory post-synaptic potential (EPSP) will depolarize, or increase the neuron’s voltage. If an EPSP is large enough, it will bring the neuron’s voltage beyond threshold (arrow in diagram) and cause it to fire an action potential, AP. From [6].

auditory, visual, *etc.*) and putatively encodes the characteristics of the stimuli into action potential sequences, called spike trains, for further computation in the brain. However, how neurons encode information into their spike trains is highly controversial [7]. Nevertheless, the method must be robust to noise. Due to the random opening and closing of ion channels and the vast number of inputs each neuron receives, the firing activity of most neurons observed *in vivo* often appears stochastic. Therefore, the average firing activity of each neuron is commonly assumed to be the important statistic for information transmission [9].

More intense input stimuli is hypothetically encoded as an increased spike firing rate, which is decoded by downstream neurons (Figure 1.3). These downstream neurons process the input and encode their response into firing rates that will be interpreted by higher level centres. Note that the existence of a firing threshold makes neural encoding inherently non-linear. Furthermore, at very low and very high stimulus inputs, the firing rate dies out or saturates due to non-linearities in spike production (Figure 1.3). But if information is transmitted in spike trains and computed in neurons, where is it stored in a neural network?




Figure 1.3: A plot of the firing rate of electrosensory neurons in weakly electric fish, *Aptereronotus leptorhynchus*, as the relative stimulus intensity, or contrast, is varied. These neurons code for the strength of electric signals in the water. Here, the stimulus is low frequency ($f_{cut} = 4$ Hz) noise. Spike times were collected for 32 ms after stimulus onset to avoid adaptation (*i.e.* a gradual decrease in the firing rate when a stimulus is presented and maintained). A linear fit (solid black line) accurately reflects the data within the “coding region” as identified by the vertical lines within the plots. The mean rate of these neurons, when no stimulus is present (dashed lines), is in the middle of their linear regime, maximizing sensitivity. Adapted from [10].

1.4 Neural Learning and Information Storage

Analogous to a computer, there are two types of information stored in an organism: immutable information about the system and environment (*i.e.* read-only memory) and current information about the system and environment that can be changed (*i.e.* random-access memory). Presumably, the anatomy and genetic code of the individual bestow some inherent information about the organism at inception, and so data on basic physiological functions are hardwired into the cerebral architecture. Although it requires this information, the brain’s fundamental purpose is not to accumulate this type of data. The brain’s function is to store, recall, and update new information about the state of the individual and its surroundings. In other words, it learns.

Even before the underlying biological mechanisms were understood, many theorists have contemplated what method the brain could use for learning. One of the most influential algorithms was first postulated by Donald Hebb in 1949 [11]. It is based on a theoretical conjecture about learning: namely, that it relies on the identification of correlations between events. For example, if you touch a hot stove element, you feel pain. To learn from this action, the brain must identify the correlation between motor

movement and pain receptor activity. If these external events had a neural representation, then learning at the neuronal level would manifest in an increase in the correlation of one neuron's activity and another. With repeated pairings, it is possible that touching a stove element would register pain regardless of the temperature of the object. This is the crux of Hebb's rule:

When an axon of cell A is near enough to excite cell B or repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased [12].

The mechanism for stable alteration of neural correlations was finally found by Bliss and Lømo in 1973 [13]. Long-term potentiation (LTP) and long-term depression (LTD) is a stable increase or decrease, respectively, of the strength of a synapse that depends on the firing activity of the pre- and post-synaptic neurons. Such synaptic modification allows the brain to store information, and changes in synaptic strengths induced by neural activity allows the brain to learn. With different synaptic weightings, neurons will respond differently to an identical input and perform a slightly different calculation. Within a network, synaptic plasticity can coordinate the firing activities of dissimilar neurons such that if one spikes, the other spikes as well. This is the fundamental biological substrate of classical conditioning reflexes [7].

In a similar argument that led to the rate coding hypothesis of external stimuli, the mean firing rates of neurons was initially assumed to be the relevant description for the firing activities of neurons to drive plasticity, and, hence, learning (see [14] and references therein). The individual spike times of neurons can then be histogrammed to create a continuous firing rate, and the correlations between pre- and post-synaptic firing rates drive synaptic plasticity. For example, the synaptic strength increases when the pre- and post-synaptic spike trains are positively correlated, and the strength decreases when the spike trains are negatively correlated. Within a network, these rate-based rules could "tune" each individual synapse to a specific value and, hence, store information. Neural models based on this hypothesis were able to learn a variety of tasks, and to stably store many different input patterns [14, 15].

However, neural computation would be delayed if every neuron was required to average its inputs over time before reacting. For survival, organisms must react quickly to a sensory stimulus, and experiments have found that an animal's response is often based on the order of one spike from their sensory and higher order neurons [9]. Averaging

over identical inputs instead of time is also not a viable solution since, despite a large number of neurons in vertebrates, there are many invertebrates that do not have an abundance of neurons but still react swiftly. Furthermore, investigations into the limits of perception of humans have observed responses based on single spikes from sensory neurons. If individual spikes are important carriers of information, then learning based on mean firing rates only may not be complete.

1.5 Spike-Timing Dependent Plasticity

In 1997, Markram and Sakmann found that LTP and LTD were profoundly affected by the timing of individual pre- and post-synaptic spike pairs on a millisecond timescale [16]. In other words, when a pre-synaptic neuron fires in close temporal proximity to the firing of a post-synaptic neuron, the synapse that connects them will permanently change its strength and the magnitude and direction of the change depends on the time delay between the spikes, not on the average spike rate. These spike-timing dependent plasticity (STDP) rules replaced rate-based learning rules as the most biophysically realistic mechanism for synaptic plasticity.

STDP rules have been subsequently discovered in a diversity of shapes, and each is presumed to bestow properties on the network it is situated in (Figure 1.4) [11]. For example, an STDP rule could increase causal relationships in a network if it strengthened the synapse when the pre-synaptic neuron fired before the post-synaptic neuron and weakened the synapse when the post-synaptic fired before the pre-synaptic neuron. This would select for connections where the pre-synaptic neuron induced the post-synaptic neuron to fire, and attenuate connections where the pre- and post-synaptic neurons have no correlation. Due to the close correspondence with his initial postulate, an STDP rule with these qualities is called a Hebbian rule (Figure 1.4A and B). On the other hand, an anti-Hebbian rule would decrease the correlation in a network (Figure 1.4D). More generally, STDP rules with LTD greater than LTP can tune the membrane potential of a cell to just under its threshold, maximizing sensitivity [17]. In addition, anti-Hebbian learning rules can stabilize positive feedback loops, enabling networks to perform as neural integrators [18].

STDP has been extended recently by Pfister and Gerstner to include spike triplets [23]. In this model, spike triplets, one pre-synaptic spike and two post-synaptic spikes, bestow an increased synaptic effect compared to that predicted by pairing alone, and the combination of pairs and triplets successfully matched experimental data from the

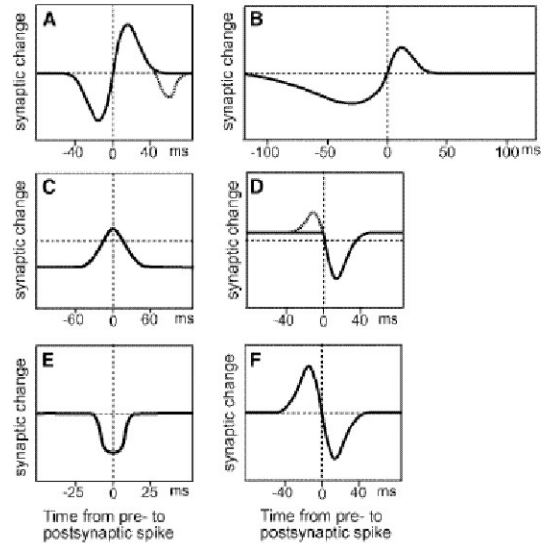


Figure 1.4: Examples of different experimental STDP rules in different cerebral areas. On the ordinate is the time delay between pre- and post-synaptic spikes, the vertical dotted line representing both spikes occurring simultaneously. **A** and **B** are classical Hebbian rules discovered in the mammalian neocortex [16] and the hippocampus [19], respectively. **D** is an anti-Hebbian rule found in a cerebellum-like structure in weakly electric fish [20]. **C** and **E** are symmetric rules discovered at the neuromuscular junction [21] and in spiny stellate neurons in the neocortex [22], respectively. **F** is a theoretical anti-Hebbian rule with associated potentiation and depression. From [11].

visual cortex; a region where STDP based solely on spike pairings failed. Investigations into bursts, brief sequences of many action potentials, have also revealed non-linear contributions of spikes beyond the pairing window [24]).

In order for STDP learning to be successful, however, the synaptic strengths must be relatively stable after plasticity has adjusted the weights to their learned distribution. This is less tractable than in rate-based models since there are far more fluctuations in the individual spike times than in the average firing rate. Learning stability also places restrictions on the types of STDP rules that should be discovered. For example, STDP rules must have equal areas of depression and potentiation; otherwise, the synapse would be eventually driven to its maximum or minimum values by the stochastic nature of neural firing [11]. Presumably, a modifiable synapse is meant to be tuned to an appropriate value by learning, and if all synapses attain identical values then synaptic plasticity is irrelevant. Nevertheless, unbalanced rules have been discovered (Figure 1.4C and E), which suggests that STDP may not fully characterize the plasticity of a synapse.

There are also difficulties in the implementation of rules based on individual spike times in a continuously active network. STDP rules are typically generated with individual pre- and post-synaptic pairings, but what is the correct rule when both neurons are producing spike trains? Should a pre-synaptic spike be paired with every post-synaptic spike? Or do post-synaptic spikes take precedence? When should spikes be grouped into triplets? Into bursts? Recent experimental work has shown that “nearest neighbour” approaches (pairing each pre-synaptic spike to the closest post-synaptic spike or vice versa) best reproduce the observed synaptic plasticity, but there are still a variety of ways to pair neighbouring spikes and each brain region appears to operate under a different convention (Figure 1.5) [25].




Figure 1.5: Example of the ambiguity of using STDP rules in an active network. i and j represent a spike train from a post-synaptic neuron and a pre-synaptic neuron, respectively. This is a common illustration of the activity of a neuron over time with a spike as a vertical bar whenever the neuron fires. Lines linking spikes between neurons i and j represent a pairing that induces plasticity with light grey lines indicating a pre-before post- pair and dark grey lines indicating a post- before pre- pair. The order of the spikes can often change the sign of plasticity (see Figure 1.4). **A** depicts a symmetric interpretation: each pre-synaptic spike is paired with the last post-synaptic spike and vice-versa. **B** depicts a pre-synaptic centred interpretation: each pre-synaptic spike is paired with the last post-synaptic spike and the next post-synaptic spike. **C** depicts a reduced symmetric interpretation: pre-synaptic and post-synaptic spikes are paired symmetrically as in **A** but only immediate pairings cause plasticity. All of these are nearest neighbour approaches, and each has been found to fit different experimental data. From [25].

1.6 Supervised Learning

Besides STDP, there are other methods hypothesized to function in the brain that could teach neurons the appropriate correlations between disparate events. Supervised learning is an alternative method that first employs a learning phase where both the input and the correct output are provided to the naïve network [7]. The network’s task is then to create an algorithm that translates input stimuli to output stimuli so novel inputs will still generate appropriate outputs. This is distinct from standard STDP systems where learning is unsupervised: the neuron, whose synapses are being adjusted, is unaware of the correct response. Supervised learning is advantageous theoretically since the learning algorithm is straightforward, and, because learning is controlled externally, the learnt correlations are inherently stable.

In neural networks, the correct output would be given by a “teacher” signal, a powerful input that will dictate the cell’s activity. For example, if the cell needs to learn to fire when two specific inputs arrive, the teacher signal will also arrive with the two inputs and forcefully depolarize the cell. The neuron must then adjust so that, after repeated pairings with the teacher, the two specific inputs alone are sufficient to fire the cell. The biological mechanism underpinning this learning is still the long-term alteration of synapses, but the presence of a teacher signal that gates the synaptic modifications creates supervised learning. This has been the prevailing model of the mammalian cerebellum to explain how it learns motor patterns [5].

1.7 Cerebellum

The unique attribute of cerebellar structures is a massive number of fibres that run parallel to each other and synapse upon an output array of cells layered perpendicularly to the fibres (Figure 1.6). These synapses are plastic and exhibit both long-term potentiation and long-term depression [26]. Each parallel fibre originates from a granule cell, a neuron that receives sensory input of different modalities from mossy fibres. Although the number of granule cells is immense, each granule cell will only activate for a specific confluence of sensory inputs, leading to very few cells being active at any one time. These granule cells putatively supply “context” to the output array of neurons. Theoretically, this so-called sparse coding of the granule cells, coupled with the maximum number of synaptic contacts between granule cells and output cells, is a favourable architecture for information transfer and learning [27, 28].




Figure 1.6: Schematic of the mammalian cerebellum. **A** illustrates a region of the cerebellum *in vivo* while **B** simplifies the cerebellar circuitry. Granule cells in the granule cell layer receive input from mossy fibres and project into the molecular layer, where their axons form a ‘T’ and create parallel fibres. These parallel fibres then form excitatory synapses onto purkinje cells, the output neurons of the mammalian cerebellum. In addition, purkinje cells receive a very strong excitatory input from climbing fibres, which synapse with a single purkinje cell numerous times. Golgi, stellate, and basket cells are local neurons also found in the cerebellum that receive parallel fibre input and form inhibitory synapses onto purkinje cells (shown only in **A**). Purkinje neurons project to the deep cerebellar nuclei that control motor activity. From [29].

In the mammalian cerebellum, each parallel fibre synapses onto ~ 400 purkinje cells, the mammalian output layer cell, and each purkinje cell receives input from $\sim 200,000$ parallel fibres [28]. Due to a unique one-to-one correspondence, each mammalian purkinje cell receives very strong excitatory input from one cell in the inferior olivary nucleus.

An important feature of the cerebellar physiology is the control of synaptic plasticity at the parallel fibre-purkinje cell synapse by the inferior olivary nucleus. If both the parallel fibre and purkinje cell are simultaneously active, the synapse will exhibit spike-timing dependent potentiation. However, if the parallel fibre input coincides with the inferior olivary nucleus’ activity, then the parallel fibre-purkinje cell synapse will exhibit spike-timing dependent depression [11]. Since the inferior olivary input is an external signal, the cerebellum exhibits a form of supervised STDP learning with the inferior

olivary nucleus as the teacher. The function of cerebellar learning was thought to detect patterns of sensory activity to coordinate motor commands [27].

Recently, a new hypothesis has emerged about the function of the cerebellum: that it operates as an adaptive filter [30, 5]. If granule cells could perform as highly selective sensory filters, then the plastic parallel fibre network would supply an ideal input framework to purkinje cells that neural learning can shape for important applications such as memory storage [31] or signal transformation [3]. However, both adaptive filtering and pattern recognition theories are still controversial, as not all of the components of either theory are discovered. For example, the delay window for synaptic depression at the parallel fibre-purkinje cell synapse is tightly constrained theoretically, but a much wider window has been found in mammals [11]. In addition, appropriately strong potentiation has yet to be found at the synapses, making unlearning or “forgetting” a pattern impossible.

Cerebellar structures have been found in other systems as well, and without the unique input from an inferior olivary nucleus analog. This has led to a reconsideration of the function of the cerebellar architecture and the importance of the IO nucleus. One hypothesis is that cerebellar structures may not be pattern detectors or learning devices but suppliers of precisely timed signals. Due to the slow transmission velocities of parallel fibres and the sequential layering of purkinje cells, the cerebellum could accurately identify temporal delays between different sensory inputs. For example, auditory localization in teleosts, a species of fish, occurs in a cerebellum-like structure [32]. Granule cells driven by different ears would fire at different times based on the location of an auditory stimulus, and the position of the purkinje cell that received inputs from both parallel fibres coincidentally would code for the relative delay [32]. Evidence of STDP at the parallel fibre-purkinje cell synapse in some species, however, has led back to the hypothesis of adaptive filtering, but, importantly, unsupervised learning is exploited instead of supervised learning to teach the synaptic strengths [4, 3].

1.8 Weakly Electric Fish

One of the organisms that employs an unsupervised learning rule in its neural sensory processing is the weakly electric fish. These fish emit electronic signal called an electric organ discharge (EOD) to sense their surroundings (Figure 1.7). Operating similarly to echolocation, this discharge interacts with objects in the water, creating amplitude modulations (AMs) of the EOD [33, 34, 35]. These signals are sensed by electroreceptors

on the fish’s skin, which linearly encode the AMs into their firing rate modulation and transmit this signal to neurons in the sensory processing centres in the brain [36, 10]. The EOD is used to both search for prey and communicate to conspecifics (*i.e.* other fish of the same species). Unfortunately, these two tasks often interfere. For example, communication signals are “global” signals, as they have a large spatial extent, and will produce a strong AM over the entire surface of the fish. This can overpower weaker localized stimuli, such as prey, that may also be present [37, 34].

Although electroreceptors always faithfully encode the amplitude modulations of the EOD, neurons in the electrosensory processing station in the fish’s brain, a cerebellar-like structure, are known to become desensitized to global predictable inputs [38]. In other words, this network identifies and eliminates extraneous signals to maximize the detection of novel stimuli, which solves the cocktail party problem. Investigations into the electrosensory system are aided by the simplicity of the input stimulus: a sinusoid. In addition, electronic signals can be easily recorded and modified for precise experimental control of the input to the neural network. Furthermore, brain homology in basic sensory processing allows results from weakly electric fish to be transferred to more advanced neural systems, such as the mammalian cerebellum and the dorsal cochlear nucleus [3]. Thus, weakly electric fish present an ideal model organism to explore cerebellar-like structures and adaptive filtering in the brain.

A previous study by Roberts and Bell has deduced how mormyrids, a species of weakly electric fish, cancel global redundant signals using a neural network [40]. These fish emit a discrete, pulsatile EOD and adaptively cancel signals that constantly recur after each pulse, putatively by an STDP rule that is found in their sensory processing centre. Roberts and Bell showed through computer simulations that the STDP rule together with the neural architecture in the mormyrid brain is sufficient to reproduce the adaptive cancellation observed *in vivo*. This is a rare study that determined how a physiological function was executed using a realistic neural network, and the present investigation follows the same approach.

In this thesis, adaptive cancellation in another species of weakly electric fish, apteronotids, is explored. These fish produce a continuous EOD, which complicates the identification and removal of predictable signals. Furthermore, because of beat signals between conspecific’s EODs, effective cancellation in apteronotids must occur at a range of frequencies, not just the pulse frequency. Analogous to mormyrids, a plasticity rule has been found in the apteronotid electrosensory processing network, but it is a unique spike-timing rule that has not been found in any other organism. Unlike previous STDP

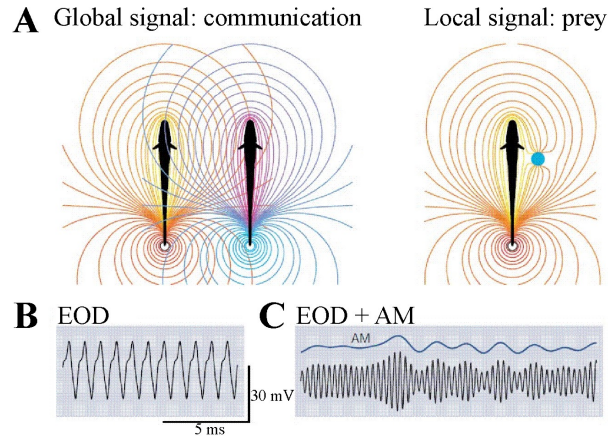


Figure 1.7: Overview of weakly electric fish. Their electric sense is used for both communication (a global signal) as well as prey detection (a local signal), shown in **A**. Coloured lines represent iso-potential surfaces. A plot of an apteronotid’s typical EOD is shown in **B**, which maintains a constant frequency and amplitude. Stimuli from the environment create amplitude modulations of the EOD carrier, shown in **C**, which are encoded by sensory neurons on the fish’s skin and transmitted to higher brain centres. Adapted from [39].

models, single pre-synaptic and post-synaptic spike pairings have no effect on the synaptic strength [41]. However, pairing pre- and post-synaptic bursts do alter the synapse, and the size of the burst (*i.e.* the number of spikes it contains) non-linearly affects the magnitude of the induced plasticity. This is a novel form of plasticity whose ramifications for network function are unknown. Nonetheless, the function of the network and its inputs are well-documented, allowing a connection between a synaptic mechanism and overall network behaviour.

Results from weakly electric fish can be extended to other cerebellar structures and to all biological systems that putatively function as adaptive filters. The study of apteronotid input cancellation also extends the concept of unsupervised spike-time dependent plasticity rules to include groups of spike events and not just single spikes. This is an opportunity to investigate the computational significance of unique plasticity rules, to gain insight into cerebellar structures, and to understand the implementation of sensory processing functions in neural networks.

1.9 Organization of Thesis

The goal of this thesis is to investigate how adaptive filtering is achieved in the electrosensory processing station in *Apteronotus leptorhynchus*. The hypothesis is that the known architecture of this neural network combined with the novel plasticity rules is sufficient to recreate the adaptive cancellation seen *in vivo*. To this end, a minimal but biophysically realistic computational model of the cancellation circuitry in the ELL is constructed. This model will be compared to experimental data collected by electric fish physiologists at the University of Ottawa. Although this structure is known to cancel global signals with a finite frequency bandwidth, for simplicity the sensory processing task that the model will be charged to reproduce is the cancellation of monotone stimuli over the frequency range where cancellation is known to occur. This study also attempts to answer the following questions:

1. Are either or both of these burst rules sufficient to reproduce the adaptive cancellation seen *in vivo*?
2. What predictions of the model can be made about the biological system?
3. What is the computational effect of multiple burst rules operating concurrently?
4. Can the model's dynamics be predicted analytically?

Chapter 2 provides a background on the biological circuitry of the ELL and the experimental data. It also identifies the protocols to obtain the data, and how the data and the model will be compared. In Chapter 3, the model is constructed and the numerical methods that simulate the model are outlined. Chapter 4 reviews the procedure to fit the model to the experimental data. Chapter 5 is a submitted manuscript that details the results of the model and its predictions about the dynamics of the biological circuitry. In Chapter 6, the behaviour of the model itself is investigated analytically to explore the effect of the burst plasticity rules and periodic cancellation feedback in general. The results of this thesis are compared to other work and placed into context in the discussion and finally summarized in the conclusion.

Chapter 2

Biological Background and Experimental Data

This chapter presents a brief background on the physiological details of weakly electric fish, their electrosensory input, and the circuitry that putatively causes sensory input cancellation. Experimental protocols used to collect the data *in vivo* are explained, the data is presented, and the statistics used to compare the data to the model are described. Experimental data was recorded by a collaborator, Dr. Gary Marsat, a postdoctoral fellow working in Prof. Maler's lab at the University of Ottawa.

2.1 Weakly Electric Fish

Apteronotus leptorhynchus is a species of weakly electric fish that continuously emits a high-frequency (600-1000 Hz) sinusoidal electric organ discharge (EOD) into its environment to sense its surroundings and communicate to conspecifics. Small objects in the environment such as prey will create spatially localized amplitude modulations (AMs) of the EOD, whereas tailbending or communication signals will induce spatially global AMs [33, 34, 35]. AMs are detected by electroreceptor afferents that densely cover the body of the fish [36, 10], and linearly encode these signals into their firing rate modulation (see Figure 1.3). Electroreceptors provide feedforward input to pyramidal cells in the electrosensory lateral line (ELL), the first electrosensory processing structure in the brain [42, 43]. Pyramidal cells then project to deeper brain centres for higher level cognitive functions.

Each pyramidal neuron receives electroreceptor input from only a small region of the

fish’s skin, called a receptive field. Based on their response to electroreceptor activity, pyramidal neurons are classified into two groups: E cells, which direct excitatory electroreceptor input, and I cells, which receive electroreceptor input through an inhibitory cell. Therefore, E cells increase their firing rate to increasing sensory stimuli, while the response of I cells is inverted.

When the stimulation is a small object exciting the receptive fields of only a few pyramidal cells (local stimulation), their firing activity will be strongly modulated by the AMs generated by the stimulus [44]. These pyramidal cells will also burst more frequently, an event implicated in increased information transfer to downstream neurons [45]. As very few neurons code for this stimulus, local stimulation is a novel event, and it is faithfully transmitted to higher brain centres. However, when the stimulation is a large AM across the entire body of the fish (global stimulation), such as two EODs of different frequencies beating together, many pyramidal receptive fields are being similarly stimulated and encoding the same information. The removal of global signals from at least some of the neurons would therefore aid in identifying local stimuli that may also simultaneously exist.

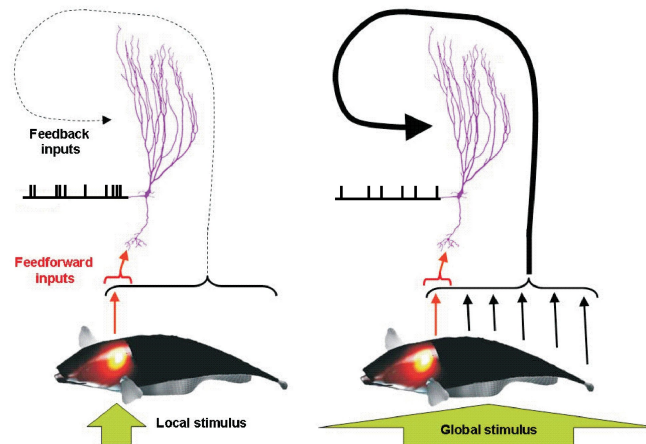


Figure 2.1: Schematic of local and global stimulation of *Aptereronotus leptorhynchus*. The purple neuron is a superficial pyramidal cell. Small, local stimuli such as water fleas, the fish’s prey, in the receptive field of a superficial cell will cause a strong modulation in the firing rate of the pyramidal cell and induce bursting. If the stimulus has the same intensity but is global in extent, the superficial cell’s firing rate is only weakly modulated and burst production is minor. In both cases, the electroreceptor input to the superficial pyramidal cell is identical. However, the feedback pathway is only active during global stimulation. This feedback putatively cancels the feedforward signal by creating a “negative image” that destructively interferes with the electroreceptor input.

Previous studies have found that a subpopulation of the E and I pyramidal cells in the ELL do adaptively cancel global, predictable signals in their electroreceptor input (Figure 2.1). These cells are called superficial cells as they are found in the most dorsal section of the ELL [46]. Cancellation mitigates the firing rate modulation of these cells compared to the modulation induced by a local stimulus of the same strength and AM frequency, which reduces the number of bursts generated and putatively changes the signal to downstream neurons. The strength of the cancellation response, however, weakens as the AM frequency increases (Figure 2.2). Although the functional reason for this cancellation decay is unknown, effective cancellation is importantly still maintained in the 0-16 Hz bandwidth, where prey signals are most likely to appear [33].

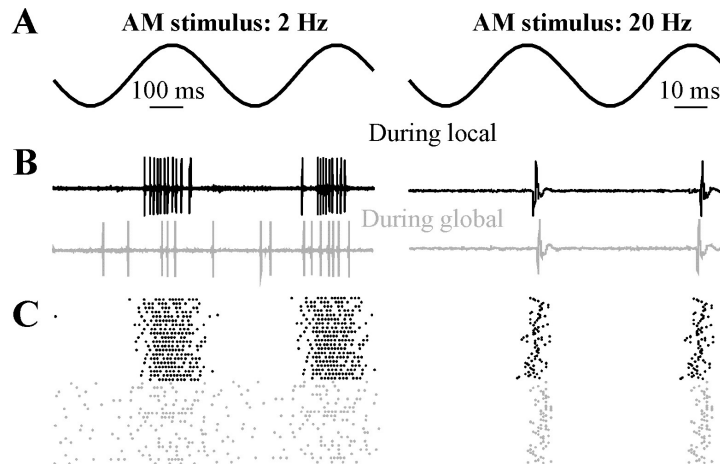


Figure 2.2: Superficial cell activity during local (black) and global (grey) stimulation. **A** depicts two input signals of different frequencies: a 2 Hz amplitude modulation of the EOD and a 20 Hz AM. **B** illustrates the typical experimental recording of the superficial neurons by an electrode during two cycles of stimulation. After many stimulus presentations, raster plots (**C**) of the firing times with respect to the phase of the input stimulus can be created. The bursty nature of the neuron’s response to local stimulation is evident. Global stimulation mitigates the fluctuations of the firing rate, but only at low frequencies: at 20 Hz, local and global responses are almost identical.

Similar to purkinje cells in the mammalian cerebellum, superficial cells also receive massive numbers of inputs from parallel fibres as well as inhibitory cells driven by parallel fibres (Figure 2.3) [47]. These parallel fibres originate from granule cells in the cerebellar-like posterior eminentia granularis (EGp). Granule cells are driven by another population of pyramidal cells, called deep cells, that are not synapsed by parallel fibres and do not exhibit a cancellation response [38]. Deep cells faithfully transmit the AM

signal regardless of the stimulus' spatial extent and project to neurons in the nucleus praemientialis (nP), which provide input to the granule cells [36].

The circuit from deep cells to granule cells and back to superficial pyramidal cells via parallel fibres is called the feedback pathway. This feedback is known to produce the adaptive cancellation response seen in superficial pyramidal neurons since chemical or surgical blockade of the pathway reverts the superficial cell firing activity to local stimulation during a global stimulus [38]. Crucially, the feedback is only active when the stimulus has a large spatial extent; local stimulation (*i.e.* prey) alone is insufficient to activate this input.

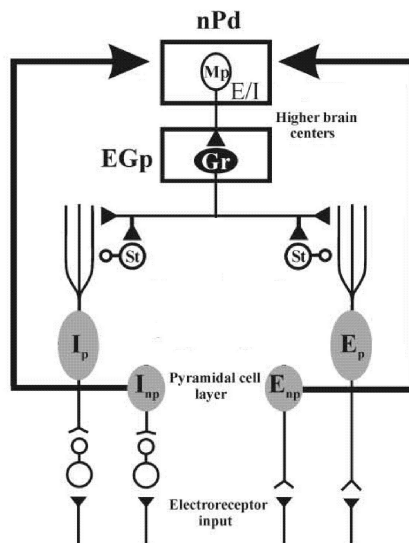


Figure 2.3: Schematic of ELL circuitry in *Apeteronotus leptorhynchus*. There are four types of pyramidal neurons in the ELL: plastic, p, and non-plastic, np, (also known as superficial and deep, respectively) as well as E and I. E neurons receive input from excitatory electroreceptors directly whereas I neurons receive electrosensory input through an inhibitory neuron and so decrease their firing rate to increases in electroreceptor input. If, and only if, the sensory input is global in extent will granule cells, Gr, become active. Granule cells project back onto superficial pyramidal cells, which putatively allow cancellation of predictable input signals. Granule cells also activate stellate cells, St, which provide disynaptic inhibition. Deep cells project to the nP and initiate the feedback pathway, but do not receive granular input themselves. Adapted from [38].

Unfortunately, due to their small size, the firing activity of the granule cells in electric fish is unknown. Nevertheless, *in vitro* and *in vivo* studies have shown that similar granule cells in mammals tend to burst to specific natural sensory input but are oth-

erwise silent [48, 49]. Furthermore, *in vitro* investigations by D’Angelo discovered that granule cells typically burst at the same phase of a periodic stimulus, regardless of the stimulus frequency [50]. This phenomenon is known as phase-locking, and granule cells in *Apteronotus leptorhynchus* likely phase-lock to the AM frequency since they are driven by deep cells whose firing rate is linearly modulated by the input stimulus.

However, even if granule cells burst synchronously, their input will not reach superficial cells at the same time because of the parallel fibre network architecture. Parallel fibres are thin, unmyelinated axons, making action potential propagation extraordinarily slow, and each parallel fibre traces out a unique distance from every granule cell to each superficial cell. Depending on the spatial location of each cell in its respective structure as well as the relative location of the structures themselves (*e.g.* same side, ipsilateral, or opposite side, contralateral, of the brain) [47], the parallel fibres that synapse onto any one superficial cell will have a variety of lengths. In addition, granule cells will not be active simultaneously because of phase differences in their input from different path lengths between deep cells and granule cells [36]. Consequently, there are a diversity of time delays in the feedback pathway (Figure 2.4).

Therefore, the cancellation problem for each superficial neuron at each AM frequency is how it can remove a feedforward input signal using a feedback pathway composed of bursting granule cell input at a multitude of fixed delays compared to the stimulus to be cancelled. The answer is thought to involve synaptic plasticity at the parallel fibre-superficial cell synapse, which was recently discovered in the centro-lateral segment (CLS) of the ELL [41]. Although single spike pairings had no effect, long-term depression was induced by pairing a 2-spike parallel fibre burst with a 2-spike superficial cell burst. Similarly, pairing 4-spike bursts also depressed the synapse, and by a much larger magnitude (Figure 2.5).

This plasticity could alter granule cell input at each delay to shape the feedback and produce a “negative image” of the electroreceptor input to cancel it. Further, continual operation of synaptic plasticity allows the network to adapt to novel global signal frequencies and strengths. Thus, the ELL putatively creates an adaptive filter to solve a simplified version of the cocktail party problem. Redundant signals are identified by their spatial extent, which activates the feedback pathway. The granule cells filter the stimulus and supply a delayed input framework to superficial pyramidal cells. Burst-induced synaptic plasticity then alters the weight of each granule to adaptively cancel the extraneous input.

To test if these components are sufficient to remove redundant stimuli and reproduce




Figure 2.4: An outline of the variety of anatomical delays in the feedback pathways. Deep cells (DP) in the left and right sides of the ELL project to the nP which in turn projects to the EGp. In the EGp, granule cells from either side propagate many parallel fibres down to the entire ELL. Blue arrows indicate input from the right ELL and red arrows indicate input from the left ELL. Solid and dashed thin lines represent parallel fibres from the left and right EGp, respectively. There are varying transmission delays due to the physical extent of the ELL as well as the different distances between a given granule cell and each SP cell it synapses upon. The points at which varying delays can occur are listed as ΔT 1 through 6.

in vivo cancellation, a model of the ELL architecture in the CLS and the firing behaviour of superficial E cells was constructed in this thesis. This model must be biologically realistic, and so was constrained to fit available experimental data.

2.2 Experimental Data

2.2.1 Neural Spiking Statistics

Since the only output of a neuron to other neurons are spike trains, this is considered the only important data to recreate in a model. However, the structure of spike trains can be very complex, and the effect of model inputs on the spike train is non-linear, making parameter fitting difficult. The critical statistics of spike trains that will be used to compare the data and the model in this thesis will be the average spike rate, the periodic modulation of the spike rate given a periodic stimulus, and the 2- and 4-spike burst rates.




Figure 2.5: *In vitro* single spike and burst-induced synaptic plasticity at the PF-SP cell synapse discovered by Harvey-Girard *et al.* [41]. The experimental method is as follows: the synaptic strength was measured initially by stimulating a parallel fibre to generate a post-synaptic potential and recording its maximum amplitude. Pre- and post-synaptic events (either single spikes, 2-spike bursts, or 4-spike bursts) were paired at different delays as numerated on the abscissa. After 100 pairings, the synaptic strength was remeasured and the difference as a percent of the original strength is plotted on the ordinate. Results are averaged over 5 recorded cells and the error bars are the standard deviation. **A** illustrates the plasticity induced by single spikes while **B** shows the plasticity induced by bursts. The dotted lines denotes no change in synaptic weight strength. Data from [41].

The output of a neuron is still considered to be encoded in the average firing rate of a neuron, as outlined in Chapter 1 [7]. Even with interest in the timing of single spikes, the mean rate is still an important metric. Furthermore, deviations from a neuron's spontaneous rate may indicate a decreased sensitivity to inputs because of the non-linear nature of spike production (Figure 1.3).

When the stimulus is periodic, it is convenient to create a post-stimulus time histogram or PSTH (Figure 2.6). This is a histogram of spike times that is synchronized to the frequency of the input and will show the average modulation of the firing rate over the period of the input from 0 to 2π [9]. For all PSTHs in this thesis, the period is discretized into 20 bins. Sometimes there is a difference between the phase of the external stimulus and the phase of the firing rate modulation of the neuron, which can be easily identified with a PSTH. This phase shift is due to the computational delay of a neuron receiving an input before firing, and perhaps also the transmission delay of the stimulus from sensory neurons to the recorded neuron.

Another statistic that will be used to compare data and the model is cancellation, which compares the strength of the firing rate modulation in two different paradigms. The strength of the modulation is obtained by fitting a sine wave to the PSTH and recording the sinusoid's amplitude (Figure 2.6). Cancellation during global stimulation

is defined as

$$\text{Cancellation} = \left(1 - \frac{\text{Amp}_{\text{global}}(f)}{\text{Amp}_{\text{local}}(f)} \right) \times 100\%$$

so that if the global modulation is identical to the local modulation, then cancellation is zero.




Figure 2.6: Example of a post-stimulus time histogram (PSTH). A neuron is stimulated with a sinusoidal input (A) in its receptive field, and its action potentials are observed using an electrode. After many stimulus cycles, the action potentials are binned according to the phase of the input stimulus when they occurred (B). The number of action potentials per bin is translated to a firing rate by dividing by the temporal length of the bin: $\Delta t = T/20$, where T is the period of the stimulus. The phase shift between the input stimulus and the neuron’s response is evident. The PSTH is often represented as a continuous function (C) for clarity. Also shown is a sinusoidal fit of the PSTH (solid grey line). The amplitude, Amp , of the fitted sinusoid is the modulation of the PSTH, which is used to calculate cancellation. The dashed grey line is the average value of the fitted sinusoid. The firing response illustrated here is from a superficial cell during 16 Hz global stimulation.

The grouping of spikes is also an important characteristic of a spike train. There is gathering evidence that bursts can operate as an alternative information event to single spikes that can aid in transmission reliability [51] or increase information channel capacity [52]. Bursts are a fast sequence of spikes that are considered a single composite event. It is often beneficial to define bursts in a system if post-synaptic neurons respond non-linearly to a brief sequence of spikes and the relative timing of spikes within each burst is unimportant.

In electric fish, bursts are known to selectively encode low frequency stimuli (0-10 Hz)

[45] such as prey to the ELL and higher order centres and to drive synaptic plasticity at the superficial cell-parallel fibre synapses [41]. Although there is no universal consensus to the definition of a burst, bursts in the model and the experimental data must be consistent with the bursts that induced plasticity experimentally. The 2-spike and 4-spike bursts that induced plasticity *in vitro* (Figure 2.5) were spike sequences with 10 ms intervals between spikes and either 2 or 4 spikes in a sequence. For generality, a 2-spike burst in this thesis was defined to be 2 spikes within 15 ms and a 4-spike burst was defined to be 4 spikes within 45 ms (Figure 2.7). However, decreasing the burst definition to a 12 ms or 10 ms inter-spike interval (*i.e.* time between spikes) had a negligible impact on model fitting and behaviour. In parallel to the mean firing rate, the 2-spike and 4-spike bursts rates are the average number of bursts of each kind that occur during stimulation. Note that a 2-spike burst is not identified within a 4-spike burst and a single spike can never be in multiple bursts.




Figure 2.7: Temporal organization of spikes. These two spike trains have the same average firing rate but are organized differently and may code different information. The burst definitions used in the model are identified in the second spike train: a 2-spike burst occurs whenever there are 2 spikes within 15 ms of each other; a 4-spike burst occurs when there are 4 spikes within 45 ms of each other.

2.2.2 Surgery and *In Vivo* Recordings

All experimental preparations and cellular recordings were performed by Dr. Gary Marsat, a collaborator in a weakly electric fish laboratory operated by Prof. Len Maler at the University of Ottawa. *Apteronotus leptorhynchus* was placed under general anaesthesia (tricaine methanesulfonate, Finquel MS222, Argent Chemical Laboratories, Redmond, WA) with oxygenated water pumped past its gills during surgery. A small incision was made at the top of the head and a portion of the skull was removed to expose the ELL region of the brain. Local anaesthetic was injected into the head opening and the fish was revived, paralyzed with curare, and placed into a holding cell in an experi-

mental tank with water between 25°C and 27°C and conductivity around 200 S (for more information, see [53, 54]). The wound was positioned above the water during all experiments.

In vivo recordings were accomplished using KAc-filled borosilicate intracellular electrodes. The anatomy of the ELL is well established, and medio-lateral and dorsal-ventral positions of the electrode were monitored to estimate which segment of the ELL (lateral, centro-lateral, centro-medial, medial) the electrode was in. Superficial pyramidal neurons were determined based on their dorsal-ventral positions, their firing rates, and their cancellation of low frequency global stimuli. All procedures were approved by the University of Ottawa Animal Care Committee.

Stimuli consisted of amplitude modulations of the fish’s EOD. Global stimulation was delivered through 2 large dipoles on opposite sides of the fish. Local stimulation was produced by first mapping the receptive field of a recorded pyramidal neuron and placing a dipole at its receptive field centre. To create amplitude modulations, the EOD of the fish is continuously sampled and the dipole emits a monotonic electric signal which has a frequency equal to the fish’s EOD plus the AM frequency to be investigated. Thus, to analyze a 20 Hz AM, the dipole’s signal is at a frequency of the fish’s EOD +20 Hz. The intensity of either stimulation paradigm created a 10-15% modulation of the fish’s EOD when measured near the surface of the cell’s receptive field. This intensity is within the linear coding region of electroreceptors (see Figure 1.3).

2.2.3 Local and global stimulation data

From experimental recordings, the mean firing rate, 2- and 4-spike burst rate, and PSTH of recorded superficial pyramidal neurons during local and global stimulation at AM frequencies of 0.5 Hz, 1 Hz, 2 Hz, 4 Hz, 8 Hz, 12 Hz, 16 Hz, 20 Hz, and 32 Hz are known (Figure 2.8 and 2.9). The amount of cancellation can be calculated from the local and global PSTHs and will also be compared between experiment and model. In addition, under spontaneous conditions (*i.e.* no AM stimulus presented), the average firing rate of the recorded neurons was 9.5 Hz.

Since superficial E pyramidal neurons in the ELL receive electrosensory input via electroreceptors, there is a response delay between the neuron’s PSTH and the input stimulus. As mimicking the transmission time was not a goal of the model, the experimental data was shifted to optimally overlap the stimulus during local stimulation at each AM frequency independently (Table 2.1). Identical phase shifts were employed to

Frequency (Hz)	Phase shift	Temporal shift (ms)
0.5	$-3\pi/10 \pm \pi/20$	-300 ± 50
1	$-3\pi/10 \pm \pi/20$	-150 ± 30
2	$-2\pi/10 \pm \pi/20$	-50 ± 10
4	$-2\pi/10 \pm \pi/20$	-25 ± 6
8	$0 \pm \pi/20$	0 ± 3
12	$0 \pm \pi/20$	0 ± 2
16	$2\pi/10 \pm \pi/20$	6 ± 2
20	$3\pi/10 \pm \pi/20$	8 ± 1
32	$5\pi/10 \pm \pi/20$	7.8 ± 0.8

Table 2.1: Phase shifts used to align the local PSTH data of superficial cells with the stimulus. The phase shifts can represent temporal delays of computation, and so the time delays are also listed. A negative phase shift implies the PSTH of the superficial cell was leading the stimulus (*i.e.* the peak of the neuron’s firing rate occurred before the peak of the stimulus). These same phase shifts were used to modify the PSTHs recorded during global stimulation.

adjust the data during global stimulation. Note that the large negative phase shifts at low frequencies is likely due to the adaptation of the electroreceptors, which would cause the firing rate to peak even as the stimulus was increasing, leading to the cell’s response leading the stimulus.

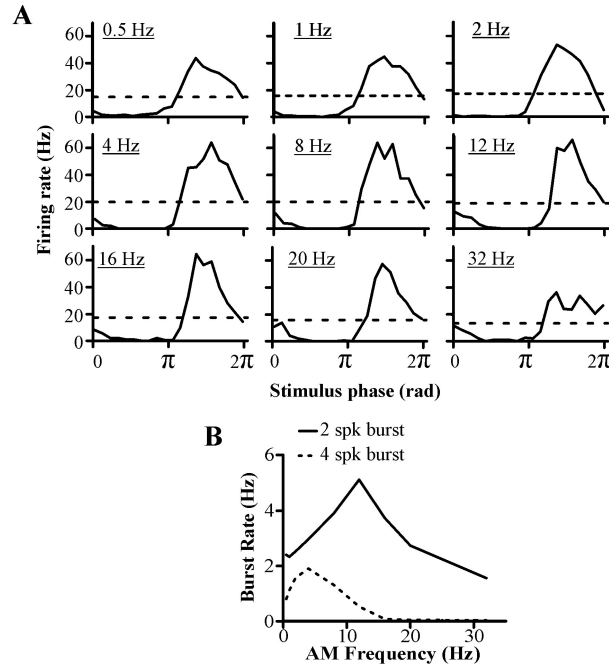


Figure 2.8: *In vivo* experimental data of the superficial cell response to local stimulation. (A) PSTHs with average firing rates (dashed lines), as well as (B) 2-spike and 4-spike mean burst rates have been recorded by Gary Marsat. The decrease in the maximum firing rate of PSTHs at frequencies below 4 Hz is indicative of electroreceptor adaptation. The poor response to 32 Hz suggests that this forcing frequency may be too high for the pyramidal neuron to track optimally. Note that the 2-spike burst rate is maximum at 12 Hz due to the selection of 4-spike bursts in the spike train, which, because of the independent spike requirement of individual bursts, decreases the 2-spike burst rate when 4-spike bursts are numerous (*i.e.* at low frequencies). If the independent spike requirement is removed, the model and experimental burst rates are inversely proportional to AM frequency (see Section 6.1.1).

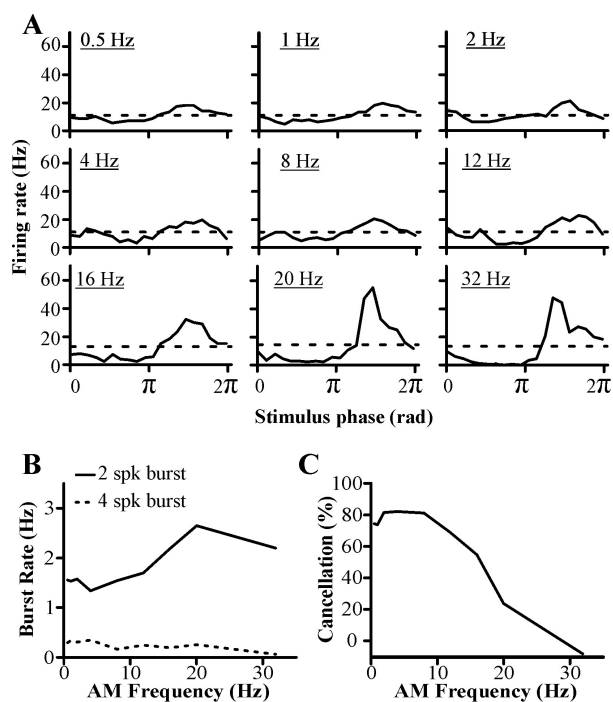


Figure 2.9: *In vivo* experimental data of the superficial cell response to global stimulation. (A) PSTHs with average firing rates (dashed lines), as well as (B) mean 2-spike and 4-spike burst rates and (C) mean cancellation data have been recorded by Dr. Gary Marsat. Experimental cancellation is effective for AM frequencies less than 10 Hz but deteriorates sharply for higher frequencies. The cancellation is negative at 32 Hz because the firing rate modulation was higher during global stimulation than during local stimulation (compare the PSTH at 32 Hz in this figure with Figure 2.8).

Chapter 3

Modelling Methods

In this chapter, the basic methodology of neural modelling in this thesis is outlined. Each element of the biological cancellation network described in the previous chapter is transformed into a mathematical framework to assemble the neural model. Details of the numerical methods to simulate the system are also provided.

3.1 Neural Modeling

Neurons actively separate charges across their cellular membrane to create a potential difference between the extracellular fluid and the cell interior. The neuron then acts as a capacitor, and ion channels in the membrane create currents that change the voltage. Due to different densities of ion channels and varying morphological features, the voltage difference across the membrane need not be the same in all regions of the neuron. However, if the dendrites of the neuron are short and have low resistance, the cell can be approximated as electronically compact and modelled as a single compartment with the following governing equation for a capacitor [7]:

$$C \frac{dV}{dt} = \sum_i I_i \quad (3.1)$$

where C is the membrane capacitance, V is the voltage difference across the membrane, and I_i is the electric current through the membrane by channel i . Most of the current into the neuron is supplied by ion channels that are selectively permeable to a single ion. The amount of current that flows through the cellular membrane depends on the voltage difference of the membrane, the concentration difference of all the permeable ions, and

the conductance of the ion channels.

Typically, the current of a channel that is permeable to only one ion is modelled as $g(E_{rev} - V)$, where g is the conductance of the channel and E_{rev} is the reversal potential of the ion, which depends on the ion concentration and the temperature. If the membrane potential equals the reversal potential, then the electrical and diffusion potentials for that ion are equal and no net current crosses the membrane (for more information, see [8]). The sum of the reversal potentials of all active ion channels, weighted by their conductances, defines the resting potential of the cell. The conductance, g , may change as a function of voltage, firing history, or other ion concentrations.

Once the cell's membrane potential reaches a certain voltage threshold, potassium and sodium ion channels with non-linear conductance relationships to the membrane voltage create the stereotyped action potential response of the neuron. The ion channel dynamics that allow the neuron to spike are well known but complex and computationally expensive [7]. Since a neuron always fires after its membrane potential is sufficiently depolarized, it is convenient to replace the firing dynamics with a voltage threshold. Whenever the membrane potential exceeds this threshold, V_{thresh} , the membrane potential is reset to its resting potential, V_r , and a spike is said to have occurred at this time. Thus, only subthreshold currents need to be reproduced, which can be simulated with different levels of detail. In the simplest model, all ion channel conductances are assumed to be constant and can be summed together to create one leak term. In this way, the cell can be modelled as a capacitor and resistor in series, which has the following governing equation:

$$C \frac{dV}{dt} = g_m(E_m - V), \quad (3.2)$$

where E_m represents the combined reversal potentials of all ion channels in the membrane, and g_m represents the total conductance. Dividing equation 3.2 by this conductance gives

$$\tau_m \frac{dV}{dt} = E_m - V. \quad (3.3)$$

In other words, the membrane potential changes to minimize the net current through the membrane, and any disturbances of the voltage around the reversal potential E_m will decay exponentially back to E_m with time constant τ_m . This approach to simulating neurons is called the leaky integrate-and-fire (LIF) model, and it is one of the simplest models of a neuron. Nevertheless, the non-linear nature of neural computation is still present in the voltage threshold. If the refractory period of a cell, τ_r , is added (a natural

firing dead time after a neuron has fired that prevents it from firing again), then the voltage is artificially held at its rest value until the refractory period is over. Non-dimensionalizing the reversal potential in terms of the voltage difference, $V_{thresh} - V_r$, and the time in terms of τ_m , the canonical form of the LIF model is obtained:

$$\frac{dV}{dt} = I - V \quad (3.4)$$

where V is the non-dimensionalized neuronal membrane potential and I is the non-dimensionalized reversal potential, commonly called the bias current. Often, V_{thresh} is set to one and V_r is set to zero. Since this equation is deterministic, the cell will fire at an exact frequency if the current is constant and greater than threshold.

A realistic neuron in the brain, however, never has a constant input; it is constantly bombarded with synaptic inputs from a variety of other neurons, as well as inputs from glial cells, movements of ions in the cell structure, and random opening and closing of ion channels in its membrane. Since only the firing rate over a timescale of seconds is usually important, it would be fruitless to model all of these microscopic mechanisms individually, and so they are often collectively modelled using Gaussian white noise. This assumes that all of these current sources are at least weakly independent, which is a fair approximation. Thus, Gaussian white noise is included in LIF models to simulate the random firing patterns of realistic neurons. Interesting effects can be obtained in a non-linear system with a subthreshold input and the addition of noise, such as stochastic resonance [55] or coherence resonance [56], where noise is actually beneficial to signal detection. However, these behaviours are beyond the scope of this work. If Gaussian white noise is added to the model, the non-dimensionalized equation becomes

$$\frac{dV}{dt} = I - V + \sigma\xi(t), \quad (3.5)$$

where $\xi(t)$ is uncorrelated Gaussian white noise ($E\{\xi(t)\xi(t')\} = \delta(t - t')$; $E\{\}$ is the expectation operator taken over different realizations of the noise process) and σ is a parameter that controls the noise intensity. Stochastic LIF models have been used previously to simulate neurons in the ELL of electric fish [57], as they achieve an appropriate balance between biological realism and computational simplicity.

3.2 Modelling the ELL

The method chosen to investigate the ELL neural network in this thesis is to create a minimal model. A minimal model is useful because it is computationally efficient, it is analytically tractable, and its results are straightforward to interpret. This approach attempts to distill each component of the original system into its basic form, and removes superfluous influences that are not critical to the system’s performance. In this section, each element of the biological network is simplified to a mathematical representation and added to the stochastic LIF model. The governing equations of the model are provided at each step to illustrate the assembly. The parameters identified in this section are fitted to the experimental data in the following chapter.

3.2.1 Superficial pyramidal neurons

The firing activity of superficial cells is fully simulated in the model because they exhibit the redundant input cancellation to be investigated in this work. For simplicity, the leaky integrate-and-fire (LIF) framework with Gaussian noise was used to model their firing activity:

$$\tau_m \frac{dV}{dt} = I - V + \sigma \xi(t). \quad (3.6)$$

As will be described below, the source of the input current and noise to the pyramidal cell in this model are electroreceptors, granule cells, and inhibitory cells in the ELL. Independently simulating each neuron in the network would be computationally uneconomical, so the firing behaviours of neurons other than the superficial cell will be approximated based on previous experimental studies. Their effect on superficial cells can then be estimated.

3.2.2 Electroreceptors

Electroreceptor afferents densely cover the skin of weakly electric fish and are the primary sensors of the EOD. These neurons phase-lock to its frequency, but, because the EOD is at a much higher frequency (600-1000Hz) than the average firing rate of the neurons ($191 \pm 81\text{Hz}$), electroreceptors exhibit “skipping:” firing always at the same phase of the EOD cycle but randomly missing some EOD periods [10, 58]. Nevertheless, moderate intensity amplitude modulations of the EOD, which carry information about the environment, are almost always linearly encoded by these afferents (Figure 1.3) [10]. The exception is

at very low AM frequencies, where the afferents adapt to a slowly modulating input to minimize changes in their firing rate [59].

Electroreceptors project to pyramidal neurons in the ELL to provide information about the external surroundings. Since each spike of an afferent is a unique, excitatory input to a superficial pyramidal cell, electroreceptor input could be modelled as a modified Poisson process during stimulation. However, many afferents converge onto one pyramidal cell and their baseline firing rate is very high (~ 200 Hz [10]), so the diffusion approximation can be used to simplify the feedforward input as a bias current and a Gaussian noise source with standard deviation σ [60]. Because the afferents are linear encoders, an amplitude modulation at a single frequency can be modelled as a sinusoidal variation of the electroreceptor firing rate and, hence, input. Afferent input is also strictly excitatory, so this input is rectified, and, since this input is passed through a synapse and a dendritic tree, the noise is low-pass filtered:

$$\tau_m \frac{dV}{dt} = -V + [I + \sigma \xi_L(t) + \kappa(f) \sin(2\pi ft)]. \quad (3.7)$$

In this model, $[\dots]$ denotes rectification and $\xi_L(t)$ denotes low-pass filtered Gaussian white noise. The contrast, $\kappa(f)$, is defined as the percent modulation of the electroreceptor's firing rate given an AM at a specific frequency. Experimentally, all AMs are presented to the fish at the same relative intensity, but, due to the adaptation of electroreceptors, the modulation of their firing rate will change at low AM frequencies. Hence, $\kappa(f)$ was identical for all frequencies except those less than 4 Hz, where afferent adaptation becomes significant. These low frequency contrasts were individually fitted.

3.2.3 Depolarizing after-potential

The burst rate of superficial neurons in the ELL is important not only for signalling to downstream neurons but also for learning at the parallel fibre-superficial cell synapse. Therefore, it is crucial to accurately mimic the bursting behaviour of these cells in the model. The superficial neuron's propensity to burst is due to a phenomenon called the depolarizing after-potential or DAP. This effect has been studied extensively [61] and has been simplified to a LIF-DAP framework [57], which was also used in this model:

$$\tau_m \frac{dV}{dt} = -V + [I + \sigma \xi_L(t) + \kappa(f) \sin(2\pi ft)] + DAP(t) \quad (3.8)$$

Briefly, after the neuron generates an action potential that propagates down the axon,

the depolarizing current in the soma will also back-propagate into the dendrites of the cell. Superficial neurons are unique in that they have voltage-gated ion channels in their dendrites as well as the soma and axon. This causes the dendritic tree to strongly depolarize, creating a post-synaptic potential. The relevant DAP input to the model will be the difference between this dendritic PSP and the somatic PSP that still exists from action potential generation. Since PSPs are often approximated by alpha functions, te^{-bt} (see Figure 1.2), the DAP is modelled as a difference in alpha functions.

Due to the slow diffusion of ions in the cell body, this current will reach the soma a short time after the neuron fires. Like all active ion channels, the dendritic channels have a refractory period, and if another action potential were to arise before the dendrites had recovered from a previous spike, then the dendrites would be unable to depolarize and the DAP would not occur.

The dendritic refractory period, r_d , is modelled as a dynamic variable that changes according to a secondary variable, b , which also controls the width of the dendritic alpha function. All spikes generated by the neuron are recorded, the most recent of which was at time t_n , and b updates whenever the neuron fires a spike. t_n^+ refers to the time just after the most recent spike was fired. This model was created by Doiron and Noonan [57] and was used in this thesis with identical parameters except A and γ have been increased four fold (see Section 4.1). The equations governing the DAP are

$$DAP(t) = \begin{cases} 0 & \text{if } t - t_n < r_s \\ \alpha s(t - t_n, \beta b(t_n^+)) - \alpha s(t - t_n, \gamma) & \text{if } t - t_n > r_s \text{ and } t_n - t_{n-1} > r_d^n \\ 0 & \text{if } t - t_n > r_s \text{ and } t_n - t_{n-1} < r_d^n \end{cases} \quad (3.9)$$

$$\frac{db}{dt} = -b/\tau + (A + Bb^2) \sum_n \delta(t - t_n) \quad (3.10)$$

$$s(t, a) = \frac{te^{-\frac{t}{a}}}{a} \quad (3.11)$$

$$r_d^n = D + Eb(t_n^+) \quad (3.12)$$

The parameters used in the above equations are listed in Table 3.1.

Parameter	Value	Parameter	Value	Parameter	Value
A	0.6*	B	2	α	20
γ	0.2*	β	0.35	τ	τ_m
r_s	$0.1\tau_m$	D	0.1	E	3.5

Table 3.1: Parameters used in the DAP model. * indicate changes from [57].

3.2.4 Granule cells

Granule cells receive input from deep cells in the ELL and are therefore driven by a periodic signal phase-locked to the predictable global AM. Due to their small size, electrophysiological recordings of granule cells in apteronotids are difficult, but observations of granule cells in other animals suggest that they would burst to periodic sensory input [49, 48, 50]. Thus, a single granule cell was assumed to burst once during each stimulus cycle and be silent at all other times.

There are a multitude of pathways from deep cells to granule cells and, via parallel fibres, to superficial cells (Figure 2.4) [47]. For example, a deep pyramidal neuron may project to the contralateral or ipsilateral nP, which may project to granule cells in the contralateral or ipsilateral EGp. Granule cells can then feed back onto the contralateral or ipsilateral ELL, and the specific location of the granule cells and pyramidal neurons will also change the length of the projections. Since the conduction velocity of action potentials in parallel fibres is extraordinarily slow (0.1 m/s [40]), different axonal lengths will translate into different temporal delays of granule cell burst input on the order of milliseconds. In addition, there are both E and I deep cells [42], and the firing modulation of these neurons will be in anti-phase. Assuming granule cells are driven by only one type, this bestows an inherent π phase shift to the feedback input.

If there are heterogeneities or noise in the granule cell population, then even synchronous deep cell input will produce a distribution of burst times from different granule cells. Hence, it was assumed that the parallel fibre feedback to superficial cells occurs equally at all possible phases of the AM stimulus, regardless of the AM’s frequency. The sum of the parallel fibre activity at any time during global stimulation is taken to be a constant input, Λ .

Since granule cells are assumed to phase-lock to the predictable stimulus, the total feedback is periodic with the same frequency. In the model, the feedback cycle is discretized into 2.5 ms segments, which means that different AM frequencies will have a different number of feedback segments (*e.g.* 100 feedback segments for a 4 Hz stimulus

and 50 segments for an 8 Hz stimulus). Each segment represents a collection of parallel fibres that stimulate the superficial cell at the same phase of the feedforward input each period. The choice of discretizing the feedback in time and not in phase is based on the assumption that parallel fibre activity is always a granule cell burst, and that a burst has a defined temporal length. Each segment, labelled ϕ , becomes active at time t_ϕ , has an inherent strength Λ , a unique synaptic weight, w_ϕ , and then inactivates at $t_\phi + 2.5$ ms. The total feedback input is a discontinuous but periodic signal composed of Λ multiplied by the synaptic strength, w_ϕ , for each segment as time moves from segment to segment during a period:

$$\tau_m \frac{dV}{dt} = -V + [I + \sigma \xi_L(t) + \kappa(f) \sin(2\pi ft)] + DAP(t) + \Lambda w_\phi, \quad (3.13)$$

w_ϕ is governed by the burst-induced depression rules found *in vitro* as well as a potentiating rule. For simplicity, each segment was taken to be both a 2-spike burst and a 4-spike burst, so that the appropriate learning rule was immediately employed if the pyramidal cell fired either of the two bursts (Figure 3.1).

A stable phase relationship for each segment and, hence, each weight, is necessary for cancellation. However, it is created from fixed temporal delays, and changing the AM frequency would require the weight distribution to relearn the appropriate values. Furthermore, for a fixed set of delays, parallel fibre activity would overlap at higher frequencies, and may not sufficiently cover the period of lower frequencies. For simplicity, it was assumed that each frequency is cancelled independently and has its own unique synaptic weights for its collection of segments. This would imply physiologically that each parallel fibre is only active for a specific frequency (or a small frequency range) and that there are sufficient parallel fibres to cover all phases for all frequencies of interest. Also, given that each parallel fibre is active for a defined period of time, this implies that there are more parallel fibres active at lower frequencies than higher frequencies. Fortunately, due to the massive number of parallel fibres feeding back into the ELL, this is plausible.

3.2.5 Pyramidal cell-parallel fibre synaptic plasticity

The strength of the parallel fibre-pyramidal cell synapse is governed by two burst-induced depression rules found by Harvey-Girard *et al. in vitro* (Figure 2.5) [41]. Pairing 4-spike bursts or 2-spike bursts with inter-spike intervals of 10 ms caused significant synaptic depression whereas single spike burst pairings had no effect. Due to the shape of the

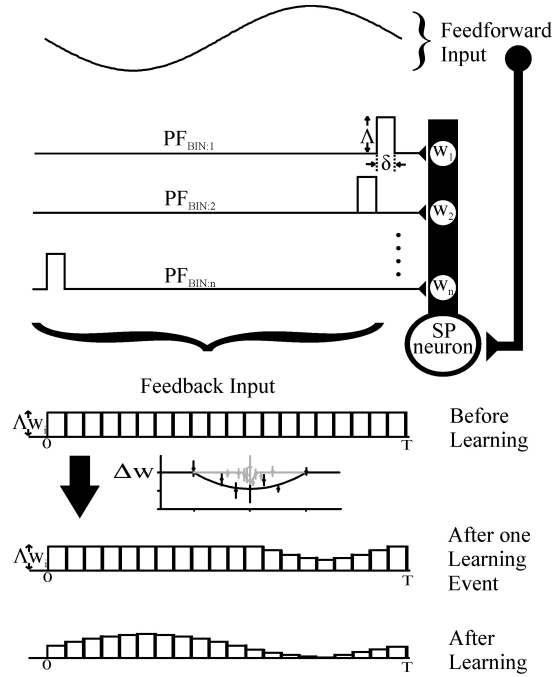


Figure 3.1: Schematic of the modelled feedback pathway. The period is divided into 2.5 ms segments ($\delta = 2.5$ ms), and no two segments are active at the same time. Each active segment is defined as both a 2-spike and a 4-spike parallel fibre burst. Whenever the superficial neuron bursts, the appropriate learning rule is immediately invoked (depending on the type of superficial burst) and changes the weights of nearby segments. After successful learning, the weight distribution should resemble a sinusoid in anti-phase to the feedforward input.

depression as function of the time delay between burst pairs, each rule is an anti-Hebbian burst-timing dependent plasticity rule.

The burst-induced plasticity was fitted by a continuous curve quadratic with the pairing delay and implemented by monitoring the firing activity of the model pyramidal cell. Whenever the neuron emits 2 spikes with an inter-spike interval of 15 ms or less, a 2-spike burst is recorded. If the cell emits 4 spikes within 45 ms or less, then a 4-spike burst is recorded. The time of the first spike in each burst is taken to be the time the burst occurred, in agreement with Harvey-Girard *et al.*'s definition [41], and the spikes in each burst must be independent (*i.e.* there cannot be a 2-spike burst in a 4-spike burst, or a 4-spike and a 2-spike burst in 5 spikes). Since the parallel fibre feedback is segmented and each segment is defined as a burst, there are parallel fibre bursts at each time t_ϕ in the model, spaced 2.5 ms apart. When the superficial cell bursts under global stimulation at time t_B , the following burst learning rule is immediately invoked for all

parallel fibre segments:

$$\text{if } |t_\phi - t_B| < L_{w_{2,4}} \rightarrow w_\phi = w_\phi + w_\phi \eta_{2,4} \left[\left(\frac{t_\phi - t_B}{L_{w_{2,4}}} \right)^2 - 1 \right] \quad (3.14)$$

where η_2 and L_{w_2} are used if the superficial cell's burst is a 2-spike burst, and η_4 and L_{w_4} are used if it is a 4-spike burst. This rule is applied to all weights whose segments began at a time t_ϕ as long as $|t_\phi - t_B| < L_w$. Beyond this range, the weights are unchanged. This is a multiplicative rule, since the amount of depression caused by a burst depends on the value of the weight at that time. This method was chosen, as opposed to additive learning, because it is more biologically realistic [25, 11].

3.2.6 Potentiating rule

It is known that redundant input cancellation in apteronotids can be unlearned once the predictable stimulus is removed [38]. However, no mechanism for long term potentiation at this synapse has been found as of yet. Nevertheless, a purely depressing rule as found *in vitro* would trivially decrease all weights to zero given sufficient time. Hence, a simple potentiating rule with a large time constant is included where all weights were increased at each time step according to

$$\tau_w \frac{dw_\phi}{dt} = w_{max} - w_\phi. \quad (3.15)$$

This rule is applied regardless of the activity of each segment or the superficial cell and is independent of stimulus frequency. This is similar to synaptic normalization constraints found in many theoretical neural network models [7], which guarantees competition among synapses. The potentiation rule here is less rigid, allowing the weights to strengthen or weaken independently. It is also more biologically plausible, as a large time constant will cause the weights to change very slowly, making this effect difficult to detect experimentally.

3.2.7 Disynaptic inhibition

Parallel fibre feedback from granule cells also activate stellate inhibitory cells that synapse on superficial pyramidal neurons (Figure 2.3), an effect called disynaptic inhibition [62]. Disynaptic inhibition is included in the model because it maintains the firing rate of the superficial neurons near their spontaneous rate during global stimulation. As in-

hibitory cells are stimulated from numerous parallel fibres, each with their own phase delay compared to the stimulus, the total input in the inhibitory cell will have a negligible frequency or phase relationship. Consequently, the effect of the disynaptic inhibitory cell during global stimulation can be modelled as the constant activation of a single type of ion channel: $g(E_{inh} - V)$. Furthermore, the reversal potential of the inhibitory synapse, E_{inh} is near the resting potential of the cell [63], which is defined to be zero in the LIF model, and there is no evidence that the parallel fibre-inhibitory cell synapse is plastic. Thus, the inhibition was modelled as a constant shunt, $-gV$, proportional to the activity of the feedback pathway:

$$\tau_m \frac{dV}{dt} = -V + [I + \sigma \xi_L(t) + \kappa(f) \sin(2\pi ft)] + DAP(t) + \Lambda (w_\phi - gV) \quad (3.16)$$

3.2.8 Summary

The model is summarized in Figure 3.2.

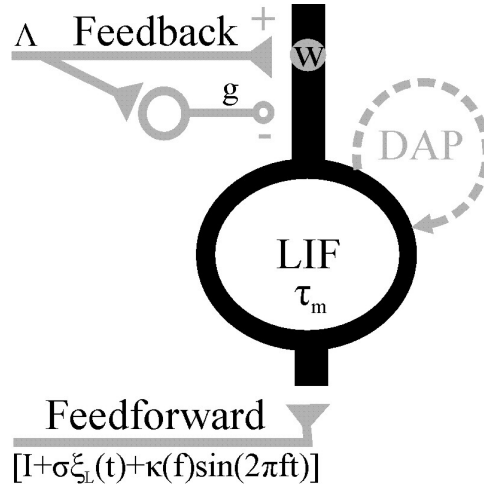


Figure 3.2: Schematic of model operation. A superficial cell is modelled with the leaky integrate-and-fire formalism. It receives feedforward input via electroreceptors, which is approximated as a rectified bias current with a sinusoidal variation and low-pass filtered Gaussian noise. The superficial cell's propensity to burst is included by installing a dendritic after-potential (DAP). The neuron receives feedback via parallel fibers that are assumed to burst and cover all phases of every frequency that the system can cancel. The strength of the parallel fibre-superficial cell synapse is adjustable and modelled as a weight. The feedback input is modified online by two burst-induced depression rules as well as by a recovery rule that govern the value of the synaptic weight at each phase of the stimulus.

3.3 Computer Simulations

All data analysis was performed within MATLAB although the model simulations were coded in C for speed. A series of random normally distributed numbers of zero mean and unit variance were generated offline for the noise, $\xi(t)$, using proprietary MATLAB programs. Low-pass filtering of the noise was accomplished using an acausal fourth order Butterworth filter, also implemented via a built-in MATLAB function. During the online simulation of the stochastic model, $\xi_L(t)$ is replaced by an element of this series at each time-step, and the elements are used sequentially. Note that the variance of the coloured noise is no longer unity after generation but f_{cut}/f_{Nyq} , where f_{Nyq} is the Nyquist frequency. For this reason, the low-pass filtered noise is divided by the square root of this ratio to obtain unit variance. The Nyquist frequency is the highest frequency that can be sampled without aliasing and is equal to half the sampling frequency. If the model is simulated in time-steps of Δ_s , then $f_{Nyq} = \frac{1}{2\Delta_s}$.

Since the AM frequencies are slow compared to the typical membrane time constants of neurons, the model neuron was assumed to behave ergodically for the AM input. Therefore, ensemble averages were replaced with time averages to calculate the mean firing rate as well as other statistics for each simulation. Unless otherwise stated, error bars on all model data were generated by simulating 1750 seconds of superficial cell stimulation, repeating the simulation 10 times, and calculating the standard deviation between the iterations. During local stimulation, the statistics of the model are recorded from the first 1750 seconds of simulation. During global stimulation, however, the model data should only be recorded when the weights have attained their equilibrium values.

3.3.1 Simulating the noise

Numerically simulating the canonical stochastic leaky integrate-and-fire model requires the Euler-Maruyama method [64]. In other words, to simulate $\dot{V} = -V + I + \sigma\xi(t)$ with a time-step Δ_s , then

$$V(t + \Delta_s) = V(t) + \Delta_s \left(-V(t) + I \right) + \sqrt{\Delta_s} \sigma \xi(t), \quad (3.17)$$

where the $\sqrt{\Delta_s}$ term multiplying $\xi(t)$ arises by the fact that the integral of Gaussian white noise is a Wiener process. However, the feedforward input must be rectified since electroreceptors are purely excitatory. Unfortunately, the Euler-Maruyama method breaks down for rectified Gaussian white noise, and the time step begins to affect the firing

rate (Figure 3.3). Consequently, Gaussian white noise in the model was replaced with low-pass filtered Gaussian noise for all computations in the thesis since its firing rate converged as the time step decreased (Figure 3.3C). In addition, low-pass filtered noise is more realistic since dendrites would smooth the current input from electroreceptor synapses to the pyramidal cell soma [65]. A time-step of 0.01 in units of τ_m was chosen for all future simulations because it is the largest Δ_s where the firing rate does not change if Δ_s is decreased further.




Figure 3.3: Effect of changing the integration time-step with different methods to model Gaussian noise when the noise is rectified. $\dot{V} = -V + [I + \sigma\xi(t)]$ is numerically approximated with (A) the Euler-Maruyama method, $V(t + \Delta_s) = V(t) + [\Delta_s(-V(t) + I) + \sqrt{\Delta_s}\sigma\xi(t)]$, or (B) the forward Euler method $V(t + \Delta_s) = V(t) + [\Delta_s(-V(t) + I) + \Delta_s\sigma\xi(t)]$. In C, $\dot{V} = -V + [I + \sigma\xi_L(t)]$, where $\xi_L(t)$ is low-pass filtered ($f_{cut} = 500$ Hz) noise, was simulated with the forward Euler method. σ is 0.1305, 1.305, and 0.759 in A, B, and C, respectively, so that, in all cases, the firing rate was 9.5 Hz when $\Delta_s = 0.01$. The time-step is then varied while all other parameters are fixed to investigate the convergence of each numerical approach. 10 simulations were run at each timestep for 175 s to generate the errorbars. $\tau_m = 7$ ms, $I = 0.58$.

3.3.2 Simulating the weights

During global stimulation, the synaptic weights of the model will evolve over time and break the ergodicity assumption of the model until transients have decayed and the weights are at their equilibrium values (assuming such values exist and are stable). To investigate whether the equilibrium weight values exist and, if so, how long it takes to converge, a weight at the peak and at the trough of the input as well as the average weight value of the entire distribution were tracked during the iteration of the model for 3 different initial conditions during global stimulation for various AM frequencies (Figure 3.4). By inspection, weights active at the same phase converge to the same final value regardless of their initial conditions, implying that a steady state of the system

does exist. Since the final distribution is the “negative image” of a sinusoid (see also Figure 6.20), the final weight values differ depending on the phase of segment the weight is associated.




Figure 3.4: Weight evolution during global stimulation. The parameters used are from optimally fitting the model and global stimulation data (the global η model), as described in Chapter 4, but with $\kappa(f)$, η_2 and η_4 constant across frequencies. The parameters are $\tau_m = 7\text{ms}$, $I = 0.58$, $\kappa(f) = \kappa = 0.39$, $\sigma = 0.759$, $f_{cut} = 500\text{ Hz}$, $\Lambda = 1$, $g = 1.44$, $\eta_4 = 0.0036$, $\eta_2 = 0.0018$, $L_{w_4} = 100\text{ ms}$, $L_{w_2} = 10\text{ ms}$, $\tau_w = 980\text{ s}$, and $w_{max} = 1.5$. Three different initial conditions (all weights initialized at 0.5, 1, or 1.5) were each run 5 times and the value of a weight at the peak of the input ($\phi = 3\pi/4$, green lines), the trough of the input ($\phi = \pi/4$, red lines), and the average weight value (blue lines) were tracked for 7000 s. Note that the input to the model is always an inverted sine wave (see Figure 2.8). Vertical dashed lines represent the epoch when statistics on the firing and burst rates are recorded during global stimulation.

Increasing the AM frequency tends to reduce the equilibrium weight values of the peak, trough and weight average, which will change the time for the weights to converge to their final values (the convergence time). Moreover, the convergence time depends on the phase of the weight and the initial weight value. Nevertheless, transients in the weight dynamics appear to have decayed for all frequencies and phases by 3500 s, which, by extension, means that the system is now ergodic. Given that there is no statistical difference in the weight distribution for periods beyond 3500 s, model data for global stimulation was always collected over the epoch between 3500 s and 5250 s (vertical

dotted lines in Figures 3.4). Since the convergence time is shorter when weights begin higher than their equilibrium values, the weights are always initialized at w_{max} to aid convergence.

Chapter 4

Parameter Fitting

The following chapter describes how the values of the local parameters, I , σ , f_{cut} , τ_m , and $\kappa(f)$, as well as the global parameters, L_{w_2} , L_{w_4} , η_2 , η_4 , g , τ_w and w_{max} , were obtained by fitting the model's behaviour to the available experimental data. Since local parameters are also used during global stimulation, local parameters were identified first, but could not by themselves produce a close fit to the experimental data. The size of the depolarizing after-potential (DAP) was then included in parameter optimization and a close match was obtained.

The global stimulation parameters were fitted to optimize three different learning paradigms: using only the 2-spike burst rule, only the 4-spike burst rule, or both. However, none of the three paradigms could effectively match the experimental data, especially at high frequencies when the model cancelled better than the *in vivo* neuron. Therefore, in agreement with biological observations in other systems, the feedback strength was reduced at high frequencies, either by altering Λ or η in the model, which produced a favourable match. For each parameter set to be optimized, the biological constraints on the parameters are identified, the effect of the parameters on the model is illustrated, and the final optimal fit is shown.

4.1 Local Stimulation

Under local stimulation, the governing equation for the model superficial cell is

$$\tau_m \frac{dV}{dt} = -V + [I + \sigma \xi_L(t) + \kappa(f) \sin(2\pi ft)] + DAP(t)$$

and the goal of parameter optimization is to closely match the model's behaviour

with the local experimental data (Figure 2.8). However, this task is complicated because the spike rate data has a non-linear dependence on the parameters due to the threshold behaviour of the leaky integrate-and-fire model. Furthermore, there is a non-linear relationship between 2-spike and 4-spike burst rates and firing rate (see Section 6.1.1 and 6.2.3).

Fortunately, biological criteria limit the phase space that needs to be explored. The parameters that can be fitted are I , the average, or spontaneous, electroreceptor input current, σ , the standard deviation of the electroreceptor input, f_{cut} , the cutoff frequency of the electroreceptor's noise, τ_m , the neuron's membrane time constant, and $\kappa(f)$, the amplitude of the electroreceptor input modulation. Physiological studies have shown that superficial pyramidal cells typically have τ_m on the order of a few milliseconds [43] and so τ_m was constrained to be between 1 ms and 20 ms. As stated in Section 3.2.2, the electroreceptor input should behave as a Poisson process but the diffusion approximation can be used to model the feedforward input as Gaussian noise with its variance, σ^2 , equal to its instantaneous input, $I + \kappa(f) \sin(2\pi ft)$.

Interestingly, experimental observations have shown that under local stimulation the variance of the electroreceptor firing rate does not vary with the stimulus input unless the electroreceptors are driven at very high AM intensities where saturation effects become significant [10]. Since all of the experimental stimulations to be modeled in this thesis were performed at low intensities, the noise variance was taken to be independent of the firing rate modulation, $\kappa(f)$, and equal to the mean input during spontaneous stimulation ($\sigma = \sqrt{I}$) for the purposes of parameter fitting. $\kappa(f)$ is itself constrained to be less than I because the maximum instantaneous firing rate of an electroreceptor under local stimulation is never more than double its spontaneous rate [10], so the maximum instantaneous input should be less than double the average input (*i.e.* $I + \kappa(f) \sin(2\pi ft) < 2I$).

The cutoff frequency of the noise does not impact the firing rate as significantly as the other parameters, especially at higher cutoff frequencies (see Section 6.1.1). Furthermore, low cutoff frequencies required $\kappa(f)$ to be greater than I to fit the locally stimulated experimental firing rates, which, as explained above, is unrealistic. Since varying the cutoff frequency from 200 Hz to 1000 Hz had a minimal effect on the firing and burst rates, f_{cut} was fixed at 500 Hz, and, it was not further altered.

To optimally fit the parameters I , τ_m , and σ , the following method was employed: (I , τ_m) pairs were identified that match the experimentally observed spontaneous (*i.e.* $\kappa(f) = 0$) firing rate of 9.5 Hz. For each pair, a 4 Hz sinusoidal input is added and the contrast of that input, $\kappa(4)$, is increased until the model firing rate matches the

experimental firing rate (20.2 Hz). This will be the local modulation strength for all frequencies 4 Hz and greater. Due to the low-frequency adaptation of electroreceptors, $\kappa(f)$ was individually fitted for each frequency less than 4 Hz [59]. PSTHs and 2- and 4-spike burst rates for different AM frequencies were calculated for each parameter set and compared to the experimental data for quality of fit (Figure 4.1).

The optimal parameters were then extrapolated ($\tau_m = 5$ ms in Figure 4.1), and the process was repeated within a narrower parameter range to maximize optimization. Although the variance of the noise, σ^2 , was fixed to be proportional to I , allowing σ to be a free parameter was investigated and found to have no significant effect on the fits (data not shown).

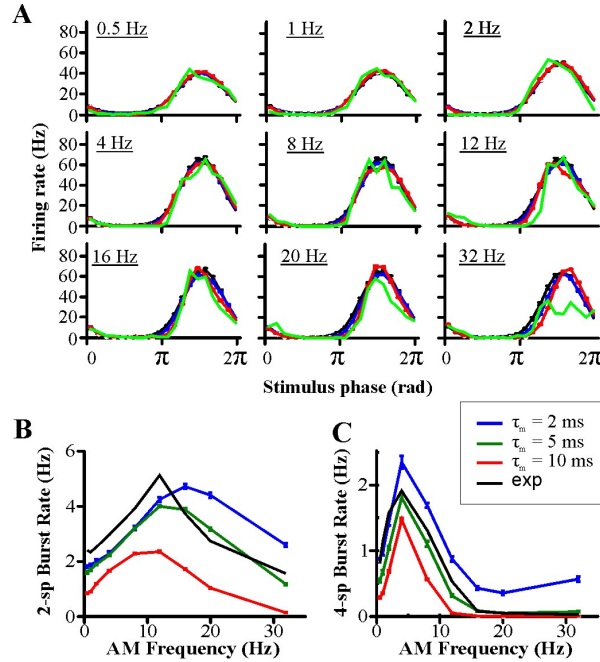


Figure 4.1: Parameter fitting to mimic experimental data under local stimulation. Model (A) PSTHs and (B) 2- and (C) 4-spike burst rates for different stimulus frequencies are compared to experimental data as the membrane time constant, τ_m , and the average electroreceptor input, I , of the model is varied. Black lines indicate experimental data. Coloured lines indicate model data. Blue lines have $\tau_m = 2$ ms, $I = 0.337$, and $\kappa(0.5) = 0.22$, $\kappa(1) = 0.23$, $\kappa(2) = 0.26$, $\kappa(f \geq 4) = 0.32$; green lines have $\tau_m = 5$ ms, $I = 0.512$, and $\kappa(0.5) = 0.25$, $\kappa(1) = 0.27$, $\kappa(2) = 0.31$, $\kappa(f \geq 4) = 0.37$; red lines have $\tau_m = 10$ ms, $I = 0.652$, and $\kappa(0.5) = 0.32$, $\kappa(1) = 0.33$, $\kappa(2) = 0.39$, $\kappa(f \geq 4) = 0.48$. In all cases, $\sigma^2 = I$. Decreasing τ_m (or I) leads to higher burst rates at all frequencies, especially high frequencies, while simultaneously not having a significant effect on the PSTHs.

This led to relatively good results, but, unfortunately, after exploring these param-

eters, including varying σ , it was found to be impossible to satisfactorily fit the firing rate and burst rates of the experimental data simultaneously. Specifically, the model had difficulty matching the experimental 2-spike burst rates at low frequencies (Figure 4.1B). Therefore, because of its direct influence on the burst probability, the depolarizing after-potential was altered to facilitate fitting.

The DAP has a number of parameters that could be augmented for this purpose, but the parameters with the greatest effect on the burst and firing rates were the somatic and dendritic widths of the DAP, labelled A and γ , respectively. Thus, A and γ were simultaneously varied to provide another unique parameter to the fitting protocol. The above method was then re-employed but with three independent variables, one of which must be held constant while pairs of the other two were varied while maintaining a spontaneous mean rate of 9.5 Hz. For example, the DAP strength was fixed while I was varied, and τ_m was chosen at each I so that the spontaneous rate was 9.5 Hz. $\kappa(f)$ can then be set by the experimental firing rate during local stimulation. This process was re-iterated for each value of the DAP strength to be investigated and further repeated in a narrower parameter range to optimize the fit. The optimal DAP size had A and γ four times their values from Noonan *et al.* [57], and the optimal local stimulation fit is shown in Figure 4.2. These local parameters are used in all global stimulation models.

4.2 Global Stimulation

The goal of parameter optimization under global stimulation is to match the global experimental data (Figure 2.9) as closely as possible. Global stimulation was simulated using the superficial cell model:

$$\tau_m \frac{dV}{dt} = -V + [I + \sigma \xi_L(t) + \kappa(f) \sin(2\pi ft)] + DAP(t) + \Lambda (w_\phi - gV).$$

Λ acts as a gating variable between local and global stimulation, and, for simplicity, was set to one and not further fitted in this section. The weight, w_ϕ , are governed by the following two equations:

$$\begin{aligned} \text{if } |t_\phi - t_B| < L_{w_{2,4}} \rightarrow w_\phi &= w_\phi + w_\phi \eta_{2,4} \left[\left(\frac{t_\phi - t_B}{L_{w_{2,4}}} \right)^2 - 1 \right] \\ \tau_w \frac{dw_\phi}{dt} &= w_{max} - w_\phi \end{aligned}$$

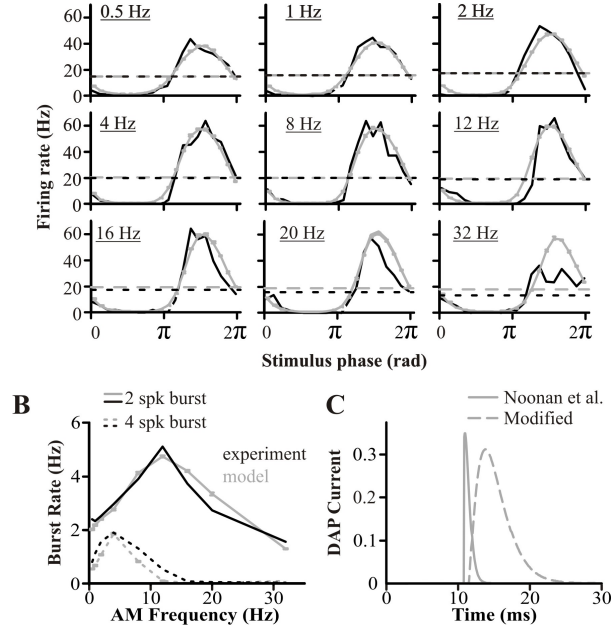


Figure 4.2: Optimal model fit of experimental data under local stimulation. The model (grey lines) and experimental data (black lines) are compared via (A) PSTHs and (B) 2- and 4-spike burst rates for different stimulus frequencies. Dashed lines in PSTHs indicate average firing rate. Solid lines and dotted lines in B represent 2- and 4-spike burst rates, respectively. The change of the DAP from Noonan *et al.* [57] is depicted in C, where the neuron model spikes at 10 ms and the DAP currents are recorded. For reference, the electroreceptor bias current, I , is 0.58, so an input current of 0.3 is small but not negligible. The optimal local parameters are $I = 0.58$, $\tau_m = 7$ ms, $\kappa(0.5) = 0.25$, $\kappa(1) = 0.27$, $\kappa(2) = 0.31$ and $\kappa(f \geq 4) = 0.39$, with DAP parameters $A = 0.6$ and $\gamma = 0.2$.

Weights were initialized at w_{max} and the system was allowed to come to equilibrium (see Section 3.4.2) before data was recorded for 1750 seconds. Weight values change according to the burst learning rules, which have parameters L_w and η for both 2- and 4-spike rules, and the potentiating rule, which have parameters w_{max} and τ_w . The LIF equation also includes global disynaptic inhibition, g , which is another global stimulation parameter.

L_{w_2} , L_{w_4} , η_2 and η_4 were obtained directly from the *in vitro* observations of burst-induced plasticity by Harvey-Girard *et al.* [41]. L_{w_2} and L_{w_4} represent the temporal width of each learning rule, whereas η_2 and η_4 represent the decrease in synaptic strength caused by a single burst pairing at zero delay. Consequently, the temporal width may be read from the 2- and 4-spike burst data directly (10 ms and 100 ms, respectively - see

Figure 4.3). However, the *in vitro* data were collected over 100 pairings and only provides the final synaptic weight value, so the values for η_2 and η_4 must be calculated. Assuming a multiplicative rule ($w \rightarrow w(1 - \eta)$ after each learning event), one can calculate the synaptic depression from a single pairing given the experimental data. For 4-spike bursts paired at zero delay, the final synaptic strength is approximately 70% of its original value after 100 pairings, so

$$0.7w_{initial} = w_{initial}(1 - \eta_4)^{100}$$

Solving for η_4 gives $\eta_4 = 0.0036$. For symmetry, η_2 was chosen as half of η_4 , which also provides a good fit of the data (Figure 4.3).

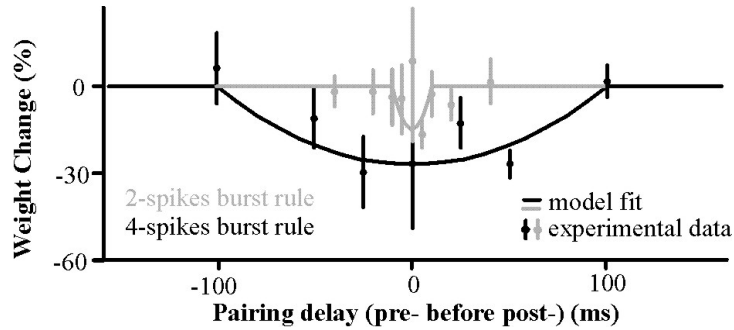


Figure 4.3: Burst-induced depression at the superficial cell-parallel fibre synapse as observed *in vitro* by Harvey-Girard *et al.* [41]. Data points are the percent synaptic weight change after 100 pairings at each pairing delay for 2-spike pairs (grey) and 4-spike pairs (black). Solid lines are calculated model fits assuming 100 pairings and a plasticity rule quadratic with pairing delay. η is proportional to the average weight decrease at zero delay while L_w is the maximum pairing delay in which a burst pair still affects the synaptic weight.

Once η_4 and η_2 are chosen, g , τ_w and w_{max} must be selected so that the appropriate PSTH shapes and burst rates are created. However, before parameter fitting can occur, one must decide which learning rules are to be used. Since it was important to ascertain whether any or all of the burst rules are necessary to mimic experimental global stimulation, this parameter optimization was accomplished for each of the 3 global stimulation paradigms - 2-spike burst learning only, 4-spike burst learning only, and both. To be biologically realistic, g , which is in units of the total membrane conductance, should be as low as possible to avoid the model relying on excessively strong disynaptic neurons. Furthermore, since potentiation has not been observed *in vitro*, τ_w should be large so that potentiation can be justified as a very slow process that is difficult to detect experimentally.

Once again, due to the complex relationships between the parameters and the burst and firing rates, a similar algorithm to find the optimal local parameters had to be employed. Different (τ_w, w_{max}, g) coordinates were chosen and compared to the firing rates, burst rates, and PSTHs of the experimentally observed neuron (Figure 4.4). For each parameter set, g was fixed by equating the firing rates of the model and experiment at 4 Hz.

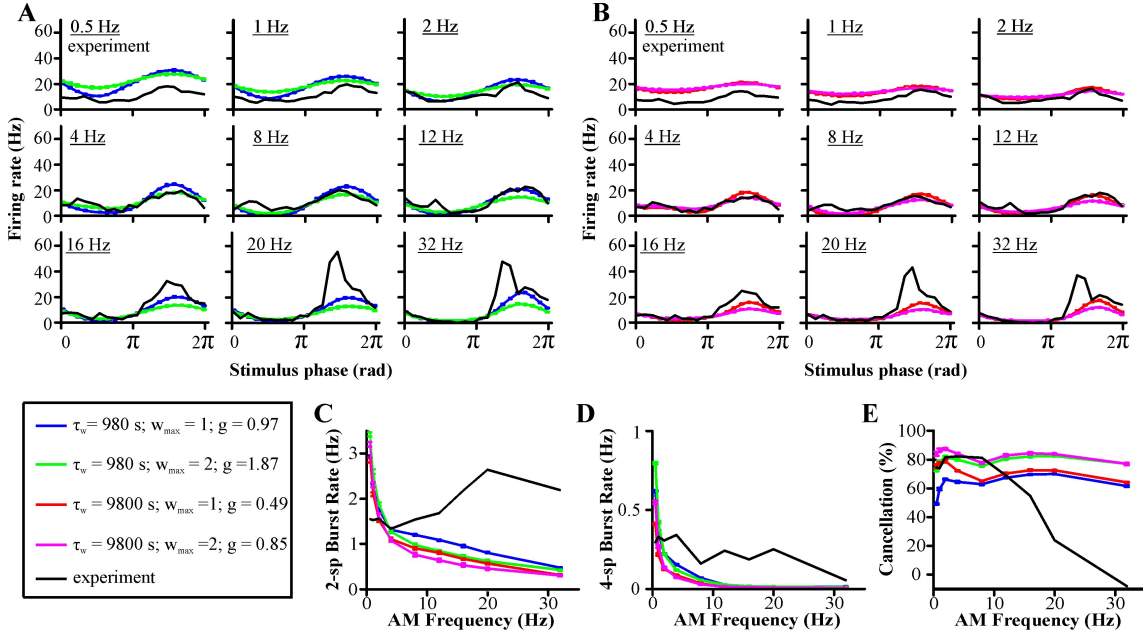


Figure 4.4: Parameter fitting to match experimental data under global stimulation when both burst learning rules are in effect. Model (A, B - two panels are used for clarity) PSTHs, (C) two-spike and (D) four-spike burst rates and (E) cancellation for different stimulus frequencies are compared to experimental data as different (τ_w, w_{max}, g) coordinates are tested. Black lines indicate experimental data. Coloured lines indicate model data. Blue lines have $\tau_w = 980$ s, $w_{max} = 1$ and $g = 0.97$; green lines have $\tau_w = 980$ s, $w_{max} = 2$ and $g = 1.87$; red lines have $\tau_w = 9800$ s, $w_{max} = 1$ and $g = 0.49$; purple lines have $\tau_w = 9800$ s, $w_{max} = 2$ and $g = 0.85$. w_{max} has the effect of increasing the firing rate in the trough of the AM input and decreasing it at the crest, which increases the cancellation at all frequencies. Increasing w_{max} also decreases the burst rates mildly. τ_w has a less pronounced effect on the model but operates similar to w_{max} : if τ_w is increased, cancellation improves and the burst rate decreases.

After optimizing each paradigm individually, the optimal fits are compared against each other and to the global data (Figure 4.5). Unfortunately, none of the three paradigms was a good match to the experimental observations. The 2-spike burst learning rule by

itself could effectively cancel high frequencies but was unable to cancel low frequencies while the 4-spike burst learning rule could only cancel low frequencies and not high AMs (Figure 4.5D). The close PSTH match of the experimental data and 4-spike burst rule only paradigm at high frequencies is not indicative of a good global fit, but simply that both are failing to cancel the feedforward signal and are reverting back to local stimulation, which is well fitted (Figure 4.2). The increase in cancellation in the 4-spike burst only paradigm at 32 Hz is due to the period being so short that two stimulus peaks now fit within the 4-spike burst definition of 45 ms, making 4-spike bursts more likely. In addition, all paradigms have difficulty matching the global experimental burst rate, where all of the model simulations show an approximately inverse relationship between burst rate and AM frequency while the experimental data does not (Figure 4.5B and C).

Nevertheless, the best fit to the *in vivo* data was produced by the paradigm that had both rules active, and was therefore selected as the best global model. Yet, with both learning rules operating simultaneously, the model actually cancels high frequency AMs better than the experimentally observed neuron despite bursting much less.

4.3 Granule Cell Activity

After it was discovered that the inclusion of both 4- and 2- spike bursts caused the model to cancel high frequencies better than the neuron *in vivo*, some of the assumptions that created the model were scrutinized. A biological mechanism that would progressively decrease the strength of the global feedback pathway at high frequencies was needed. For example, perhaps the disynaptic inhibition, g , is not constant at high frequencies as it was assumed, but weakens, possibly due to low-pass filtering properties. However, local and global experimental data at 20 Hz are nearly identical, suggesting that the feedback pathway is negligible by this frequency, and it is difficult to believe that disynaptic inhibition can exactly balance the excitatory input from parallel fibres so as to completely nullify the feedback input.

Another assumption concerning the feedback pathway was that each granule cell supplies one burst per period for all frequencies, even 32 Hz. Observations by D'Angelo *et al.* [50] reveal that this assumption is false in the rat cerebellum: *in vitro* rat granule cells forced by a periodic stimulus tended to burst once per cycle until the forcing frequency was greater than 10 Hz, after which they fired a few spikes per stimulus period in a stochastic manner with an average firing rate of about 10 spikes per second. Assuming homologous systems, these observations suggest that learning is stifled at high frequencies

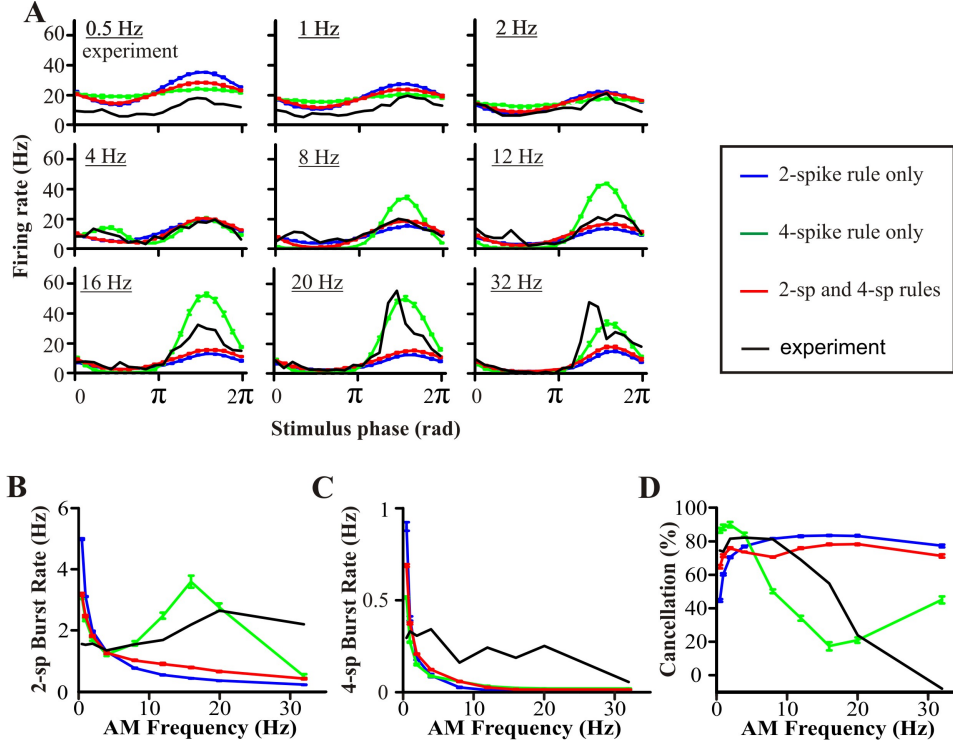


Figure 4.5: Optimal global stimulation fits for three different paradigms: 2-spike burst-rule-only (blue), 4-spike burst-rule-only (green), and both simultaneously operating (red). Model (A) PSTHs, (B) 2-spike and (C) 4-spike burst rates as well as (D) cancellation of the local stimulation response for each paradigm are compared to experimental data. The optimal parameters for the 2-spike burst rule are $\tau_w = 4900$ s, $w_{max} = 1$, and $g = 0.87$; for the 4-spike burst rule, $\tau_w = 9800$ s, $w_{max} = 2$, and $g = 1.04$; for both rules in effect, $\tau_w = 980$ s, $w_{max} = 1.5$, and $g = 1.44$.

in the electric fish ELL not because of insufficient post-synaptic superficial cell bursts, but of insufficient pre-synaptic parallel fiber bursts.

This information was included in the model, but its implementation was difficult because the model never recreates granule cells directly, only their burst events. Initially, the decrease of granule cell bursts at high frequencies was included by decreasing and individually fitting the strength of the feedback signal, Λ , after 10 Hz, in accordance with D'Angelo's observations (Figure 4.6). This had the benefit of simplicity, as Λ directly tunes the strength of the global feedback in the model (Equation 3.16), and when Λ approaches zero clearly the model will fire as it did locally. Good cancellation, PSTH, and burst rate matches to the experimental data were achieved with this new model (Figure 4.6).

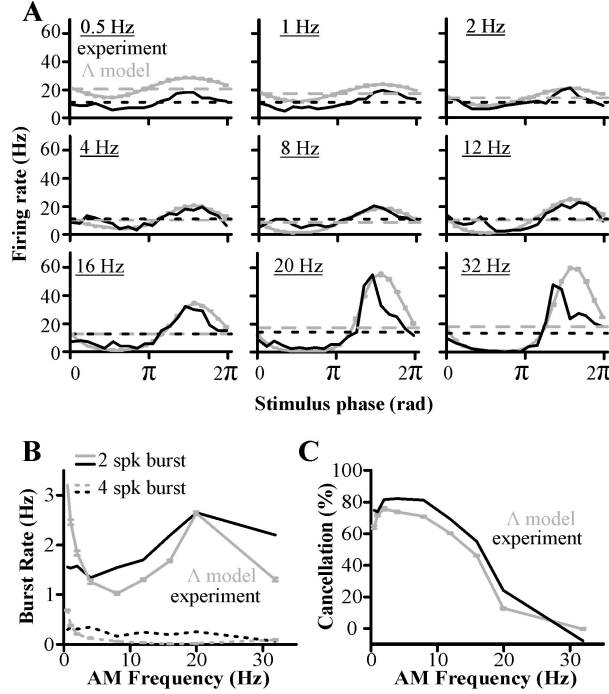


Figure 4.6: Model behaviour under global stimulation when implementing the hypothetical decay of granule cell bursting at high frequencies with a decrease in Λ . Λ is optimally fitted for frequencies above 10 Hz; the curve of the optimal Λ as a function of frequency resembles an exponential decay (Figure 4.8). (A) PSTHs, (B) 2-spike and 4-spike burst rates, and (C) cancellation data are compared between the model (grey lines) and experimental data (black lines). Dashed lines in (A) indicate the average firing rates of model (grey) and experiment (black). Solid lines and dotted lines in B represent 2- and 4-spike burst rates, respectively.

On the other hand, ignoring single spikes and only changing the strength of the feedback implies that granule cells either burst on a cycle or are silent, which is not what the experimental observations suggest. In fact, rat granule cells maintained a constant firing rate even when forced by high frequencies [50]. This indicates that the number of bursts per second decreases but the average feedback input, both directly and disynaptically, is approximately constant, as long as the same input current always results from a given number of spikes per second, regardless of their organization or individual timing. Thus, it is not biologically plausible to decrease Λ to explain the decay in cancellation.

There are other methods, however, of plausibly decreasing the feedback pathway's strength at high frequencies. Another approach was to decrease η , the learning rate,

to take into account the probability of a parallel fibre bursting when it is active, which before was unity (it always bursted), but now is less. For example, if parallel fibres fire bursts every other cycle on average, then decrease η by half. This concept had the benefit of including single granule spikes by maintaining Λ and only changing the proportion of granule cell bursts. Although η_2 and η_4 correspond to different bursts and have different values, they were decreased by the same percent at each frequency to fit the feedback decay. η values were then fitted at each AM frequency after 10 Hz, which produced a close match to the experimental data (Figure 4.7). This is the final edition of the model during global stimulation and optimally reproduces the experimental data. In this thesis, this model will be referred to as the global η model.

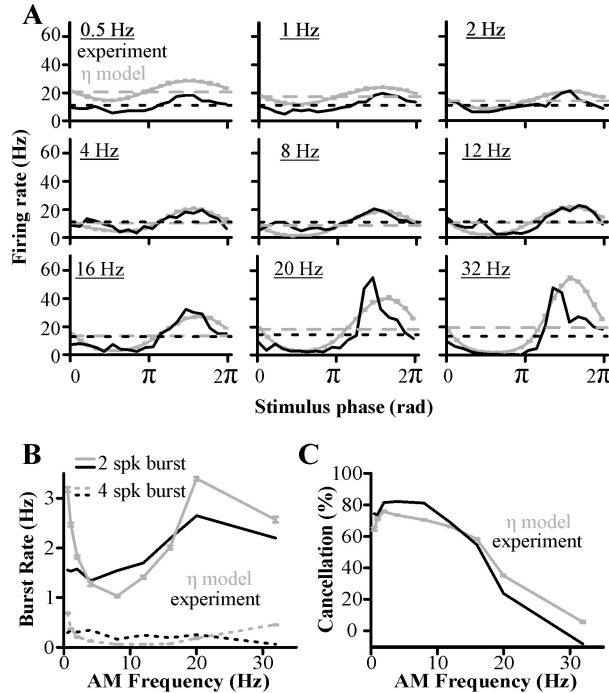


Figure 4.7: Model behaviour in the global η model, which includes the hypothetical decay of granule cell bursting at high frequencies with a decrease in η . The parameters are $\tau_w = 980$ s, $w_{max} = 1.5$, $L_{w_4} = 100$ ms, $L_{w_2} = 10$ ms, $I = 0.58$, $\sigma = 0.759$, $\tau_m = 7$ ms, and $f_{cut} = 500$ Hz. η_4 and η_2 are 0.0036 and 0.0018, respectively, for AMs less than 8 Hz but are individually fitted for higher frequencies. The curve of the optimal η as a function of frequency resembles an exponential decay (Figure 4.8). (A) PSTHs, (B) 2-spike and 4-spike burst rates, and (C). Cancellation data are compared between the model (grey lines) and experimental data (black lines). Dashed lines in (A) indicate the average firing rates of model (grey) and experiment (black).

The optimal η values led to a roughly exponential decrease of η at high frequencies

(Figure 4.8), which seems realistic given the monotonic decrease in burst rates with increasing frequency during local stimulation (see Section 6.1.1 and Discussion). In reality, there is some biological support to decrease Λ by a small amount at higher frequencies, which would help the η decay model better match experimental data. Parallel fibre synapses on both superficial cells and inhibitory neurons undergo short term facilitation, which means that a parallel fibre burst will provide more excitatory input than the sum of its individual spikes alone [66, 67]. Thus, as the burst rate decreases at high frequencies, so does the average feedback input, Λ , even while the average firing rate is constant. Nevertheless, since the goal of this thesis is to produce a minimal model, only the learning rate, η , was decreased to account for fewer granule cell bursts as it could satisfactorily fit the experimental data.

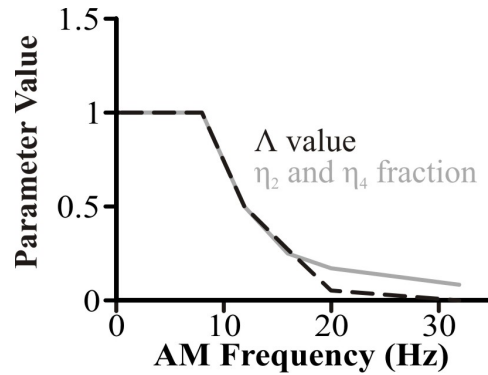


Figure 4.8: Λ and η values at high frequencies found to optimally fit the experimental data. Since Λ equals 1 under low frequency global stimulation, the black dashed line plots the actual values of Λ . η_2 and η_4 are 0.0018 and 0.0036, respectively, under low frequency global stimulation and the η values shown in grey are normalized to those values. Both η_2 and η_4 are decreased equally. The decay of either parameter represents the hypothetical decreasing burst rate of granule cells that supply the feedback to superficial cells.

Chapter 5

Frequency-tuned cerebellar channels and burst-induced LTD lead to the cancellation of redundant sensory inputs

This is a copy of a manuscript submitted to the journal of Neuroscience, written in collaboration with Dr. Gary Marsat, Prof. Len Maler, Prof. André Longtin and Dr. Erik Harvey-Girard. This manuscript briefly reviews the biological and modelling background outlined in Chapter 2 and 3 and explores the model's behaviour when different burst rules (2-spike, 4-spike, or both) are employed. The implications of the success of the model are discussed, and experimental data is presented that supports the hypothesis of independent frequency channels in the parallel fibre feedback. Parts of the discussion in this manuscript are repeated and elaborated in the discussion section of this thesis.

5.1 Contribution Statement and acknowledgements

KB did the modeling with input from GM, AL, LM and EHG.

GM did the *in vivo* experiments with input on design from KB and LM.

KB, GM, LM wrote the Ms with input from AL and EHG.

This work was supported by the Canadian Institutes for Health Research.

5.2 Abstract

For optimal sensory processing, the CNS must extract novel, unpredictable signals from the redundant sensory input in which they are embedded, but the detailed cellular and network mechanisms that implement such selective cancellation are presently unknown. Using a combination of modeling and experiment, we characterize in detail a cerebellar circuit in weakly electric fish, showing how it can carry out this computation. We use a model incorporating the wide range of experimentally estimated parallel fiber feedback delays and a burst-induced LTD rule derived from *in vitro* experiments to explain the precise cancellation of redundant signals observed *in vivo*. Our model demonstrates how the back-propagation-dependent burst dynamics adjusts the temporal pairing width of the plasticity mechanism to precisely match the frequency of the redundant signal. The model also makes the unexpected prediction that this cerebellar feedback pathway must be composed of frequency-tuned channels; this prediction is subsequently verified *in vivo*, highlighting a novel and general capability of cerebellar circuitry.

5.3 Introduction

One of the most influential theories of cerebellar function describes it as a learning device. The initial theories by Marr [28] and Albus [27], and later physiological investigations by Ito [26], argue that synaptic connections between parallel fibers and Purkinje cells are modified by experience. A central component of this theory is that an error signal originating from climbing fibers selects parallel fiber (PF) inputs by weakening (via LTD) those inputs whose timing is not appropriate for the required motor output. Although this theory is still intensely debated, an elegant body of work on a cerebellar-like structure in mormyrid electric fish confirmed the importance of PF timing and synaptic depression in shaping an adaptive filter that cancels expected sensory input [40, 3, 4]. Previous work has shown that a similar process takes place in the gymnotiform fish that have independently evolved an active electrosensory system. Our study in gymnotiform fish builds on these findings to reveal new properties of cerebellar circuits.

Apteronotus leptorhynchus is a weakly electric fish that continuously emits a high-frequency (600-1000Hz) sinusoidal electric organ discharge (EOD) into its environment to sense its surroundings and communicate with conspecifics. Small objects such as prey create spatially localized low frequency (<16 Hz) amplitude modulations (AMs) of the EOD [35, 34, 33] while tail bending or the presence of same sex conspecifics will

generate spatially global AMs in the same low frequency range. These AMs are linearly encoded by numerous electroreceptor afferents [36, 10] that project to various types of pyramidal cells within the electrosensory lateral line lobe [43, 42]. Pyramidal cells also have spiny apical dendrites in receipt of PF feedback emanating from a caudal mass of cerebellar granule cells (eminencia granularis posterior, EGp [47]. ELL pyramidal cells can be differentiated into superficial (SP), intermediate and deep varieties [46]. SP cells have extensive apical dendrites and express the NR2B subunit of the NMDA receptor as well as Ca^{2+} store proteins, while deep cells have only small apical dendrites and do not express these proteins [43, 68, 69]. SP cells respond selectively to the low frequency (<20 Hz) component of broadband noise stimuli delivered locally to their receptive field (RF) centers [44]; the response consists of both isolated spikes and spike bursts generated by complex somato-dendritic dynamics [70, 61, 71, 72]. SP cells filter out the low frequency component of the same broadband stimulus when it is delivered globally. In contrast, deep cells respond to low frequency signals delivered either locally within their RF centers or globally [38]. Earlier studies have shown that deep cells provide feedback (via cerebellar granule cells and parallel fibers, see Supplemental Figure S1) to the apical dendrites of SP cells and that this feedback permits the adaptive cancellation of global low frequency signals [38]. SP cells can therefore eliminate the global low frequency AMs due to, *e.g.* tail bending or low frequency communication signals, yet still extract the local AMs that signal the presence of prey and other local objects. Behavioral studies [33] suggest that this cancellation mechanism is essential for prey capture in the absence of vision.

PF-SP cell synapses are subject to correlative LTD in which depression occurs when both pre and post-synaptic burst firing occur within a certain time window [41]. We used this *in vitro* plasticity rule, together with the estimated distribution of conduction delays in the feedback pathway, to construct a minimal model of the cancellation mechanism. While our model relies on a bursting mechanism specific to cortical [73] and electrosensory [71] pyramidal cells, our methodology of embedding bursting dynamics in a delayed feedback loop that includes cerebellum and a correlative learning rule can clearly be applied to other systems. Our state-of-the-art cancellation model combines realistic integrate-and-fire dynamics, back-propagation (DAP)-based bursting, intrinsic noise, periodic external stimulation, feedback with a range of delays, feedback-dependent shunting, and a burst-spike-timing-based correlative learning rule. Our model approach required several stringent conditions to be met: 1) the canceling feedback had to be frequency specific, 2) the propensity to burst, and especially to produce large bursts, had to

Parameter	Value	Parameter	Value	Parameter	Value
V_{thresh}	1	σ	0.759	η_4	3.6×10^{-3}
V_{reset}	0	f_{cut}	500 Hz	η_2	1.8×10^{-3}
τ_m	7 ms	g	1.44*	L_{w_4}	100 ms
τ_{ref}	0.7 ms	τ_w	980 s	L_{w_2}	10 ms
I	0.576	w_{max}	1.5		

Table 5.1: Parameter values for the model’s Equations 5.1, 5.6, and 5.7. * $g = 1.5$ if only the large-burst rule is used and $g = 1.66$ if only the small- burst rule is used (see Results).

decrease with signal frequency, while 3) the duration of the temporal window of synaptic plasticity had to depend on burst size, and thus on signal frequency. Experimental data confirmed all of these predictions.

5.4 Materials and Methods

5.4.1 LIF model

This model of a superficial E pyramidal cell in the centro-lateral segment (CLS) of the ELL is based on the leaky integrate-and-fire framework and replicates the model by Noonan *et al.* [57]. Specifically, the voltage evolves according to the equation:

$$\tau_m \frac{dV}{dt} = -V + [I + \sigma \xi_L(t) + \kappa(f) \sin(2\pi ft)] + DAP(t) + \Lambda(w_\phi - gV) \quad (5.1)$$

When the membrane potential, V , crosses the threshold, V_{thresh} , a spike is recorded and V is reset to V_{reset} . V is maintained at V_{reset} for an absolute refractory period, τ_{ref} , after which V continues to evolve according to Equation 5.1. Here, τ_m is the membrane time constant, I is the input current when the receptor afferents have no stimulus input, $\xi(t)$ is a low-pass filtered Gaussian white noise source with variance 1 and cutoff frequency f_{cut} , and σ is the standard deviation of this noise. During local stimulation the strength of the feedback, Λ , is set to 0 since local stimulation does not drive the feedback. During spontaneous activity, the strength of the stimulus, $\kappa(f)$, which can depend on the particular frequency component driving the system, is set to 0. Other terms are explained below. Parameter values for the model are summarized in Table 5.1.

5.4.2 Modeling the bursting dynamics

Bursting in pyramidal cells is due to back-propagating dendritic spikes. DAP(t) in Equation 5.1 represents the depolarizing after-potential, an injection of current into the soma of the neuron after an action potential is fired due to active channels in the cell's dendrites [61]. This effect has been modeled previously in E superficial cells by Doiron, Noonan and colleagues [61, 57]. Their model was used in this paper with minimal parameter changes: A and γ have simply been increased to match the bursting behavior observed *in vivo*.

Briefly, after the cell fires ($V = V_{thresh}$), and assuming the previous firing time is not too recent, the cell will receive a DAP, *i.e.* a small current injection a short time later. This extra stimulation is modeled as a difference in alpha functions $s(t, a)$ (Equation 5.4 below): one generated by the soma, and the other by the dendrites. If, however, the interval between this spike time and the previous spike time is less than the refractory period of the dendrite, r_d , then the DAP is inactive for the current spike. r_d is modeled as a dynamic variable that changes according to a secondary variable, b , which also controls the width of dendritic alpha function. All spikes generated by the neuron are recorded, the most recent of which was at time t_n , and b updates whenever the neuron fires a spike. t_n^+ refers to the time just after the most recent spike was fired. The equations governing the DAP [57] are

$$DAP(t) = \begin{cases} 0 & \text{if } t - t_n < r_s \\ \alpha s(t - t_n, \beta b(t_n^+)) - \alpha s(t - t_n, \gamma) & \text{if } t - t_n > r_s \text{ and } t_n - t_{n-1} > r_d^n \\ 0 & \text{if } t - t_n > r_s \text{ and } t_n - t_{n-1} < r_d^n \end{cases} \quad (5.2)$$

$$\frac{db}{dt} = -b/\tau + (A + Bb^2) \sum_n \delta(t - t_n) \quad (5.3)$$

$$s(t, a) = \frac{te^{\frac{t}{a}}}{a} \quad (5.4)$$

$$r_d^n = D + Eb(t_n^+) \quad (5.5)$$

The parameters used in the above equations are listed in Table 5.2.

Parameter	Value	Parameter	Value	Parameter	Value
A	0.6*	B	2	α	20
γ	0.2*	β	0.35	τ	τ_m
r_s	$0.1\tau_m$	D	0.1	E	3.5

Table 5.2: Parameters used in the DAP model. * indicate changes from [57].

Frequency (Hz)	κ	Frequency (Hz)	κ	Frequency (Hz)	κ
0.5	0.25	4	0.39	16	0.39
1	0.27	8	0.39	20	0.39
2	0.31	12	0.39	32	0.39

Table 5.3: Value of $\kappa(f)$ for the different stimulation frequencies.

5.4.3 Local stimuli inputs (feedforward)

The electroreceptor afferent input could be modeled as a Poisson process whose mean value changes as a function of the phase of the input frequency. However, since many afferents converge onto one pyramidal cell and since their baseline firing rate is very high (~ 200 Hz [10]), we can simplify the input as a bias current and a Gaussian noise source. Since the electroreceptor input is strictly excitatory, the modeled feedforward input is rectified and [...] in Equation 5.1 symbolizes rectification. The values of $\kappa(f)$ were fitted for each stimulus frequency to match firing and burst rates with experimental data (Table 5.3). The lower κ values at frequencies < 4 Hz reflect the adaptation of the electroreceptor [59].

5.4.4 Parallel fibre inputs (feedback)

A single parallel fiber is assumed to fire at the same unique phase of each stimulus cycle and be silent at all other times (see Results below for justification). Due to the distribution of delays, the total parallel fiber input is assumed to be continuous. In the model, the parallel fiber input was modeled as a current input (of strength Λ in Equation 5.1) and was discretized into 2.5 ms non-overlapping segments with each segment having a unique synaptic weight. The number of segments was chosen so that the entire period of each stimulus frequency was completely covered with parallel fiber input (*i.e.* 100 parallel fibers bins for 4 Hz, 50 bins for 8 Hz, etc.). This created a constant input before learning began. All weights were initialized at w_{max} .

During learning, burst-induced depression decreased the synaptic weight according to the rule recently identified *in vitro* [41]. The strength of the burst-induced depression rules, η_2 and η_4 , are kept constant across frequencies in the initial version of our model (Figure 5.4) but vary with higher frequencies in the final version (see Results and Figure 5.5). The burst induced depression of the synaptic weights is governed by Equation 5.6. At each time step, the weights slowly relax back to w_{max} with a time constant of τ_w according to Equation 5.7:

$$\text{if } |t_\phi - t_B| < L_{w_{2,4}} \rightarrow w_\phi = w_\phi + w_\phi \eta_{2,4} \left[\left(\frac{t_\phi - t_B}{L_{w_{2,4}}} \right)^2 - 1 \right] \quad (5.6)$$

$$\tau_w \frac{dw_\phi}{dt} = w_{max} - w_\phi \quad (5.7)$$

The response of the model was always quantified after weight values came to equilibrium. The disynaptic inhibition produced by the parallel fibers is modeled as an extra shunting conductance, $-gV$ in Equation 5.1, that also varies with the PF feedback strength Λ . The disynaptic inhibition is kept constant across phase, as we have no experimental evidence that LTD occurs at these synapses (either on the input or output of the inhibitory interneurons).

5.4.5 Burst definition and learning rule

Through the analysis of *in vivo* recordings, we determined that the inter-spike intervals within a burst are typically <15 ms in SP cells (data not shown) and used that criteria to identify bursts in the SP cell model. Bursts were categorized as small (2 or 3 spikes) or large (4 or 5 spikes) while less frequent even larger bursts were divided into smaller bursts for the purpose of applying the burst-based learning rule. For example, a 6-spike burst is counted as a large 4-spike burst followed by a small 2-spike burst (note that large bursts have priority: we do not divide a 6-spike burst into three 2-spike bursts). The timing, t_B , of the first spike in a burst determines the center of the depression function that is applied to the weights according to Equation 5.6. The strength and width of the depression rules are different for small and large bursts (η_2 , L_{w_2} and η_4 , L_{w_4} respectively).

5.4.6 Quantitative analysis

All model data for local stimulation were generated by collecting data over 1750 simulated seconds and the averaged response per period of stimulus frequency was calculated. For

global stimulation, the system was allowed to come to equilibrium for 3500 seconds and then data was tabulated during the next 1750 seconds. There was no change in results if the model was allowed to come to equilibrium for 5250 seconds. As the feed-forward transmission delays of the receptor afferents were not included in the model, the *in vivo* PSTH data was shifted in phase to optimally overlap the experiment's PSTH for each stimulus frequency. The same phase shifts were used in both local and global PSTH plots. The disynaptic inhibition was fitted such that the model's average firing rate for 4 Hz matches the data's average firing rate.

Cancellation at a given frequency is defined as

$$\text{Cancellation} = \left(1 - \frac{\text{Amp}_{\text{global}}(f)}{\text{Amp}_{\text{local}}(f)} \right) \times 100\%$$

The amplitude at a given frequency was defined as the amplitude of a sine wave fitted to the model or experimental data's PSTH induced by that stimulus frequency.

5.4.7 *In vivo* electrophysiology

Details of the surgery and recording techniques are as described previously [54, 53]. Briefly, craniotomy is performed under general anesthesia. During the experiment, the fish is awake but paralyzed with curare and locally anesthetized. Single-unit extracellular recordings from superficial pyramidal cells of the CLS of the electrosensory lateral line lobe were performed.

Stimuli consisted of amplitude modulations of the fish's own electric field. The stimulus was delivered through two large global electrodes placed on each side of the fish thereby achieving a global stimulation. For local stimulation, a small dipole was placed in the center of the cell's receptive field; the distance between the dipole and the skin was adjusted to maximally stimulate the whole receptive field of the cell while avoiding stimulation of receptors outside the classical receptive field. The intensity of both local and global stimuli were adjusted so that the modulation was 10-15% of the fish's own electric field as measured near the cell's receptive field. In the experiment revealing the frequency selectivity of the feedback pathway, the training stimulus consisted of both a normal global stimulus to which a strong local stimulus of the same frequency and phase was added. As a result, the cell receptive field was stimulated at a stronger intensity than what would result from a normal global stimulus. The paradigm was shown to induce plasticity in the feedback pathway [74], allowing the strength of the feedback to adjust

to the strength of the feedforward input. This training stimulation at a single frequency was presented for 3 min after which 15 s stimuli of two different frequencies (the training frequency and another one) were interleaved and played in a normal global stimulation configuration. All experimental procedures were approved by the University of Ottawa Animal Care Committee.

5.5 Results

The goal of our model was to replicate the cancellation of low frequency redundant signals observed *in vivo* during global stimulation (Figure 5.1). Contrary to stimuli restricted to the pyramidal cell’s RF center (classic RF), global stimuli effectively drive both the classical RF and, via a cerebellar (EGp) feedback pathway, also activate a non-classical RF [38, 44]. By comparing responses to local and global stimuli, we can quantify the canceling effect of the feedback. Under local stimulation, the cell responds strongly above its baseline rate at the stimulus peak but far below baseline at the trough. During global stimulation, the differential response is mostly eliminated for frequencies below 10 Hz (e.g. 2 Hz in Figures 5.1A-B, left panel). However, this cancellation deteriorates drastically between 10 and 20 Hz (Figures 5.1A-B right panel, C). The main question motivating our modeling effort was to determine whether the *in vitro* LTD rule [41] together with known anatomical and physiological constraints could explain this cancellation.

5.5.1 Model replicates experimentally observed responses in the absence of feedback

We used the leaky integrate-and-fire formalism to model SP cells but added a burst-inducing depolarizing after-potential (LIF-DAP model) that incorporates the effect of back-propagating dendritic spikes. This model is a reduction of biophysically realistic models [61] and has previously been shown to capture the essential details of SP cells burst discharge and electrosensory stimulus encoding [57]. Baseline discharge was modeled by delivering a bias current, I , to the LIF-DAP model (Figure 5.2A) with noise added so as to match the SP cell’s mean baseline discharge (mean 9.5 Hz, SD = 3.1, N = 9 neurons). SP cells are linearly modulated by sensory input [38, 10]. Thus, direct sensory input to the SP cells was modeled by sinusoidally modulating the bias current. Since electroreceptors adapt to AM < 4Hz [59], the strength of this modulated input was adjusted at these frequencies to match the experimentally observed firing and burst

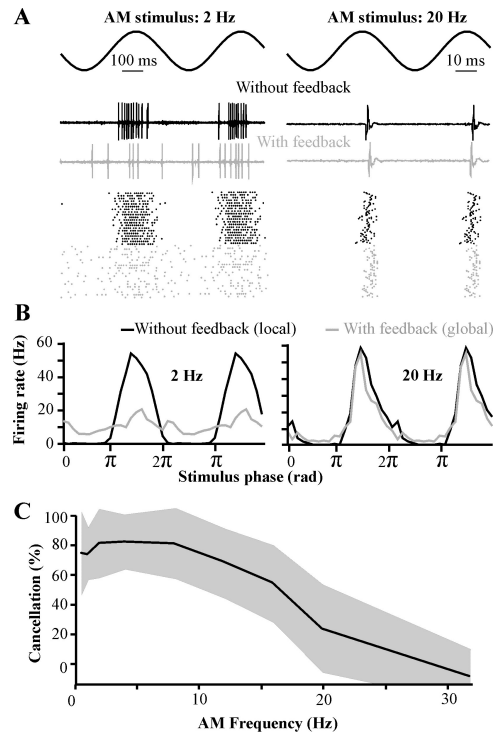


Figure 5.1: Properties of the system to be replicated by the model: local stimuli (no feedback) elicit a strong phased-locked response whereas global stimuli - which do recruit the cerebellar feedback - elicit attenuated responses to low frequency sine waves. This cancellation is due to feedback inputs. **(A)** Examples of the responses of one neuron to sinusoidal stimuli of low or high frequencies (2 Hz or 20 Hz; top trace). The spike trains elicited by two cycles of the stimulus are displayed (middle traces) along with raster plots of the responses to many cycles (bottom trace; note that the raster plots displayed for the two consecutive cycles are the same; the data are duplicated in this fashion to improve visualization). **(B)** Average firing rate of superficial E-cells of the centro-lateral segment of the ELL ($n=9$) elicited by one cycle of the stimulus (the data is duplicated on an adjacent cycle for clarity). **(C)** Cancellation of the stimulus response due to feedback as a function of stimulus frequency (mean \pm S.D.; $n=9$). Cancellation is calculated by first fitting a sine wave (with the same frequency as the stimulus) through the mean firing rate data (e.g. see panel **B**) in order to quantify the amplitude of the response to local and to global stimuli. Cancellation is then calculated as the ratio of these amplitudes (see Materials and Methods).

rates under local conditions (see Material and Methods).

As shown in Figure 5.2C, our LIF-DAP model accurately matched the observed responses of SP cells to local stimulation for frequencies up to 32 Hz despite neglecting the effects of feedforward disynaptic inhibition [63] or possible nonlinear biophysical dynamics.

The DAP dynamic causes both the cell’s burst rate and number of spikes/burst to vary with stimulus frequency; our model correctly captured the experimentally determined relationship between stimulation frequency and burst rate and spikes/burst (Figure 5.2B). These bursts are important for both stimulus encoding [75, 45, 53] as well as for induction of LTD at the synapse between PF and pyramidal cells [41]. Note that in the latter case, small and large bursts result in different associative time windows for LTD (see below). Therefore, in order to parallel the *in vitro* experiments, bursts were divided into small (2 or 3 spikes) and large (4 or 5 spikes, see Materials and Methods). The burst dynamics of our LIF-DAP replicated the *in vivo* bursting propensity for both small and large bursts during local stimulation (Figure 5.2B).

5.5.2 Incorporating a feedback that cancels low frequency coding

The cerebellar feedback pathway to the ELL originates from non-plastic cells of the ELL which project bilaterally to the nucleus praeminentialis (nP). nP cells project bilaterally to the granule cells of the EGp and these granule cells in turn project bilaterally to the SP cells of the ELL [36, 47]. This produces a wide range of feedback delays dependent on the length and conduction velocity associated with each possible trajectory. An additional source of delay arises from the spatial extent of the EGp ($\sim 1200 \mu m$ [76]); cells at one edge of the EGp will incur an additional delay of 12 ms (given a conduction velocity of 0.1 m/s [40]) compared to cells at the other edge. The delays resulting from the various possible trajectories vary from ~ 20 to >70 ms. Furthermore, the origins of the feedback pathway, the deep cells of the ELL, have a variety of phase relationships with the stimulus depending on the side of the body they represent and on their functional type. Indeed, these cells can be divided in two types: E-cells, responding throughout the peak of the stimulus, and I-cells, responding throughout its trough [46, 53].

Considering this variety of delays and phase relationships, and assuming that granule cells of the EGp phase-lock to their input [50], PF can provide feedback to SP cells at all possible phases of the low frequency AMs that are canceled during global input. In

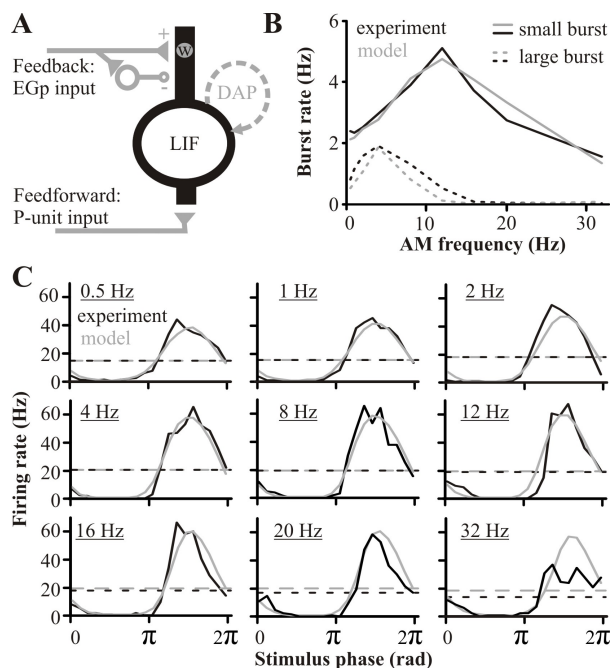


Figure 5.2: Model exploration: replication of pyramidal cells properties when stimulated locally (no feedback) and constraints imposed on the model during global stimulation (with feedback). **A**) Schematic of the main components of the model: a leaky integrate-and-fire (LIF) model cell receiving feedforward input from the receptors and feedback inputs from the EGp. A depolarizing after-potential (DAP) was added to replicate the neuron's burst propensity. The feedback inputs have a direct excitatory component delivered through plastic synapses (of weight w) and a disynaptic non-plastic inhibitory component. This feedback input is active only when modeling the response to global stimuli. **B**) Comparison of *in vivo* and modeled pyramidal cell responses during local stimulation: burst rates as a function of stimulus frequency. Bursting is quantified by dividing the spike trains into small (2 or 3 spikes) or large bursts (4 or 5 spikes); longer bursts are taken as combinations of small and large bursts (see Materials and Methods). **C**) Post-stimulus time histograms comparing the response of *in vivo* ($n=9$ cells) and modeled pyramidal cells elicited by local stimuli of different frequencies. Dashed lines represent the average firing rate.

addition, granule cells have been shown to phase-lock to periodic inputs and to burst to sensory stimuli and be silent elsewhere [48, 50, 49, 77]. We therefore assumed in our model that (1) the array of PFs provide feedback to SP cells at all possible phases of the stimulus and (2) EGp granule cells respond in a bursty manner, phase-locked to low-frequency AM signals.

The feedback pathway provides direct excitatory and disinaptic inhibitory input to SP cells [62] and both were replicated in our model. The excitatory input was modeled as a feedback strength, Λ , multiplied by the synaptic strength of the PF-SP cell synapse, w , that can be altered by the burst induced plasticity identified experimentally [41]. The inhibition is also proportional to the feedback strength but is otherwise constant because no long term plasticity has been found at these synapses. The reversal potential of the inhibitory synapse (GABA-A receptors) is near the SP cell’s resting potential [78], so inhibition was modeled as a constant membrane conductance increase, g . The overall strength of the excitatory feedback was adjusted to be slightly larger than inhibition, consistent with *in vivo* and *in vitro* physiological studies [74, 78]. Under global stimulation, the firing activity of SP cells is modeled by the same Equation 5.1 given in the Materials and Methods, but the strength of the feedback Λ is set to 1.

The strength of the synapses between PFs and SP cells were modeled as weights that span the stimulus period (2.5 ms bin widths; see Materials and Methods) and were initially all equal since the array of PFs provide inputs at all possible phases of the stimulus. If the SP cell bursts in close temporal proximity to the firing time of PFs, the corresponding synaptic weight will depress according to the depression rule identified *in vitro* [41]. The different rules identified for small and large bursts, are based on *in vitro* experiments that paired 2-spike bursts or 4-spike bursts pre and post-synaptically. A variety of pre and post-synaptic burst sizes can occur *in vivo* and the plasticity rule might therefore continuously interpolate between rules based on different burst sizes. For the sake of simplicity and to remain consistent with *in vitro* data, we restrict our model to the experimentally defined small burst and large burst categories. The corresponding plasticity rules were modeled as quadratics (Figure 5.3) having different depression strengths and delay ranges. To prevent complete depression of all PF feedback, we included a recovery process by which all weights slowly relax back to their default value in an activity and frequency-independent manner.

We initially only took into account large bursts and the plasticity they elicit where the pairing time window for LTD is ± 100 ms. By itself, this rule was only effective for canceling AM frequencies < 5 Hz because the plasticity time window was too broad to

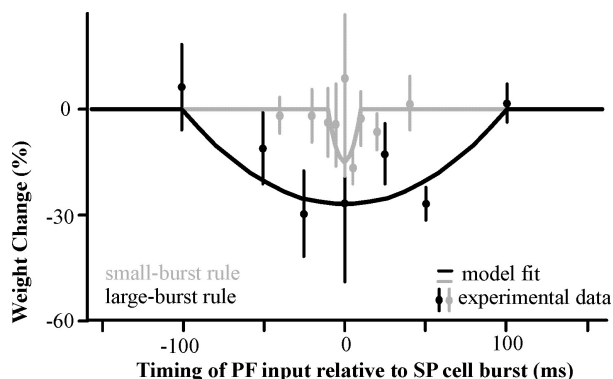


Figure 5.3: Burst-induced depression measured *in vitro* (data points) and applied to the parallel fiber-pyramidal cell synapse of the model (lines) for small bursts and large bursts (gray and black, respectively). The measured weight change is the result of 100 pre- and post-synaptic burst pairings. Note that the large-burst rule is not equal to the sum of two small-burst rules.

selectively target the positive phase of the AM of higher frequencies (Figure 5.4A). For example, for an 8 Hz signal, a burst at the crest of the stimulus will depress weights 75 ms later, at its trough, and therefore reduce the excitatory feedback required to cancel the trough. Trying to use the narrower, small-burst rule by itself -counting longer bursts as several small bursts- also failed to replicate the experimental results (Figure 5.4A), in part because it canceled high frequencies too much (see below), but also because it did not cancel low frequencies enough.

The inclusion of both burst rules acting simultaneously in the model allows cancellation to be effective at a much larger range of frequencies, accurately reproducing the cancellation observed in SP cells at frequencies below 16Hz (Figure 5.4A). Therefore, the non-linear dependence of the learning rule time window on the size of the burst is essential to adequately shape the feedback synaptic strength: the temporal precision associated with the small burst is required to cancel high frequencies while the strong impact of large burst is required to cancel low frequencies. However, discrepancies between the model and the experimental data remain at very low (0.5 and 1 Hz) and at high frequencies (20 and 32 Hz). At low frequencies, the model cancels as well as the real pyramidal cells but the average firing rate is slightly higher in the model (Figure 5.4B). This minor discrepancy may be due to intrinsic [79] or feedforward inhibitory [80] adaptation dynamics which were not included in our model.

The more serious discrepancy is that our model permits cancellation of not only low, but also high frequencies (16-32 Hz), which is not seen *in vivo* (Figure 5.4B). Yet, under

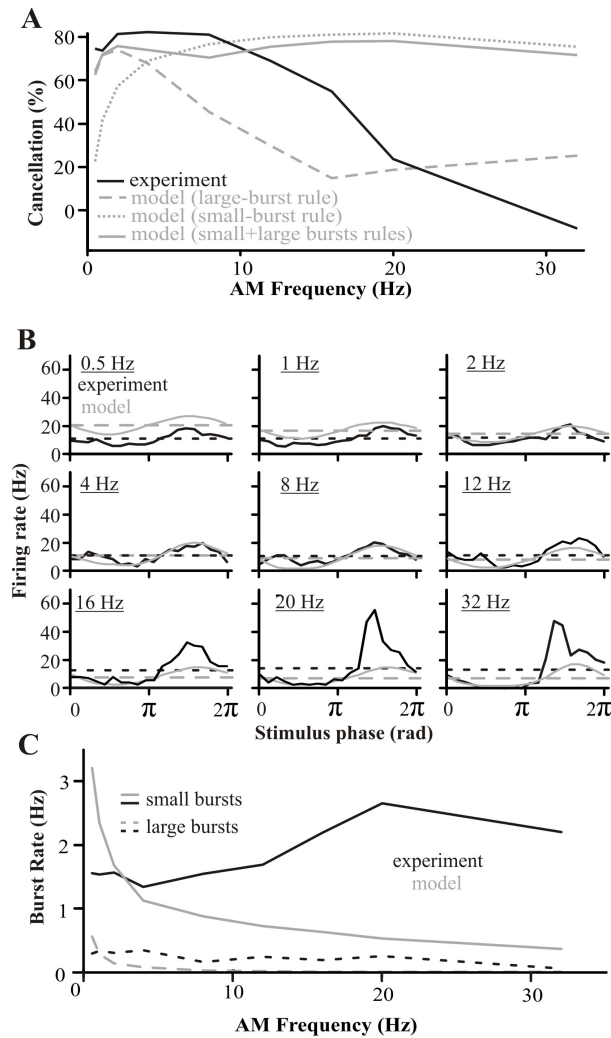


Figure 5.4: Comparisons between model and experimental responses when the model learned either with the large-burst rule, the small-burst rule or both. Note that, in this figure, parallel fiber activity is as yet unconstrained (compare to Figure 5.6). **A**) Cancellation performance of the model compared to experimental data. **B**) PSTH of the model and *in vivo* responses (global stimuli) when both large and small burst rule are used while the model is learning. Dashed lines represent average firing rate per second. **C**) Burst rates in model and *in vivo* responses when both large and small burst rule are used while the model is learning. Note the difference in small burst rate at high AM frequencies.

global stimulation, the model actually bursts less than that observed *in vivo* (Figure 5.4C) implying that cancellation is failing *in vivo* despite SP cells producing enough bursts to drive plasticity. The difference may come from the fact that our model thus far assumed that each burst in the SP cell is matched by a burst in all the PFs representing a given phase. *in vitro* studies of mammalian granule cells have shown that they produce spike bursts in response to sinusoidal current injection but that the number of spikes per burst drops off with increasing frequency, and, after 10 Hz, the cells failed to produce bursts on every cycle (Figure 2A of [50]). It is therefore likely that, at higher AM frequencies, a burst in the pyramidal cells will not be matched by a burst in every PF representing a given phase and thus will not elicit depression in all the corresponding synapses.

We incorporated the declining probability of a pre-synaptic burst into the learning rule by optimally decreasing the parameter η as a function of frequency to reproduce the experimental data (see Figure 5.5). We found that an approximately exponential decrease in η above 8 Hz was sufficient to obtain the decrease in cancellation at higher frequencies seen *in vivo* (Figure 5.6). Further studies will be required to test whether the EGp granule cell burst size is inversely related to mossy fiber input frequency as predicted by our model.




Figure 5.5: Frequency dependence of learning. The decreasing propensity for granule cells to burst at higher frequencies is implemented by changing the learning rule strength η . For each stimulus frequency, we determined the value of η that resulted in an optimal fit between model and experimental results. We plot here η as a percent of the original values (0.0036 and 0.0018 for large- and small-burst rules respectively).

5.5.3 Independent parallel fibre frequency channels revealed *in vivo*

We have experimentally demonstrated that good cancellation of global frequencies < 16 Hz can be achieved and our model faithfully captures this result. A key element

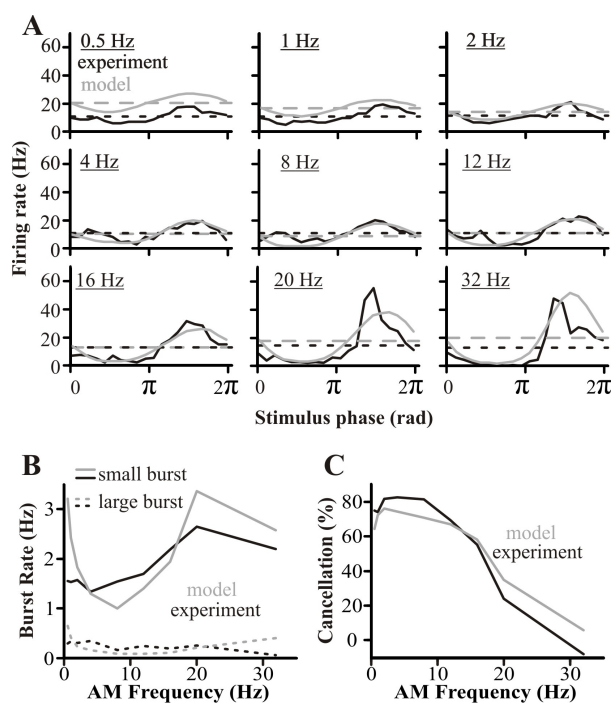


Figure 5.6: Comparison of the model with *in vivo* data showing its ability to replicate the feedback induced cancellation. This model now includes the frequency dependence of granule cell bursting. **(A)** PSTH of the model and *in vivo* responses to global stimuli of different frequencies. Dashed lines represent average firing rate per second. **(B)** Burst rates in model and *in vivo* responses. Bursting responses were segregated into small and large bursts as described previously (see also Materials and Methods). **(C)** Cancellation measured for responses of this final version of the model correspond well to the *in vivo* data.

of our model - the existence of PFs representing all phases for any stimulus frequency - is dependent on anatomical and physiological data demonstrating a wide range of PF feedback delays to SP cells. This implies, however, that a delay associated with a particular phase for one frequency may be at a different phase for another frequency. For example, a PF whose delay is associated with the peak of a 4 Hz AM will be associated with the trough of a 12 Hz AM. To keep the phase definition of each PF consistent, our model assumed that the weights learned at one frequency are independent of learning at another frequency. **Our model therefore predicts the existence of independent frequency-tuned PF feedback channels to allow the burst LTD rule to select the right time delays necessary to cancel each frequency.**

To test for the possible existence of independent frequency channels in the feedback pathway *in vivo*, we took advantage of the fact that the synaptic plasticity between parallel fibers and pyramidal cells also allows adjustment of the feedback strength relative to the feedforward input [74]. Two sinusoidal stimuli (2 Hz, 8 Hz) were presented individually to an E-type SP cell both locally and then globally in order to assess the initial degree of cancellation; cancellation was achieved for both frequencies (Figure 5.7). Following a protocol introduced by Bastian [74], we then paired the global presentation (G) of one frequency (8 Hz in the example shown in Figure 5.7) with the local presentation of the same frequency (L) AM at an intensity such that the summation of both local and global AMs strongly drive the neuron to burst discharge. The response of the cell to the G+L signal diminished as the number of presented cycles increased, presumably due to the strong evoked burst discharge inducing plasticity at PF synapses. Following the G+L protocol, the original stimuli were again presented globally to measure the change in cancellation. For the stimulus at the same frequency as the G+L pairing, the SP cell's spiking response was, for the first 30-60 s, out of phase with the stimulus: the trough of the stimulus now elicited spiking while its peak silenced the cell. This response is diametrically opposite to that expected from the anatomy and physiology of an E-cell and cannot be due to the direct input from electroreceptors [81]. As previously shown by [38], this reversed response is due to the temporarily modified PF feedback synapses, which provide a feedback input too strong relative to the feedforward input and causes an over-cancellation of the response. Most importantly, the global response to presentation of the 2 Hz AM was unchanged and cancellation remained similar before and after the training with 8Hz; additional experiments revealed that the opposite protocol (pairing of local and global AMs for the 2 Hz stimulus) produced exactly the same frequency specific change in the feedback (data not shown). **We conclude that our prediction**

was verified: PF feedback cancellation input to SP cells is frequency specific.

5.6 Discussion

We have shown that superficial pyramidal neurons of the electric fish, *A. leptorhynchus* can cancel spatially redundant sinusoidal signals (AMs) with frequencies < 16 Hz. It is notable that this frequency range matches the spatially localized AMs generated as a fish scans past small prey [33], as well as the spatially global AMs self-generated by tail movements or by the presence of same sex conspecifics. The cancellation mechanism therefore permits SP cells to detect prey and, presumably, the fish to capture prey even when prey-induced local electrosensory signals are contaminated by global AMs with overlapping frequencies. Previous work had established that the cancellation mechanism involves feedback projections onto SP cells that emanate from non-plastic deep pyramidal cells and reach SP cells via cerebellar granule cells [38]. We developed a minimal model of the cancellation mechanism constrained by well-established neurophysiological and anatomical data: (a) the spiking response of SP cells to stimulation of their receptive field centers; (b) the dynamics and AM frequency-dependence of the SP cells burst mechanism; (c) the known delays in the PF feedback pathway; (d) the frequency-dependent adaptation of electroreceptor afferents; and (e) *in vitro* PF-SP cell synaptic plasticity rules: burst-dependent correlative LTD with a plasticity pairing window that is proportional to the number of spikes per burst.

The model we have developed is minimal in the sense that we have not incorporated additional physiological data that might improve its performance. For example facilitation of PF-evoked EPSPs has been observed *in vitro* [66],[67]; thus, the presynaptic enhancement of PF-evoked EPSPs can drive SP cells and local interneurons more strongly. This effect would improve the cancellation at low frequencies and contribute to the cancellation decay at high frequencies by slightly decreasing the average feedback strength as PF bursts become rarer. However, the inclusion of different forms of PF plasticity and plasticity of disynaptic inhibitory neurons makes resulting models far more complex and the connection between biology and model assumptions less transparent.

Remarkably, our minimal model was quantitatively consistent with the response of SP cells to stimulation of their RF center and to the development of cancellation following stimulation of the RF center and non-classic RF - the source of the feedback that drives cancellation. The *in vitro* experiments had shown that large bursts produced long plasticity windows $> \pm 50$ ms while the small burst windows were $< \pm 10$ ms

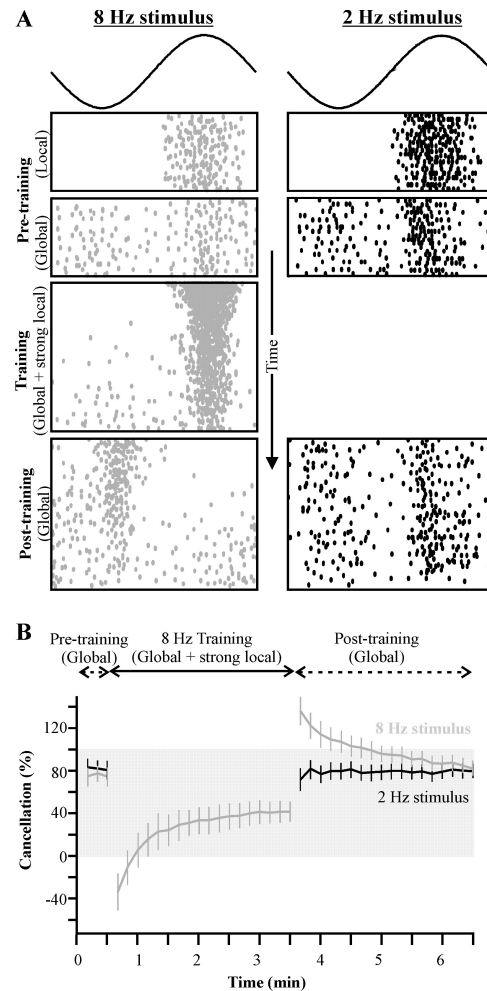


Figure 5.7: Frequency-specificity of parallel fiber inputs demonstrated *in vivo*. **(A)** Raster plot showing, as a function of time, the response of a typical cell to the different steps of the experimental protocol: pre-training baseline for two AM frequencies delivered either locally (no feedback) or globally (thus recruiting feedback); training stimuli of a single frequency where a local stimulus is added to a global one; post-training responses to global stimuli revealing the effect of the induced plasticity on the response to the two frequencies. **(B)** Quantification of the mean (\pm s.e.; $n=11$ cells) canceling impact of the feedback. Values below zero mean that the response is stronger than the uncanceled baseline (local stimulation). Values above 100 indicate an over-cancellation, *i.e.* the phase of the response is shifted 180°; the higher the value, the stronger the response in anti-phase.

[41]. We found that the small bursts were sufficient to produce excellent cancellation for high frequency inputs but the plasticity window was too narrow to cancel low frequency stimuli. Conversely the large temporal window associated with large burst could not achieve cancellation of high frequency AMs. The use of different temporal plasticity windows for bursts of different sizes, and the dependence of burst size on the temporal frequency of the stimulus were critical to adequately shape the canceling feedback across AM frequencies.

The fact that lower frequencies AMs elicit bursts with more spikes is by no means surprising since these larger burst can be expected to lead to more calcium entry in the post synaptic cells via postsynaptic NMDA receptors [82, 68]. Depending on the type and amount of calcium binding and sequestering proteins [83, 84] and their competition with LTD promoting second messenger systems [68], this variation in Ca^{2+} influx might affect the duration of the LTD window. It may be useful to look for such dynamic matching of LTD pairing windows in other sensory systems with a large range of input frequencies, *e.g.* in the vestibular system [85].

An additional constraint was necessary in our model to explain the deterioration of cancellation at AM frequencies >8 Hz. Mammalian granule cells (*in vitro*) only burst consistently on every cycle for frequencies ≤ 10 Hz and the number of spikes/burst appears to be inversely proportional to frequency [50]. Incorporating such a frequency-dependent drop of PF bursting produced an excellent quantitative fit between model and experimental data. Our model thus predicts that, similarly to rat cerebellum, bursting in EGp granule cells of electric fish decreases with stimulus frequency as described in figure 5.5. Our result suggest that frequency-dependent burst size is a general consequence of the intrinsic properties of cerebellar granule cells [50] in combination with the dynamics of granule cell circuitry [86].

A major discovery presented here is the frequency tuning of the PF feedback. Recording of electrosensory afferent input to the EGp has revealed input tuned to high frequency AMs but none to the frequency range important for cancellation [87, 88]. We therefore hypothesize that frequency tuning is generated within the EGp granule cell network itself. The intrinsic dynamics of mammalian granule cells induce resonance in the theta range (3-12Hz [50]) and the local Golgi cell inhibitory network is also associated with low frequency oscillatory activity of granule cells [86]. It has even been suggested that long-term plasticity in the cerebellar granule cell layer can induce “temporal matching” appropriate to the timing function of granule cells [89, 50, 90]. The temporal precision of granule cell discharge has been proposed to be important for the adaptive regulation

of oscillatory motor output [89] and may even be implicated in the frequency tuning of the vestibulo-ocular reflex (VOR) [85]. Frequency specificity of the VOR is evident in the motor output, but there is no experimental evidence pinpointing the neurons and circuits responsible for this frequency tuning; and existing models of such tuning do not explicitly posit the granule/Golgi cells as its source.

In contrast, our results suggest that frequency tuning originates in the cerebellar granule cell network. Furthermore, apart from the initial observation from Lisberger *et al.* [85], the functional consequences of frequency tuning of cerebellar circuitry remains an unexplored issue. Therefore, a conceptually novel aspect of our study is to relate the frequency tuning capacity of granular cell networks (as suggested, for example, in [86]) to the frequency specificity of cerebellar learning by showing how frequency channels in parallel fibers are involved in learning at the PF-SP synapse in our system. The EGp is an inverted cerebellar region [47] with its granule cells directly accessible to patch recording [77], optical Ca^{2+} imaging or experimental perturbation. The EGp may therefore be a convenient preparation to investigate the cellular and circuit mechanisms that generate frequency tuned and precisely timed granule cell/parallel fiber output.

There have been two general approaches to cerebellar function. One emphasizes the function of PFs in a timing role for cerebellum [91, 32] while the other describes the cerebellum as an adaptive filter [5]. It is clear from our results that these viewpoints are not contradictory. The PF feedback to ELL is presumably phase-locked to low frequency sensory input. We hypothesize that the synaptic and intrinsic dynamics of SP cells then make use of this precisely timed input to implement an adaptive filter that cancels redundant low frequency signals. It appears that exactly the same principle may be operative in another well-studied cerebellar-like structure - the ELL of mormyrid fish. The EOD of mormyrids consists of brief pulses with long and variable inter-pulse intervals and the EOD command nucleus sends corollary discharge to the EGp that, as in *A. leptorhynchus*, provides feedback to the ELL. Roberts and colleagues have demonstrated that this circuit is also used as an adaptive filter to cancel redundant sensory input [40, 3, 4] and theoretical studies have implicated the mormyrid PFs in conveying the precise timing of the EOD corollary discharge to ELL neurons [40]. The PF-SP cell synapse analog in the mormyrid ELL obey an anti-Hebbian spike timing dependent plasticity (STDP) rule [20] instead of correlative burst-induced rules. Presumably the mormyrid STDP rule is matched to the pulsatile nature of the EOD that simplifies the cancellation problem to one of causal matching of the time delay of sensory input following a corollary discharge.

Instead, apteronotids operate their EOD continuously, and, for cancellation to occur, feedback to the ELL must synchronize directly with global stimuli of arbitrary frequency. Furthermore, these fish cannot rely on a corollary discharge input to guide the timing of the feedback inputs. We propose that the need to cancel multiple frequencies requires independent frequency channels and multiple learning rules. Our results indicate that the precise timing of sensory or motor associated signals is a general function of parallel fibers of the cerebellum and related structures such as the ELL or the dorsal cochlear nucleus [92, 3, 93]. The local circuitry and dynamics (synaptic and intrinsic) of the PF target varies, however, so as to implement the appropriate adaptive filter required for each system.

In summary, our study brings two new concepts to the theory of cerebellar function. First a precise matching between the temporal window of plasticity and the burst dynamics can permit adaptive shaping of PF inputs by synaptic plasticity. Second, and most importantly, cerebellar granule cells are capable of providing frequency filters of incoming mossy fiber input. We suggest that granule cell frequency tuning may prove central to a better understanding of cerebellar circuits in general.

5.7 Figures and Legends

Chapter 6

Analytics

This chapter explores the behaviour of the model and how it could be predicted analytically. Local and global stimulation are investigated individually since different techniques must be utilized in each scenario. Under local stimulation, the model is similar to a stochastic leaky integrate-and-fire neuron, and the first-passage time (FPT) approach to its firing rate can be employed. However, the model neuron in this thesis is not as simple as the LIF archetype and the FPT analytical solution must be expanded.

During global stimulation, the weight distribution and the burst rates create a negative feedback loop that can be solved self-consistently. This leads to a derivable equilibrium condition if the input current is not modulated and to a power law dependence of some of the average output statistics with the global parameters if the input is modulated. Finally, the model is analyzed to determine if it exhibits a biologically realistic form of contrast invariance (*i.e.* cancellation is maintained as the relative intensity of the stimulus is varied) during global stimulation. For all plots, error bars of the model data were generated and included, but, due to their small size, often appear as dots in the figures.

6.1 Local Stimulation

6.1.1 Behaviour

To gain some insight into the behaviour of the model during local stimulation, Figure 6.1 illustrates the effect of changing the parameters of the local model on two statistics of the output: the firing and burst rates. The parameters are varied by multiples of the optimal local parameters identified in Chapter 3 and clearly show the inherent non-

linearity between the inputs to the model and the output firing rate. Although the details of all of these relationships are not investigated here, note that varying the cut-off frequency of the low-pass filtered noise has a minimal effect on the model (blue line in Figure 6.1A). This is the reason f_{cut} was not systematically optimized during parameter fitting.




Figure 6.1: General behaviour of the model under local stimulation. The effect of varying each model parameter individually on the firing rate while forced at a 4 Hz AM frequency is shown in **A**. For each line, only one parameter is changing, as identified by the legend, and the abscissa identifies the change in the parameter from its initial value. The initial parameter values are $\kappa(4) = 0.39$, $\sigma = 0.76$, $I = 0.58$, $f_{cut} = 500$ Hz, and $\tau_m = 7$ ms and this set is plotted at parameter multiple = 1. The effect of changing the AM frequency on the firing and burst rates is plotted in **B**, where $\kappa(f) = 0.39$ for all frequencies and the burst rates are calculated independently. The relationship between the firing rate and 2-spike (green) and 4-spike (red) burst rates are shown in **C** when there is no AM frequency (solid lines) and when there is a 4 Hz AM frequency and $\kappa(4) = 0.39$ (dashed lines). The firing rate was varied in **C** by changing the bias current, I , but all other parameters are equal to the initial parameter set in **A**.

Despite a constant mean input, the firing rate of the model decreases as the frequency of the input increases (Figure 6.1B). This is due to the leaky integrator nature of the stochastic LIF model, which essentially changes the voltage, V , proportional to the mean of the past input, averaged over a period of time proportional to τ_m . Since I is less than one, the model is in a subthreshold regime and requires noise to bring the voltage to threshold. Even at the peak of the sinusoidal stimulus, $\kappa(f) + I$ is always less than one, so noise is necessary to cause the neuron to fire. As τ_m increases, more of the AM input and the noise is being averaged over, which decreases the fluctuations of the noise and decreases the mean firing rate (purple line in Figure 6.1A). Similarly, increasing the AM frequency with a constant τ_m causes the same affect (Figure 6.1B).

Since the burst rates are proportional to the firing rate, both burst rates also decrease with increasing AM frequencies (Figure 6.1B). Unlike in Figure 4.2, the burst rates

decline monotonically with the AM frequency in Figure 6.1B because the 2- and 4-spike burst rates were calculated independently. This removes the condition of mutual exclusivity of spikes between the two burst types. The definition of a burst is based on a threshold mechanism, so the interaction between the burst rates and firing rate is inherently non-linear. This is evident in the burst rate to firing rate curve (Figure 6.1C), but, interestingly, the existence of a forcing frequency appears to linearize this relationship.

6.1.2 Theoretical Analysis

The goal of local stimulation analysis is to predict the firing rate modulation over the period of the AM frequency given the parameter values of the model. As with predicting the firing rate of any leaky integrate-and-fire model neuron, the problem can be recast as a first passage time problem. However, to account for rectification and low-pass filtering, the mean and variance of the noise must be slightly altered. In parallel to the model's initial construction (Chapter 3), each feature of the model will be studied and added sequentially to generate the complete approximation. The depolarizing after-potential (DAP) will be treated separately at the end, as it is a source of non-linear feedback that severely complicates the analysis.

Gaussian noise

The ELL model is based upon the stochastic leaky integrate-and-fire model, whose canonical equation is

$$\frac{dV}{dt} = -V + I + \sigma\xi(t), \quad (6.1)$$

with a reset voltage, V_r , threshold voltage, V_{th} , $\xi(t)$ is Gaussian white noise, and time has been non-dimensionalized in units of the membrane time constant, τ_m . There is also an absolute refractory period, τ_r , in non-dimensionalized time. The firing rate, R , of this model is a first passage time (FPT) problem whose solution in real time is (in Hz) [60]:

$$R = \left[\tau_m \tau_r + \tau_m \int_{\frac{I-V_{th}}{\sigma}}^{\frac{I-V_r}{\sigma}} e^{x^2} \operatorname{erfc}(x) dx \right]^{-1}, \quad (6.2)$$

where $\operatorname{erfc}(x)$ is the complementary error function, $\frac{2}{\sqrt{\pi}} \int_x^\infty e^{-t^2} dt$. The effects on the firing rate due to changes in the mean or variance of the input are accurately predicted

by this equation, given a small enough time-step (Figure 6.2). This is the fundamental equation that the following theoretical analysis will be built upon.




Figure 6.2: Comparison of the model's firing rate with only Gaussian noise and an input bias to the theoretical firing rate, Equation 6.2, for different time step sizes as either (A) the input bias or (B) the standard deviation of the noise is varied. In A, σ is fixed at 0.5, while in B, I is fixed at 0.5. As the time step decreases, the model gets closer to the theoretical rate. Red line was simulated with $\Delta_s = 0.1$, the green line was simulated with $\Delta_s = 0.01$, and the blue line was simulated with $\Delta_s = 0.001$. Black line (almost superimposed on the blue line) is the theoretical firing rate using Equation 6.2. τ_m is 7 ms.

Low-pass filtered noise

Given that low-pass filtered noise must be used in the model, the firing rate of the neuron with low-pass filtered noise but without rectification was investigated analytically. Thus, the governing equation is $\tau_m \dot{V} = -V + I + \sigma \xi_L(t)$, where $\xi_L(t)$ is low-pass filtered noise, and the forward Euler method can be used to simulate it (see Figure 3.3). However, Equation 6.2 is for the firing rate of a stochastic LIF neuron and does not hold for a LIF with filtered noise input (Figure 6.3).

To model low-pass filtered noise, the variance of the noise in Equation 6.2 must be changed. If σ^2 is the variance for Gaussian white noise, then the variance of the low-pass filtered noise is

$$\sigma_{low}^2 = \sigma^2 \frac{f_{nyq}}{f_{cut}} \Delta_s = \frac{\sigma^2}{2\tau_m f_{cut}} \quad (6.3)$$

where the $\frac{f_{nyq}}{f_{cut}}$ term arises from renormalizing the variance of the low-pass filtered noise and Δ_s arises from using the forward Euler method to numerically simulate the equations instead of the Euler-Maruyama method. By changing the variance in Equation 6.2 to match Equation 6.3, the analytical firing rate is close to the firing rates of the model

when I is varied, but deteriorates when σ is altered (Figure 6.3).




Figure 6.3: Comparison of the model firing rates with the analytical firing rates when the LIF is driven by low-pass filtered noise ($f_{cut} = 500$ Hz) as (A) the mean input current, I , and (B) the standard deviation of the noise, σ , is varied. The governing model equation is $\tau_m \dot{V} = -V + I + \sigma \xi_L(t)$. Analytical firing rates are calculated using Equation 6.2 with the variance of the noise equal to σ^2 (solid black lines) or to $\frac{\sigma^2}{2\tau_m f_{cut}}$ (dashed black lines). Model simulations were performed with 3 different time steps, 0.1, 0.01, and 0.001, corresponding to the red, green, and blue lines, respectively, but the convergence of the simulation for different time-steps leads to overlapping coloured lines. In A, σ is held constant at 0.5, and in B, I is held constant at 0.5. In this and all remaining figures, $\Delta_s = 0.01$ and $\tau_m = 7$ ms.

Inspection of Figure 6.3A suggests that the analytical firing rate calculated using Equations 6.2 and 6.3 is misaligned from the model's data by a fixed input current. Furthermore, the progressive deterioration of the fit between theoretical and model firing rates as σ is varied hints that the input correction is a function of σ . Investigations with different cut-off frequencies as well lead to the phenomenological mean input correction of

$$I_{low} = I - \frac{50\sigma}{f_{cut}}, \quad (6.4)$$

which, together with Equations 6.3 and 6.2, accurately predict the firing rate despite variations in the input current, standard deviation of the noise or cut-off frequencies (Figure 6.4). This is an *ad hoc* solution of the effect of low-pass filtered noise input to an LIF model, and the only evidence that Equations 6.3 and 6.4 are appropriate is the success of the estimate (Figure 6.4). In addition, the mean input correction is to first order, and the deterioration between the theoretical firing rate and the model's firing rate at very high σ and low f_{cut} is likely due to higher order terms in the correction that are not taken into account.




Figure 6.4: Comparison of the model firing rates with the analytical firing rates as **(A)** the mean input current, I , and **(B)** the standard deviation of the noise, σ , is varied for different cut-off frequencies. The governing model equation is $\tau_m \dot{V} = -V + I + \sigma \xi_L(t)$. Blue, green, and red solid lines correspond to model simulations at cut-off frequencies of 250 Hz, 500 Hz, and 1000 Hz, respectively, while broken lines are theoretical estimations using Equations 6.2, 6.3, and 6.4 with dashed, dotted, and dot-dashed corresponding to approximating the model's firing rate for cut-off frequencies of 250 Hz, 500 Hz, and 1000 Hz, respectively. In **A**, σ is held constant at 0.5, and in **B**, I is held constant at 0.5.

Rectification

As described in Section 3.4.1, rectifying the input of a stochastic integrate-and-fire neuron with Gaussian white noise causes the time-step of numerical integration to affect the firing rate. Therefore, Gaussian white noise must be low-pass filtered first. Furthermore, rectification complicates the analytical approximation since the input distribution ceases to be Gaussian around the mean input, I , but becomes a mixed distribution that has a continuous Gaussian probability density for values greater than 0, and a discrete probability mass at 0. In other words, if Y is the feedforward input to the model, then the probability density of Y is

$$f(y) = \begin{cases} 0 & y < 0 \\ \delta(y) \int_{-\infty}^0 \frac{1}{\sqrt{2\pi}\sigma} e^{-(y-I)^2/2\sigma^2} dy & y = 0 \\ \frac{1}{\sqrt{2\pi}\sigma} e^{-(y-I)^2/2\sigma^2} & y > 0 \end{cases} \quad (6.5)$$

where $\delta(y)$ is the Dirac delta function.

In general, the first passage time derivation of the firing rate that produced Equation 6.2 will not hold when the noise is rectified. However, if the probability mass at $y = 0$, is small enough, then Equation 6.2 should still approximate the firing rate. Certainly when the mean is large compared to the variance, then the distribution of Y will be far from 0 and retain most of its Gaussian character. For intermediate values of the mean

and variance when there is a non-negligible probability mass at $y = 0$, then Equation 6.2 could be used if the distribution of the input is still assumed to be Gaussian but with the mean and variance of Y as parameters instead of I and σ^2 . The mean value of Y is

$$\begin{aligned} E\{Y\} &= \int_{-\infty}^{\infty} yf(y)dy \\ &= 0 + \int_0^{\infty} y \frac{1}{\sqrt{2\pi}\sigma} e^{-(y-I)^2/2\sigma^2} dy \\ &= \frac{I}{2} \operatorname{erfc}\left(\frac{-I}{\sqrt{2}\sigma}\right) + \frac{\sigma}{\sqrt{2\pi}} e^{-\frac{I^2}{2\sigma^2}} \end{aligned}$$

where $E\{\}$ is the expectation operator. The expected value of Y^2 is

$$E\{Y^2\} = 0 + \int_0^{\infty} y^2 \frac{1}{\sqrt{2\pi}\sigma} e^{-(y-I)^2/2\sigma^2} dy$$

If we make the change of variables $x = \frac{y-I}{\sqrt{2}\sigma}$ then the integral becomes

$$\begin{aligned} E\{Y^2\} &= \frac{1}{\sqrt{\pi}} \int_{\frac{-I}{\sqrt{2}\sigma}}^{\infty} (I + \sqrt{2}\sigma x)^2 e^{-x^2} dx \\ &= \frac{1}{2}(I^2 + \sigma^2) \operatorname{erfc}\left(\frac{-I}{\sqrt{2}\sigma}\right) + \frac{\sigma I}{\sqrt{2\pi}} e^{-\frac{I^2}{2\sigma^2}} \end{aligned}$$

and so $Var\{Y\} = E\{Y^2\} - E\{Y\}^2$ can be calculated. Using the mean and standard deviation of Y as I and σ and then applying the low-pass filtered noise corrections produces the analytical estimate of rectified input to an LIF model with low-pass filtered noise:

$$I_{rect} = E\{Y\} - \frac{50\sqrt{Var\{Y\}}}{f_{cut}} \quad (6.6)$$

$$\sigma_{rect}^2 = \frac{Var\{Y\}}{2\tau_m f_{cut}} \quad (6.7)$$

This provides a close fit to the model data (Figure 6.5). For a small noise intensity,

neglecting the rectification of the input is acceptable (blue line in Figure 6.5A), but, as the variance increases, the analytical fit without accounting for rectification deteriorates. At high mean inputs, the fits tend to improve since the distribution of the input is less influenced by rectification. The analytical approximation including rectification of the input provides a much better fit to the numerical simulations, although small differences do develop at low mean input values, where the assumption that the input still has a Gaussian distribution no longer holds (Figure 6.5B, red and green lines).




Figure 6.5: Comparison of model firing rates with analytical firing rates as the mean input current, I , is varied for different values of σ . The governing model equation is $\tau_m \dot{V} = -V + [I + \sigma \xi_L(t)]$ (*i.e.* rectification is applied). Blue, green, and red solid lines correspond to model simulations with $\sigma = 0.5$, 1, and 1.5, respectively. Broken lines are analytical approximations, with dashed, dotted, and dash-dot corresponding to $\sigma = 0.5$, 1, and 1.5, respectively. In **A**, the analytical approximation is based on the low-pass filtered correction approach (Equations 6.2, 6.3, and 6.4) and assumes no change in the firing rate due to rectification. In **B**, the analytical approximation is based on the mean and variance correction of the rectified input (Equations 6.6 and 6.7) with Equation 6.2. In both cases, $f_{cut} = 500$ Hz.

Sinusoidal input

During local stimulation, the mean input current, I , is sinusoidally modulated by $\kappa(f)$ at the AM frequency f . Since the value of σ used in the model to mimic pyramidal neurons in the ELL is 0.759 and the mean input is less than one, the firing rate curve is between the blue and green lines in Figure 6.5B, which necessitates the analytical rectification approximation.

To track the modulation of the firing rate over one AM cycle, the period is divided into 20 bins. These bins are small enough such that the sinusoidal input current can be considered approximately constant within each bin, and it was assumed that each bin can be analyzed independently. This approximation is aided by the separation in timescale

between the membrane time constant and the AM period, the latter being slower. For each bin, the mean input is $I + \kappa(f) \sin(2\pi f t_i^{bin})$, where t_i^{bin} is the time at the centre of the bin i , and Equation 6.2 is used with the corrections for rectification and low-pass noise filtering (Equations 6.6 and 6.7) to find the firing rate. The firing rate in each bin is then plotted against t^{bin} to approximate the PSTH (Figure 6.6).

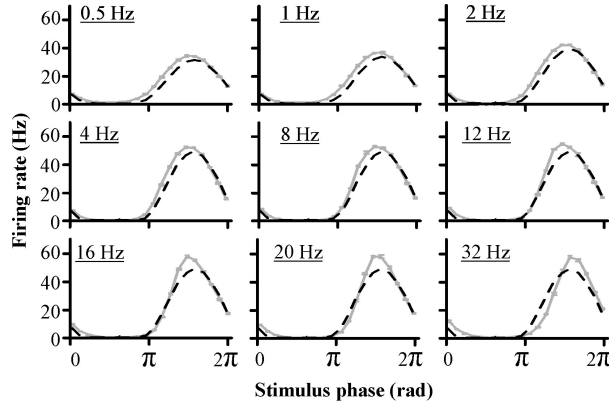


Figure 6.6: Comparison of the model's firing rate modulation (grey lines) and the analytical approximation (dashed black lines) during local stimulation without an active DAP. The governing model equation is $\tau_m \dot{V} = -V + [I + \sigma \xi_L(t) + \kappa(f) \sin(2\pi f t)]$. Contrast was decreased at AM frequency < 4 Hz to account for electroreceptor adaptation: $\kappa(0.5) = 0.25$, $\kappa(1) = 0.27$, $\kappa(2) = 0.31$ and $\kappa(f \geq 4) = 0.39$.

The analytical expression is able to closely match the model's firing rate modulation, although there are some very minor differences, especially on the rising side of the stimulus. This is due to the supposition that the input can be binned and analyzed independently: if the neuron has not fired recently, as in the trough of the input, it has a higher probability to fire when current appears, which falsifies this assumption. Further, while the analytical firing rate is identical for all AM frequencies that have equivalent contrasts, the model's firing rate does change as the AM frequency increases. Interestingly, even though the average firing rate over a period decreases with AM frequency (Figure 6.1B), the maximum instantaneous firing rate over a period increases with AM frequency (compare the analytical firing rate and the model's firing rate in Figure 6.6 at 8 Hz and 20 Hz). The reason for this is unknown.

Depolarizing after-potential

Finally, the depolarizing after-potential can be included in the analytical approximation. Since the size of the DAP input, averaged over time, is proportional to the model's firing

rate, the DAP is a form of positive feedback that is internal to the cell - as opposed to feedback via external circuitry. Incorporating the DAP, the mean input to the model becomes

$$I_{eff} = I + \lambda R(I_{eff}) \quad (6.8)$$

where $R(I_{eff})$ is the firing rate of the model at the new effective input current, I_{eff} , and λ is the strength of the feedback, which has yet to be determined. The firing rate as a function of the input current can be found from Equation 6.2 using the rectified and low-pass filtered corrections (Equations 6.6 and 6.7) without the DAP but the stable firing rate given $\lambda \neq 0$ must be found numerically.

To find the stable firing rate, Equation 6.8 can be transformed into a recursive map, *i.e.* $R(I_{i+1}) = R(I + \lambda R(I_i))$, or, defining R_{i+1} as $R(I_{i+1})$, then $R_{i+1} = R(I + \lambda R_i)$. If this map is stable, then after sequential iterations the firing rate should be constant: $R_{i+1} \approx R_i$. If the firing rate was linear, *i.e.* $R(I_i) = mI_i$, where m is a parameter, then the recursive map becomes $R_{i+1} = mI + m\lambda R_i$, whose stable value of I is

$$I_s = \frac{I}{1 - \lambda m}. \quad (6.9)$$

Clearly, as λ approaches $1/m$, I_s goes to infinity. Thus, for strong positive feedback, the system is unstable, regardless of the initial parameters. As the firing rate does not escape to infinity, the DAP must be weak in this sense. Extending this analysis to the non-linear case of the stochastic LIF, the slope of the recursive map is still the important parameter for stability. Since the firing rate is monotonically increasing (Figure 6.5), the slope is always greater than zero, and, in its linear region, the slope is approximately 100 Hz per unit of input current, which implies that only very small feedback strengths will be stable or the firing rate will escape to infinity.

This effect does not actually occur in the stochastic LIF-DAP model or in reality because a single feedback parameter λ is an oversimplification of the DAP dynamics. In particular, if the neuron fires too quickly then the dendrites may still be in their dendritic refractory period, which would prevent the DAP from initiating. This limits the feedback of the DAP at high firing rates.

To model the effect of the DAP analytically, the strength of its feedback must be known. For simplicity, a constant feedback strength, λ_{DAP} , was assumed to be a good fit for the DAP dynamics within the range of input currents for local stimulation. The value of λ_{DAP} was obtained by matching the firing rate of the model and the stable firing rate of the rectified input approximation when both had a mean input current of

0.97 and a noise standard deviation of 0.759 but no sinusoidal current modulation. This corresponds to the input current at the peak of an AM cycle for frequencies greater than 2 Hz (since $I + \kappa(f \geq 4) = 0.58 + 0.39 = 0.97$) where the DAP has the largest effect. A feedback strength of $\lambda_{DAP} = 0.0012$ was found to be optimal. This is indeed less than the reciprocal of the slope of the firing rate curve, indicating its stability. To recreate the local PSTH, the input current was binned as before, but now Equation 6.8 is employed and the firing rate and input current are re-iterated for each bin until a stable rate is reached (Figure 6.7).




Figure 6.7: Comparison of the model’s firing rate modulation (grey lines) and the analytical approximation (dashed black lines) during local stimulation when the DAP is active. The governing model equation is $\tau_m \dot{V} = -V + [I + \sigma \xi_L(t) + \kappa(f) \sin(2\pi ft)] + DAP(t)$. Contrast was decreased at AM frequency < 4 Hz to account for electroreceptor adaptation. The model PSTH in grey is identical to the PSTH of the optimal local fit (see Figure 4.2).

A comparison of the model PSTH in Figure 6.6 and 6.7 show that the DAP increases the firing rate by about 10 Hz at the peak of the input, which was a critical feature during local parameter fitting. Thus, the theoretical approximation of the firing rate was able to accurately reproduce the firing rate modulation of the model during local stimulation.

6.2 Global Stimulation

6.2.1 Qualitative investigation of \bar{w}

Under global stimulation, the feedback circuit is active and the parallel fibre synaptic weights, w_ϕ , supply input to the neuron, which affects the neuron’s firing rate and burst

rates. The burst rates, in turn, change the weight values through the 2-spike and 4-spike plasticity rules. This establishes a negative feedback loop between the weights and the burst rates, and when the individual weights have reached a stable value, learning has converged and the statistics of the model are stable. The following equations govern the model during global stimulation:

$$\tau_m \frac{dV}{dt} = -V + [I + \sigma \xi_L(t) + \kappa(f) \sin(2\pi ft)] + DAP(t) + \Lambda(w_\phi - gV) \quad (6.10)$$

$$\text{if } |t_\phi - t_B| < L_{w_{2,4}} \rightarrow w_\phi = w_\phi + w_\phi \eta_{2,4} \left[\left(\frac{t_\phi - t_B}{L_{width_{2,4}}} \right)^2 - 1 \right] \quad (6.11)$$

$$\tau_w \frac{dw_\phi}{dt} = w_{max} - w_\phi \quad (6.12)$$

Due to the complex relationship between the 2- and 4-spike burst rates and the weight distribution of the model, the feedback system is nonlinear. Nevertheless, the burst rates increases monotonically with input current, and the weights decrease monotonically with increasing bursts (since the rule is purely depressing), so a single stable fixed point is expected. To obtain some insight into the behaviour of the weights under global stimulation, the average, \bar{w} , of all the weight values (each one being associated with one phase of the sinusoidal modulation) was tracked for 1750 s when the system was forced by a 4 Hz AM frequency (Figure 6.8).

The average weight value converges quickly to near its equilibrium value but, due to noise in the neuron model, continues to fluctuate around it (Figure 6.8A). At higher resolutions, \bar{w} appears to follow a near random walk around its equilibrium value, albeit with different magnitudes for an increment or decrement (Figure 6.8B and C). \bar{w} slowly increases via the potentiating rule, which is a continuous function of time, and sporadically decreases via the depression rules whenever the cell bursts, which are discrete, random events.

Although the timing of each burst is stochastic, the effect of each superficial cell burst on \bar{w} is deterministic and can be calculated using the burst learning rules (Equation 3.14; also see Figure 4.3). The update rule for weights is multiplicative, and consequently each weight at each phase will change by a different amount, depending on the weight's initial value. However, if the spread of the weight distribution is small, \bar{w} can be substituted for the weight at each phase and the impact of one burst on \bar{w} will depend on the integral of the burst rule:

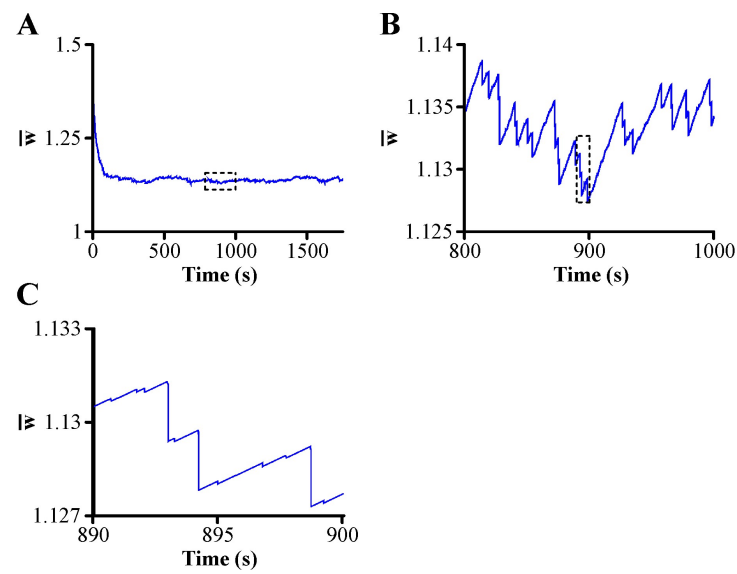


Figure 6.8: Typical simulation of the average weight value during a single run of global simulation with a 4 Hz AM when all weights are initialized at 1.5. In **A**, the average weight value is tracked for 1750 s. In **B**, the segment between 800-1000 s is examined in finer detail, and finally in **C**, the period between 880-900 s is analyzed. In **C**, large step decreases are from 4-spike bursts, while small decreases are from 2-spike bursts. Boxes in **A** and **B** specify the area to be examined in the following panel. Parameters are as in the global η model.

$$\begin{aligned}
\text{Impact} &= \int_{|t_\phi - t_B| < L_w} w_\phi \eta \left[\left(\frac{t_\phi - t_B}{L_{w_{2,4}}} \right)^2 - 1 \right] dt_\phi \\
\text{Impact} &\approx \bar{w} \eta \int_{-L_w}^{L_w} \left(\frac{t^2}{L_w^2} - 1 \right) dt \\
\text{Impact} &\approx \bar{w} \eta \frac{4}{3} L_w,
\end{aligned}$$

where η_4 and L_{w_4} or η_2 and L_{w_2} would be substituted for η and L_w to determine the impact of a 4-spike or 2-spike burst, respectively.

Since the weights are not continuous but discretized, the sum of the weights will decrease by approximately this amount every time a superficial cell burst occurs, in a Riemann sum estimation of the integral. The time each segment (or weight) is active is the size of the partition of the Riemann sum, which is 2.5 ms or $T/n_w(f)$, where T is the period of the AM frequency and $n_w(f)$ is the number of feedback segments, which is also the number of weights, at that frequency. Thus, the change in the sum of the weights is equal to

$$\begin{aligned}
\text{Impact} &\approx \Delta \left(\sum_\phi w_\phi \right) \frac{T}{n_w(f)} \\
\Rightarrow \Delta \left(\sum_\phi w_\phi \right) &\approx \frac{4\bar{w}\eta L_w n_w(f)}{3T}.
\end{aligned}$$

The change of the average weight value, \bar{w} , is simply the change in the sum divided by the total number of weights, $n_w(f)$:

$$\Delta \bar{w} \approx \frac{4\bar{w}\eta L_w}{3T} \tag{6.13}$$

For a 4 Hz AM frequency and global η model parameters, this corresponds to a decrease of $0.00192\bar{w}$ for 4-spike bursts, and $9.6 \times 10^{-5}\bar{w}$ for 2-spike bursts, which is small compared to \bar{w} . Since the 2-spike burst are more common than the 4-spike bursts, the weights appear to follow a random walk with potentiation and 2-spike burst depression almost in balance, with an occasional sharp correction by the 4-spike burst. Note that

for a given burst rule, as the AM frequency increases, the relative effect of each burst on the average weight value increases. Although it was derived when the difference in weight values at different phases is small, Equation 6.13 should still approximately hold when the weights have a non-negligible variation. This will be critical in understanding the inverse relationship between the burst rates and the AM frequency under global stimulation.

6.2.2 Qualitative investigation of the weight distribution

The dynamics of the weight distribution were partially analyzed in Section 3.4.2 to ascertain the time until the system is in equilibrium. To investigate the weight dynamics more fully, a weight at the peak and at the trough of the input as well as the \bar{w} were tracked during the simulation of the global η model (Figure 6.9). For comparison, Figure 3.4 is also plotted, where the same global η model is used all frequencies have identical values of κ , η_2 and η_4 values (see Figure 6.10).

Comparing these two paradigms reveals the impact of the learning rate and the AM frequency on the weight behaviour. Decreasing the learning rate reduces the impact of each burst; hence, more bursts are required to maintain the equilibrium weight value. This decreases the variance of the final value of each weight, illustrated by the “thickness” of the lines in Figure 6.9 and 6.10 (*i.e.* compare 32 Hz). If η is kept constant, then the variance of the traces increases with AM frequency, as bursts become rarer (see Figure 6.1) and more powerful (see Equation 6.13) at higher frequencies (*e.g.* compare the thickness of blue traces at different frequencies in Figure 6.10). Further, reducing the learning rate decreases the difference between final weight values at different phases (*i.e.* between red and green lines at high AMs in Figure 6.9) because the impact of a relative difference in burst rates at different phases of the input is diminished. This leads to a more constant global feedback input, and to poorer cancellation as seen in the global η model.

Comparing the two figures at low AM frequencies also clarifies the effect of changing κ on the weight dynamics, as Figure 6.9 includes fitting for electroreceptor adaptation while Figure 6.10 does not. Since κ is the modulation of the input that the feedback must cancel, decreasing κ reduces the difference in weight values because the feedback needs a smaller modulation to cancel a smaller input.




Figure 6.9: Weight evolution during global stimulation in the final global model. Three different initial conditions (all weights initialized at 0.5, all weights initialized at 1.0, and all weights initialized at 1.5) were each run 5 times and the value of a weight at the peak of the input ($\phi = 3\pi/4$, green lines), the trough of the input ($\phi = \pi/4$, red lines), and the average weight value (blue lines) were tracked for 7000 s. The global η model is used in these simulations: at low frequencies, $\kappa(f)$ is decreased to account for low-frequency electroreceptor adaptation and at high frequencies η_2 and η_4 are decreased to approximate the decreasing burst rate of granule cells (see also Figure 4.7).

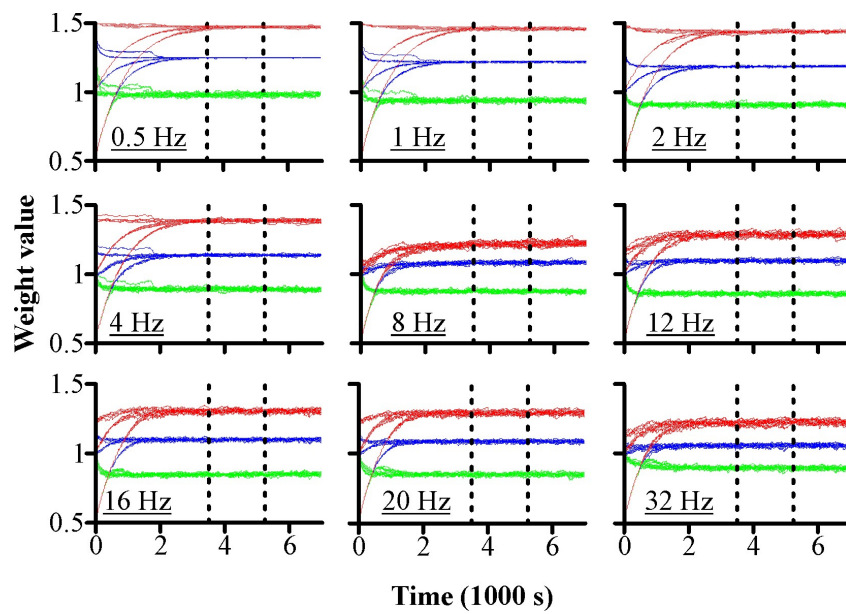


Figure 6.10: Weight behaviour during stimulation of the global η model when $\kappa(f)$, η_2 and η_4 are constant. This is the same plot as Figure 3.4. Three different initial conditions (all weights initialized at 0.5, all weights initialized at 1.0, and all weights initialized at 1.5) were each run 5 times and the value of a weight at the peak of the input ($\phi = 3\pi/4$, green lines), the trough of the input ($\phi = \pi/4$, red lines), and the average weight value (blue lines) were tracked for 7000s.

6.2.3 Theoretical Analysis

The goal of this theoretical analysis is to predict the firing rate, the burst rates and the weight distribution given the initial conditions of the systems (the parameter values and the AM frequency) under global stimulation. This goal should be possible since the burst rates affect the weight distribution through the learning rules and the weights affect the burst rates by providing input to the model neuron. After transients have decayed, the final burst rates and weight distribution of the system must, therefore, be at the equilibrium position where these two effects balance. For simplicity, only 2-spike bursts will depress the weights in this analysis so a single burst rate, Br , can be defined, and the global model parameters used here are identical to the optimal 2-spike burst model in Section 4.2 (Figure 4.5). η is also constant across AM frequencies.

During global stimulation, the equilibrium weight distribution is a collection of the stable weight values at each phase of the AM cycle. These final weight values also influence the firing and burst rates that occur at that phase. In a similar approach to the local theoretical analysis of the PSTH, each AM cycle is divided into 20 bins. These bins are only used for theoretical analysis, as the model still has $n_w(f)$ weights at each frequency. Within each bin, the average weight value, \bar{w}_i^{bin} can be calculated by assuming the input stimulus is constant inside the partition and that the bins are independent of each other. Therefore, the equation for each bin is

$$\tau_m \frac{dV}{dt} = -V + [I + \sigma \xi_L(t) + \kappa(f) \sin(2\pi f t_i^{bin})] + DAP(t) + \Lambda(w_i^{bin} - gV), \quad (6.14)$$

where t_i^{bin} is the time in the middle of that particular bin. Also, since the input within each bin is not modulated, all of the weights, w_i^{bin} , will converge to the same value and only \bar{w}_i^{bin} is important. Note, however, that the weights are still periodic with the AM frequency even within a bin.

The weights and the burst rules complete a feedback loop in the model. To formulate the analytical approximation to global stimulation, the firing rate within a single bin will be investigated when stimulated by a fixed input current (*i.e.* a fixed weight value) with disynaptic inhibition, g . Given the firing rate within a bin, the 2-spike burst rate can be obtained, which specifies the average weight value, \bar{w}_i^{bin} , that will be stable at that input current. In equilibrium, the input above the bin's feedforward input current must be equal to \bar{w}_i^{bin} , so these equations form a system that can be solved self-consistently.

After numerically finding the solution in a single bin, the approach will be applied to

all 20 bins that each represent a different phase of the input and, hence, have a different bias current (different t_i^{bin}). This will lead to different equilibrium values of the average weight and the firing and burst rates in each bin, which is used to construct the analytical distribution of the weights, the PSTH, and the average 2-spike burst rate over one AM cycle. For simplicity, the weight input was assumed to be part of the feedforward input during the analytical calculations, even though the weights are not actually rectified with the electroreceptor input in the model.

Firing rate in each bin

For a given constant weight value, the firing rate within each bin can be approximated using the same technique that was used during local stimulation. However, Equation 6.14 is not in the appropriate form for Equation 6.2 since disynaptic inhibition is now involved. The $-\Lambda gV$ term must be dealt with by dividing all terms by $1 + \Lambda g$, which changes the equation to

$$\frac{\tau_m}{1 + \Lambda g} \frac{dV}{dt} = \frac{-(1 + \Lambda g)V + [I + \sigma \xi_L(t) + \kappa(f) \sin(2\pi f t_i^{bin})] + DAP(t) + \Lambda (w_i^{bin})}{1 + \Lambda g}. \quad (6.15)$$

The membrane time constant, the bias current, and the standard deviation of the noise can then be renormalized to include the $1 + \Lambda g$ term and the firing rate can be predicted using Equations 6.2, 6.6, and 6.7. Alternatively, the mean and variance can remain fixed as long as the voltage threshold is multiplied by $1 + \Lambda g$. If the DAP is deactivated and the weights are fixed at zero, this method provides a good fit of the model during constant input (Figure 6.11).

The disynaptic inhibition is a form of shunting inhibition, which leads to subtractive gain control, a result found elsewhere [94], and the theoretical approximation matches the model's firing rate very well.

After the increased inhibition is incorporated into the theoretical analysis, the DAP may be approximated using the recursive map approach shown previously. However, the DAP is weaker now that its strength has been reduced by $1 + g$ (Equation 6.15). Furthermore, the instantaneous firing rate during global stimulation is never as high as it is during local stimulation, which limits the impact of the DAP. Because of its minor effect, the DAP was excluded in the theoretical estimation of global stimulation.




Figure 6.11: R-I curve of the model with different values of inhibition. This is an example of the effect of inhibition on the firing rate within one bin given a constant input. The model's governing equation is $\tau_m \dot{V} = -(1 + \Lambda g)V + [I + \sigma \xi_L(t)]$. Blue, green, and red lines correspond to the model's firing rate with $g = 0, 0.5,$ and 1 as the bias current, I , is varied from 0 to 2. Dashed black lines are the theoretical firing rates. In this and all subsequent plots in this Section, $\sigma = 0.759, \tau_m = 7$ ms, $\Lambda = 1,$ and $f_{cut} = 500$ Hz.

Burst Rate

Once the firing rate is known, it can be converted into a 2-spike burst rate. Since the definition of a burst involves a threshold non-linearity, the mathematical relation between the burst rate and the firing rate would be difficult to calculate analytically. Nevertheless, the model data of the firing and burst rates can be fitted for a constant input current to approximate the burst function (Figure 6.12). A quadratic fit matches the data well and the resulting function is

$$Br = 0.0071R^2 + 0.055R - 0.22 \quad (6.16)$$

where R is the firing rate and Br is the 2-spike burst rate. The norm of the residuals for this fit was 0.2 ($n = 30$).

\bar{w}^{bin} as a function of the burst rate

Thus, for a given weight value, the 2-spike burst rate can be calculated. The final missing relationship before the weight in each bin can be predicted is how the burst affects the weight value. Observing Figure 6.8, the average weight value quickly decreases until the stochastic burst rate and the deterministic potentiation rule are in equilibrium. Without 4-spike bursts, there would be far less movement around the steady-state \bar{w}^{bin} , and the balance between the depression caused by each burst and the potentiation between bursts would be more exact. Schematically, the weight dynamics might look like Figure 6.13.

Although the calculations in this section are for \bar{w}^{bin} and, therefore, at a constant




Figure 6.12: A plot of the 2-spike burst rate as a function of the firing rate (Br-R curve) for a constant input (no AM frequency). Data points (grey dots) were collected from global stimulation at a constant input with $g = 0.87$ and I ranging from 0.5 to 1.5 to recreate the range of firing rates seen during global stimulation. The data was fitted and a quadratic matches the data well (dashed black line).




Figure 6.13: Schematic of the dynamics of the average weight value at equilibrium. Bursts occur at time t_i and t_{i+1} , which depress the value of the average weight by an amount χ_i and χ_{i+1} , respectively. The weight values then slowly potentiate according to $\tau_w \dot{w} = w_{max} - w$. A difference map can be defined that transforms the average weight just after burst i to the average weight just after burst $i + 1$: $\bar{w}(t_i^+) \rightarrow \bar{w}(t_{i+1}^+)$.

input, the superscript *bin* will be dropped for clarity until the end of the derivation. Thus, the average weight value within a bin at time t is defined as $\bar{w}(t)$ and the depression of $\bar{w}(t)$ after burst i is equal to χ_i . Between bursts, $\bar{w}(t)$ increases via the potentiation rule. In reality, each weight within the distribution will increase via the potentiation rule differently, depending on the value of each weight. However, since χ_i is very small compared to the $\bar{w}(t)$ (see Section 6.2.1) and the input has no modulation other than noise, the difference between weights is negligible and the $\bar{w}(t)$ can be assumed to increase via the potentiation rule. This system can be analyzed as an iterative mapping of weight values just after a burst occurs. If a burst arrives at time t_i , then

$$\bar{w}(t_i^+) = \bar{w}(t_i^-) - \chi_i \quad (6.17)$$

where t_i^- is defined as the time just before the burst, t_i^+ is defined as the time just after the burst, and χ_i is the effect of one burst on the average weight value at time t_i . Between bursts, the weights are only under the influence of the potentiation rule, so $\bar{w}(t_{i+1}^-)$ can be found by solving the homeostatic rule:

$$\begin{aligned} \tau_w \frac{d\bar{w}}{dt} &= w_{max} - \bar{w} \\ \Rightarrow \bar{w}(t) &= w_{max} + (\bar{w}(t_0) - w_{max}) e^{\frac{-(t-t_0)}{\tau_w}} \end{aligned}$$

With the initial weight as $\bar{w}(t_i^+)$, the final weight as $\bar{w}(t_{i+1}^-)$, and the time difference $t_{i+1}^- - t_i^+ = \Delta t$, then $\bar{w}(t_{i+1}^-) = w_{max} + (\bar{w}(t_i^+) - w_{max}) e^{\frac{-\Delta t}{\tau_w}}$ and substituting this into Equation 6.17 gives

$$\bar{w}(t_{i+1}^+) = w_{max} + (\bar{w}(t_i^+) - w_{max}) e^{\frac{-\Delta t}{\tau_w}} - \chi_i \quad (6.18)$$

If this system is in equilibrium, then $\bar{w}(t_{i+1}^+) = \bar{w}(t_i^+) = \bar{w}^{bin}$, which is the equilibrium average weight value within a bin. Technically, as $\bar{w}(t_i^+)$ is the lowest value the average weight will obtain between bursts, the mean (*i.e.* averaged over time) \bar{w}^{bin} is slightly higher than $\bar{w}(t_i^+)$. However, the size of χ_i is much smaller than the average weight value, so the difference is negligible. The time difference Δt is the time between successive bursts during equilibrium, so the reciprocal of Δt is just the equilibrium burst rate, Br . During the actual simulation, the bursts will not arrive exactly Δt apart due to the noise in the system but if the weights stay near their equilibrium value, then Δt should be the mean time between bursts. In addition, χ_i is the effect on the average weight value by one burst, which was calculated in Equation 6.13. At equilibrium, all the weights must be

at \bar{w}^{bin} before the burst since the input is constant. Therefore, $\chi_i = \chi = \bar{w}^{bin} \frac{4\eta_2 L_{w_2}}{3T}$ and solving for \bar{w}^{bin} gives

$$\bar{w}^{bin} = \frac{w_{max} \left(1 - e^{-\frac{1}{\tau_w Br}}\right)}{1 - e^{-\frac{1}{\tau_w Br}} + \frac{4\eta_2 L_{w_2}}{3T}} \quad (6.19)$$

This equation relates the equilibrium burst rate, Br , to the equilibrium weight value, \bar{w}^{bin} during a constant input, and this equation holds given the parameters of the optimal 2-spike burst-rule-only model (Figure 6.14).




Figure 6.14: A plot of \bar{w}^{bin} (*i.e.* the average steady-state weight value in a bin) versus the equilibrium 2-spike burst rates when the feedback is periodic at frequencies of 1 Hz (solid blue lines), 4 Hz (solid green lines), and 16 Hz (solid red lines) at different constant inputs (*i.e.* no sinusoidal forcing). The model equation is $\tau_m \dot{V} = -V + [I + \sigma \xi_L(t)] + DAP(t) + \Lambda (w_i^{bin} - gV)$. The model presented in this figure separates the effect of the weights repeating at different frequencies from the effect of a sinusoidal input, although both occur simultaneously when an AM drives the complete model. I was varied from 0.2 to 1 but otherwise the model has parameters from the 2-spike burst-rule-only model: $\eta_2 = 0.0018$, $L_{w_2} = 10$ ms, $\tau_w = 4900$ s, $w_{max} = 1$, and $g = 0.87$. Also plotted is Equation 6.19 (dashed lines) as it predicts \bar{w}^{bin} from the mean burst rate for each frequency.

Sinusoidal input

Summing all these individual relationships together, we have a set of equations that must be self consistent: the weight value in each bin must create a firing rate (via Equation 6.2) that produces a burst rate (via Equation 6.16) that generates a weight value (via Equation 6.13) that, in equilibrium, is identical to the initial weight value used in Equation 6.2. This can be solved using a root finding algorithm and must be iterated for each of the 20 bins in an AM cycle. The result of this theoretical analysis is depicted in Figures 6.15 and 6.16.

Clearly, the analytical approximation is a good match to the model's behaviour, at

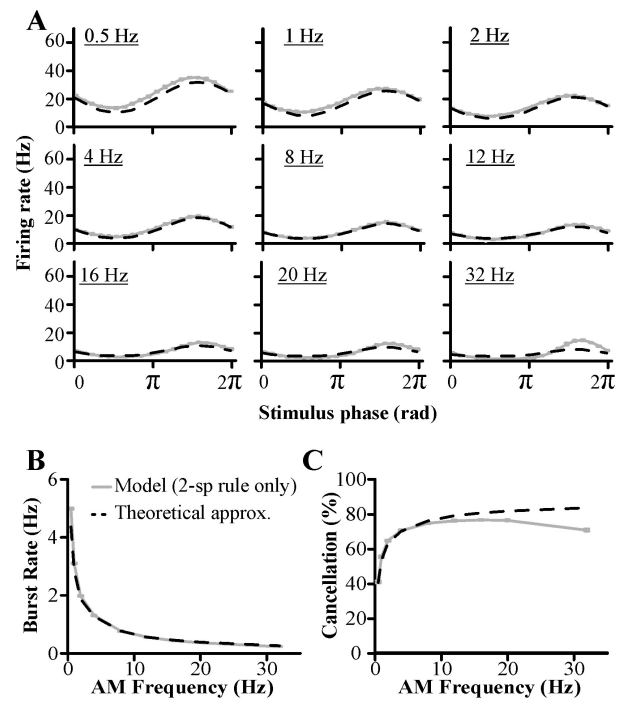


Figure 6.15: Comparison of the (A) PSTH, (B) burst rate, and (C) cancellation of the model (grey solid lines) and the analytical approximation (dashed black lines) during global stimulation with only the 2-spike burst rule active. The analytical approximation was generated by solving Equations 6.2, 6.16, and 6.19 self-consistently within individual segments of the AM period. The parameters are as in the optimal 2-spike burst-rule-only model.

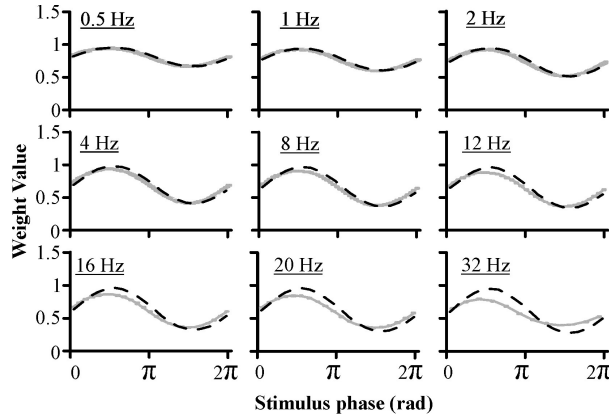


Figure 6.16: Comparison of the mean equilibrium weight distribution between the model (solid grey lines) and the analytical approximation (dashed black lines) for different AM frequencies with the 2-spike burst-rule-only model. The weight at each phase, w_ϕ , was individually tracked during the recording epoch (3500-5250 s) and its mean calculated. The model was simulated 10 times and the standard deviation of the mean of w_ϕ between iterations generated the error bars. Since each weight is active for 2.5 ms and only 1 weight is active at a time, there are fewer weights operational as the period decreases. Thus, there are fewer data points per period at higher frequencies.

least for low frequency AMs. Akin to the local stimulation fit, the global theoretical fit slightly deteriorates at high AMs. With the period always divided into 20 bins in the theoretical analysis, at high frequencies this leads to bins of small time windows, and the memory of the LIF neuron and the size of the burst depression rule will cause the bins to lose their assumed independence. Furthermore, the definition of a burst is 2 spikes within 15 ms, and that time threshold spans multiple theoretical segments at high frequencies. Thus, the generation and effect of a burst can no longer be approximated to be confined to one bin, and the analytical fit degenerates. The dependence of the firing rate at the peak of the input on the AM frequency, observed during the theoretical analysis of local stimulation (Figure 6.7), may also help explain the increased firing rate of the model compared to the analytical approximation at high frequencies during global stimulation.

6.2.4 Power laws

It would also be useful to know if the mean output quantities, averaged over the AM period, could be predicted without partitioning the period and solving the self-consistent equations within each bin. Unfortunately, although Equation 6.19 accurately predicts \bar{w} given the burst rate during a constant input (*i.e.* in a bin), it systematically underesti-

mates \bar{w} when the input is a sinusoid (Figure 6.17).




Figure 6.17: A plot of \bar{w} (averaged over weights at different phases) and the average burst rate, calculated over an AM cycle, for the model during global stimulation with the experimental AM frequencies (solid grey line). Each data point is from a different frequency with 0.5 Hz corresponding to the highest burst rate and 32 Hz corresponding to the lowest burst rate. All parameters are as in the optimal fit for a 2-spike burst-rule-only model (see Figure 4.5). Also plotted is the estimated \bar{w} from the burst rate using Equation 6.19 (dashed black lines), which assumes the modulation of the input does not affect the average weight.

Since Equation 6.19 was successful in predicting the model when the input was not modulated but is imperfect when the input is modulated, the system must have a non-linear dependence on $\kappa(f)$ and the AM frequency even during global stimulation. This is not surprising given the number of non-linear relationships between variables derived in the previous section. What is surprising is that log-log plots reveal a simple power law relationship between the average output statistics and the AM frequency (Figure 6.18).

The burst rate slightly deviates from this trend in part because it is related to the firing rate quadratically (Equation 6.16) and the firing rate does have a power law dependence. Upon further investigation, there are power laws among other global parameters as well:

$$\bar{R}(\eta_2 L_{w_2} \tau_w f)^x = C_1 \quad (6.20)$$

$$\bar{w}(\eta_2 L_{w_2} \tau_w f)^y = C_2, \quad (6.21)$$

where the exponents, x and y , are a nonlinear function of the inhibition. These invariants hold for $f < 20$ Hz, $L_{w_2} < 20$ ms, and $\tau_w < 19600$ s, beyond which power law relationships begin to breakdown (Figure 6.19). These limits affect the size of the fluctuations of the weight values around their equilibrium values, which suggests that, in order for these equations to hold, the fluctuations must be small. The limit on L_w also implies that the addition of 4-spike burst depression could render these equations invalid, which is indeed




Figure 6.18: The power law dependence between the AM frequency and the (A) average firing rate, (B) average burst rate, and (C) average weight value of the 2-spike burst-rule-only model during global stimulation. All parameters are as in the 2-spike burst optimal fit (see Figure 4.5). Averages are taken over the AM cycle. The norm of the residuals for the firing rate fit, the burst rate fit, and the average weight fit are 0.09, 0.2, and 0.004, respectively.

the case (data not shown).

The two power laws likely arise from the interplay between the firing rate and weight values derived previously but with a correction term for sinusoidal input. Combining these two equations gives $\bar{R}\bar{w}^z = C_3$, *i.e.* that the average firing rate depends only on the average weight value. This is interesting since the average weight is not directly dependent on the firing rate and the firing rate does not vary inversely with the input. Unfortunately, the exact origin of these power laws within the global feedback is not fully understood, nor is the precise reason for the breakdown when the system's parameters are sufficiently large.

6.3 Contrast Invariance

Built into the model is the concept that each AM frequency the fish encounters is separately and independently cancelled. This is provided by the separate frequency channels in the parallel fibre feedback pathway. However, the strength, or contrast ($\kappa(f)$ in the model), of any given AM frequency will vary in the electric fish's natural environment as, for example, an apteronotid moves closer or farther away from a conspecific, and it would be unrealistic if every contrast of each AM frequency also required an independent weight set for learning. In order for the model to maintain its biological relevance, the weight

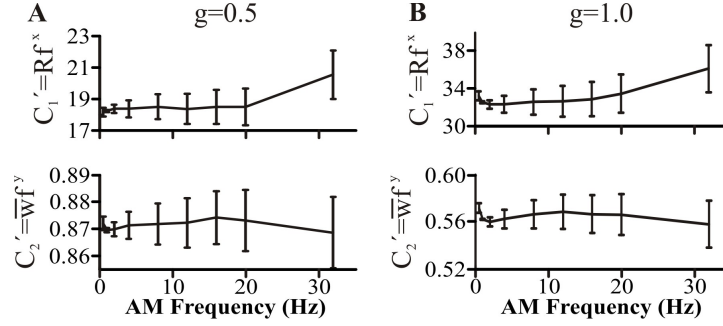


Figure 6.19: A plot of the power laws as a function of frequency for (A) $g = 0.5$ and (B) $g = 1.0$. For simplicity, the parameters in Equations 6.20 and 6.21 that are not varied are included in the constant, hence C'_1 and C'_2 . The parameters are $\eta_2 = 0.0018$, $\tau_w = 9800$ s, $w_{max} = 1.5$, $L_{w_2} = 10$ ms, $\kappa(f) = \kappa = 0.39$ (constant across frequencies), $I = 0.58$, $\sigma = 0.759$, $\tau_m = 7$ ms, and $f_{cut} = 500$ Hz. The error bars were generated by running the model 10 times at each frequency and using error propagation to determine the error on the invariants. In A, the exponents are $x = 0.40 \pm 0.02$ and $y = 0.14 \pm 0.01$. In B, the exponents are $x = 0.43 \pm 0.02$ and $y = 0.076 \pm 0.004$.

distribution that learned to cancel one contrast must be able to cancel different contrasts. This phenomenon is known as contrast invariance: that learning at any contrast converges to the same weight distribution for a given AM frequency. Unfortunately, there is no experimental data to align the behaviour of the model when stimulated at different contrasts; this is a theoretical constraint given the impracticality of the alternative.

In the model, it was assumed that the strength of the feedback, Λ , is proportional to $\kappa(f)$. This assumption is realistic in that deep cells will be less modulated if the contrast is lower, leading to less granule cell activity, reducing the global feedback input. As in the global η model, Λ is kept at one when $\kappa(f)$ varies with AM frequency to account for electroreceptor adaptation. Nevertheless, changes in the external stimulus intensity, simulated in this section, will change Λ and $\kappa(f)$ by the same per cent at all frequencies. Thus, if the contrast is half of its original, experimentally fitted value (*e.g.* $\kappa(0.5) = 0.25/2$, $\kappa(1) = 0.27/2$, $\kappa(2) = 0.31/2$, or $\kappa(f \geq 4) = 0.39/2$), then Λ will be 0.5. The global η model (both 2- and 4-spike bursts with η decaying at high AMs - see Figure 4.7) was investigated to see if the weight distribution learned at one contrast was identical to the distribution learned at lower contrasts at the same AM frequency (Figure 6.20). Only contrasts lower than the original contrasts were explored because *in vivo* superficial cell cancellation deteriorates at all AM frequencies for signal intensities much greater than that used to generate the experimental data (Figure 1.3).

Clearly, the weight distributions learned at different contrasts to cancel the same

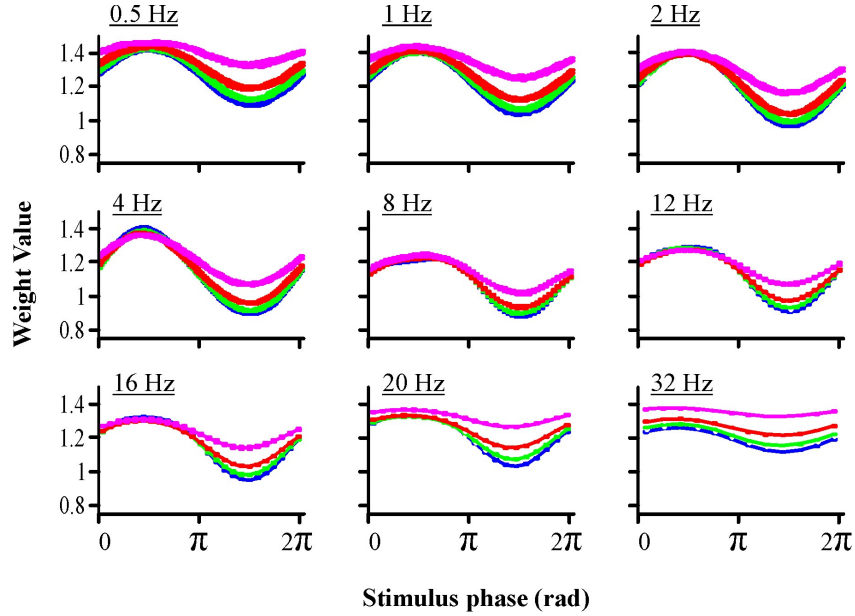


Figure 6.20: Average weight distributions learned at different contrasts with the global η model. The weight distributions were obtained by an identical procedure to Figure 6.16. Blue, green, red, and purple correspond to contrasts of 100%, 75%, 50%, and 25% of the original contrast value, respectively. The original (100%) contrast values, $\kappa_0(f)$, obtained from fitting to experimental data (see Section 4.1), are $\kappa(0.5) = 0.25$, $\kappa(1) = 0.27$, $\kappa(2) = 0.31$, $\kappa(f \geq 4) = 0.39$.

frequency are similar but not identical. However, some separation can be expected, as lower contrasts produce less bursting and less of a difference in bursts between stimulus peak and trough, which will affect the variation of the weight distribution. Further, the difference in the weight distributions is non-linear and appears to be based on the relative contrast difference rather than the absolute contrast difference: the separation between 75% and 100% contrast is the smallest since the contrast is only increased by four thirds, whereas the contrast doubles between 25% and 50%, creating the largest difference in weight distribution. The inability to cancel a sudden large change in contrast may be realistic, as a conspecific suddenly decreasing or increasing the distance to a fish may carry information that should not be removed. Nevertheless, small changes in contrast should not affect the superficial cell's behaviour, or the global feedback system would never be able to learn each new contrast quickly enough to have any kind of meaningful cancellation.

To assay whether small variations in contrast affect the model superficial neuron's behaviour, the following experiment was simulated: the model first learned to cancel a

global AM frequency at a given “learned” contrast, $\kappa_\ell(f)$, and then the contrast was adjusted to other levels while keeping the weight distribution constant (although Λ adjusts immediately with κ) to recreate a sudden change in contrast, $\Delta\kappa(f)$, before learning has time to correct it (Figure 6.21).

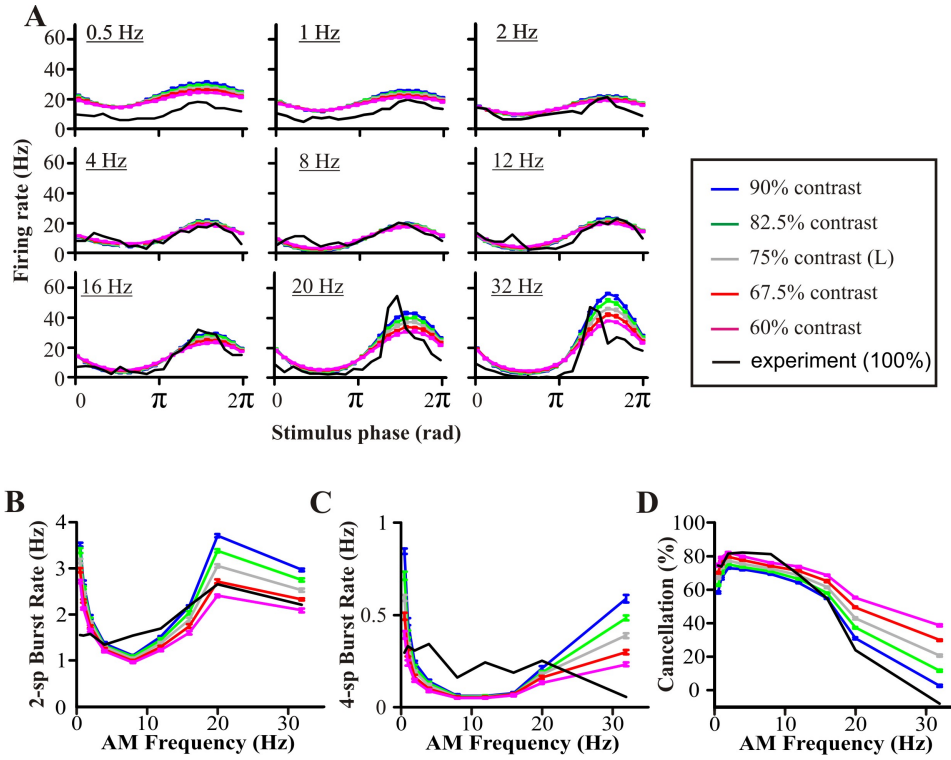


Figure 6.21: Investigation of weak contrast invariance in the model. The average weight distribution learned at 75% contrast ($\kappa_\ell(f) = 0.75\kappa_0(f)$) shown in Figure 6.20 was used in each simulation here. The weight distribution was fixed during the simulation by removing the depression and potentiation rules. $\kappa(f)$ was changed to 90% (blue), 82.5 % (green), 75% (grey), 67.5% (red), and 60% (purple) of its original value, $\kappa_0(f)$, at each frequency and the model was simulated for 1750 s. $\kappa_0(f)$ values are listed in Figure 6.20.

At a given AM frequency, when the contrast suddenly decreases from $\kappa_\ell(f)$, the cancellation improves and the burst rates decline, and a sudden increase in contrast has the opposite effect. In addition, the differences caused by the contrast shift appear to be linear, at least in the neighbourhood around the learned contrast: the change in cancellation and burst rates are approximately equal regardless of the sign of the contrast step. Further, doubling the step size doubles the change in the output statistics. Since the alteration of the cancellation and burst rates for low AM frequencies (< 16 Hz) are minimal for $\Delta\kappa(f) = \pm 15\%$ when $\kappa_\ell(f) = 0.75\kappa_0(f)$, the model is robust to small

changes in the contrast and exhibits this weaker form of contrast invariance.

Contrast invariance breaks down, however, at high frequencies. Since there is no trend in the difference in output statistics with the AM frequency before 16 Hz, this effect can be attributed to the decrease of η alone, which hinders the feedback from cancelling the input. Thus, the feedforward input drives the neuron strongly, and changes in the contrast have a greater effect. Additionally, the size of $\Delta\kappa(f)$ that can be tolerated depends on the relative difference in contrast (*i.e.* $\Delta\kappa(f)/\kappa_\ell(f)$), not the absolute difference. Contrast shifts proportional to the relative difference from the learned contrast produce equivalent results, regardless of $\kappa_\ell(f)$ (compare Figure 6.21, Figure 6.22, and Figure 6.23).

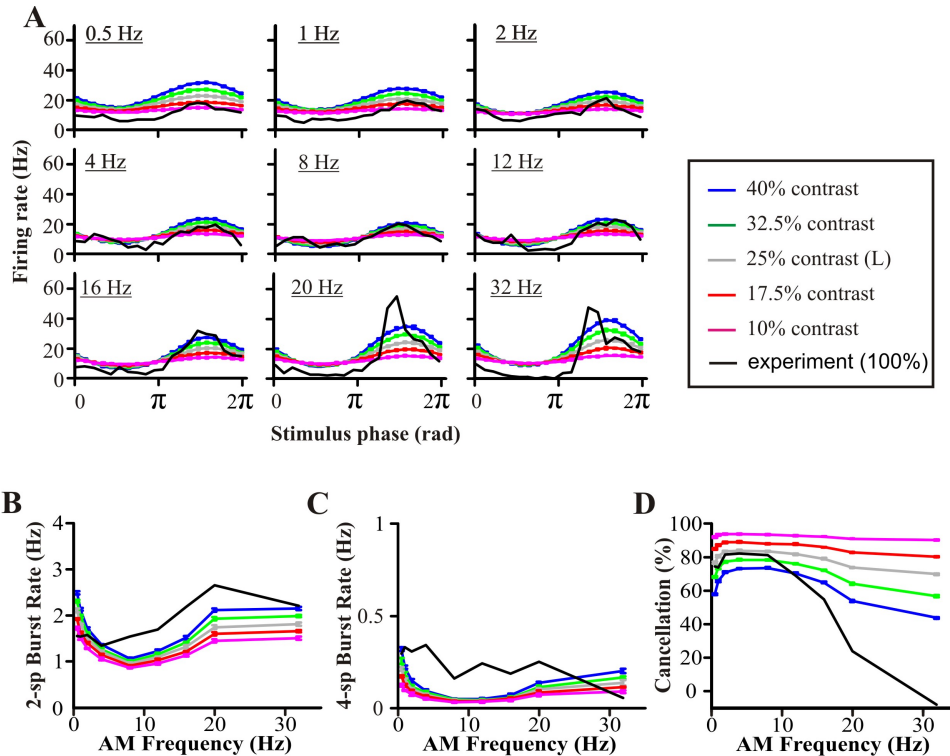


Figure 6.22: Investigation of weak contrast invariance in the model. The average weight distribution learned at $\kappa_\ell(f) = 0.25\kappa_0(f)$ as shown in Figure 6.20 was used in each simulation. The weight distribution was fixed during the simulation by removing the depression and potentiation rules. $\kappa_\ell(f)$ was changed to 40% (blue), 32.5% (green), 25% (grey), 17.5% (red), and 10% (purple) of $\kappa_0(f)$ and the model was simulated for 1750 s.

Thus, the value of $\Delta\kappa(f)$ that can be tolerated while keeping the cancellation approximately constant depends on $\kappa_\ell(f)$. This is an ideal situation given the fish's natural environment. Remember that electric fish sense their environment with an electric organ

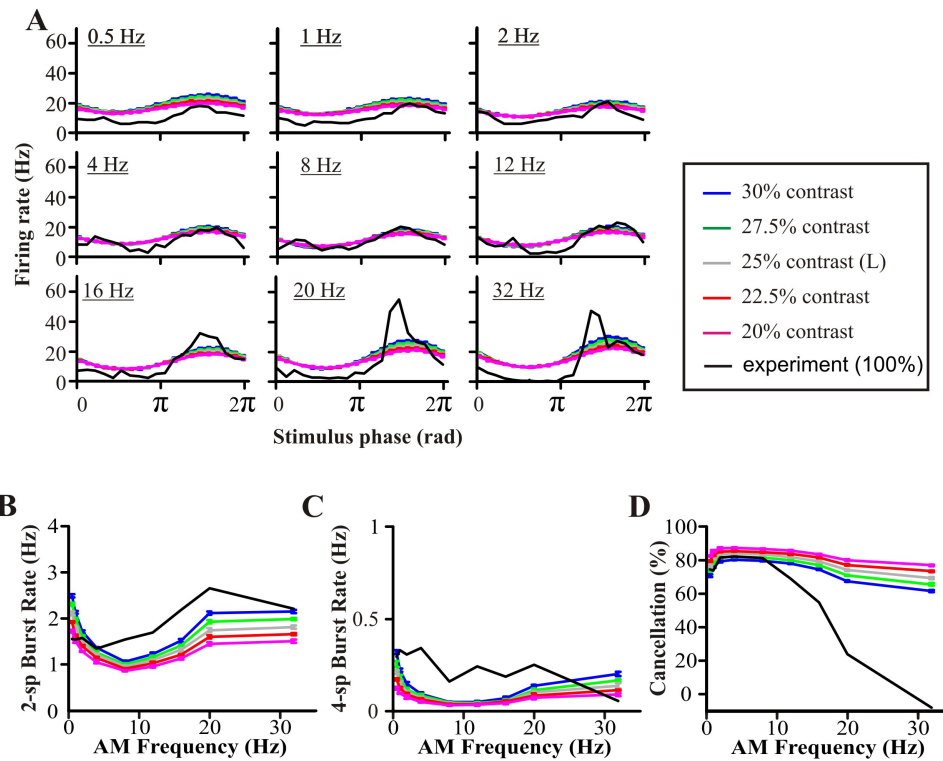


Figure 6.23: Investigation of weak contrast invariance in the model. The average weight distribution learned at $\kappa_\ell(f) = 0.25\kappa_0(f)$ was used in each simulation. The weight distribution was fixed during the simulation by removing the depression and potentiation rules. $\kappa(f)$ was changed to 30% (blue), 27.5% (green), 25% (grey), 22.5% (red), and 20% (purple) of $\kappa_0(f)$ and the model was simulated for 1750 s. This is an identical relative change in contrast compared to that in Figure 6.21.

discharge (EOD) and AM modulations of the EOD occur when two weakly electric fish are in close proximity. The strength of that signal, or its contrast compared to either fish's baseline EOD, depends approximately inversely on the distance between the fish (Figure 6.24).




Figure 6.24: A plot of the contrast of the amplitude modulation between two fish as a function of the distance between them. The fish were of equal size and arranged with their long axis parallel to each other. The amplitude modulations and the EOD were recorded by electrodes in the water. The experiment was conducted by Ginette Hupe in Professor John Lewis' laboratory in the Biology department at the University of Ottawa.

When two fish are swimming together, there will always be some fluctuations in the distance between them due to locomotion. However, the size of spatial jitter is independent of their separation. Suppose that there is a certain tolerable level of change to the cancellation without adverse effects on the fish's sensory processing. Weak contrast invariance can keep the cancellation of an AM within a tolerable interval regardless of the distance between the fish despite these spatial fluctuations. For example, looking at Figure 6.24, when a neighbouring fish is close, the contrast will be large and so will the changes in contrast due to a spatial disturbance. However, the global cancellation system is robust to these perturbations because the relative change is small. Conversely, when the fish is far away, the contrast is small and sensitive to variations, but a given spatial perturbation at this distance will induce only small variations. Thus, the global η model exhibits weak contrast invariance, which could operate to minimize the fluctuations in the cancellation of an amplitude modulation induced by a neighbouring fish with a realistic noisy position.

Chapter 7

Discussion

Electric fish are able to identify both spatially diffuse and spatially localized electrosensory signals simultaneously, even when they occupy the same low frequency bandwidth. *Apteronotus leptorhynchus* accomplishes this task by using a neural network that acts as an adaptive filter [4], preferentially cancelling redundant global sensory signals while allowing local, novel signals to transmit. This network is found in the ELL, the first electrosensory processing structure in the fish's brain. Superficial cells in the ELL receive descending input from granule cells in a cerebellar-like region that attenuates the cell's response to global low frequency stimuli but leaves the response to local signals unchanged [38].

A previous study has found that the synapse between parallel fibres, which supply the feedback input, and superficial cells undergo two types of burst-induced depression [41]. In this thesis, an adaptive filter is putatively realized in the ELL by using a cerebellar network to supply frequency-specific delay lines and an appropriate set of plasticity rules to shape the input into a negative image of the predictable signal. In addition, the statistics of the model that recreated this adaptive filter network were accurately predicted using theoretical analysis.

7.1 Both 2- and 4-spike burst rules are required for optimal cancellation

The feedback pathway in the ELL of *Apteronotus leptorhynchus* has one primary function: to cancel the amplitude modulation of any predictable global stimulus in superficial pyramidal neurons. Furthermore, the network is adaptable, and every global AM fre-

quency is independently cancelled. The ELL accomplishes this task by employing plastic synapses (*i.e.* synapses that can learn) between the pyramidal cell and a frequency-specific feedback input. The plastic synapses are modified by at least two different types of burst-induced depression rules as well as a potentiating rule, which learn to cancel the predictable stimulus.

Importantly, both burst learning rules are required for optimal cancellation across AM frequencies. The 4-spike burst rule strongly depresses the weights but has a large temporal width of ± 100 ms that is not temporally precise enough to cancel high frequency stimuli (see Figure 4.5). For example, a burst at the crest of an 8 Hz stimulus would still depress weights in the trough, which is 62.5 ms away. The 2-spike burst rule, on the other hand, causes only mild synaptic depression, but 2-spikes are more common and its burst rule is temporally very precise, which allows the parallel fibre network to cancel high frequency AMs. However, by itself, the 2-spike burst rule is unable to cancel low frequencies, as it is too weak (see Figure 4.5).

Cancellation is complicated by the interaction between the burst impact and AM frequency. Cancellation requires granule cell feedback to be phase-locked to the stimulus so the weights can have a fixed phase relationship to the input. Since each granule cell bursts for a fixed length of time, the number of active weights at each AM frequency changes. On the other hand, the time window of each burst is constant (L_w). These two factors change the proportion of the total feedback that each burst affects as a function of stimulus frequency, leading to the relationship: $\Delta\bar{w} \approx \frac{4\bar{w}\eta L_w}{3T}$ (see Section 6.2.1). Therefore, fewer bursts are required to maintain a given average weight value as the AM frequency increases, and the burst rate should be inversely proportional to the frequency (Figure 4.5).

However, the burst rate is non-linearly dependent on the average weight and does not decrease enough to exactly counteract the increased impact of the burst. Thus, the weights overall depress at high frequencies, leading to a lower firing rate, and overall potentiate at low AM frequencies, leading to a higher firing rate (Figure 4.5 and 5.3). This is a fundamental problem for any neural network that has periodic feedback modulated by frequency independent learning rules (both depressing and potentiating), because spike and burst production are always non-linearly dependent on neural inputs.

This feature might be the reason why burst-induced plasticity occurs at this synapse and not plasticity based on single spikes. With a learning event (burst or single spike or otherwise) having different effects at different frequencies due to the periodic nature of the feedback, a single learning rule is insufficient to maintain a constant level of synaptic

depression. A solution is to recruit other learning events that become more probable as the learning event impact wanes. Higher order burst rules (3-3, 4-4, 5-4, etc) are ideal for this role since each rule will uniquely affect different frequencies and are only likely at low frequencies when the strength of each learning event is small. A synapse based on single spike plasticity does not have the flexibility to add these rare but powerful learning events.

In addition, pyramidal cell bursts have been shown to selectively encode low frequency stimuli (0-16 Hz), such as prey, to higher brain centres [45]. Burst-driven cancellation, shown in this paper to be optimal below 16 Hz, would minimize bursts in superficial cells induced from predictable stimuli, reducing the noise in the putative burst channel. This is an optimal configuration to detect low-frequency stimuli even during a global predictable stimulus and strongly suggests that bursts are a separate and important information messenger.

7.2 Biological Realism

The goal of the computational investigation into the redundant input cancellation network of electric fish was to produce a biophysically realistic model that closely matches the experimental data. Since the observations of the network's behaviour were only macroscopic quantities like the firing rate of the neurons, as opposed to the membrane potentials or currents, there was some choice as to the complexity of the model. The approach to modelling in this thesis has been minimalistic in order to elucidate the components that are critical to the network's function. As such, some realism was sacrificed for simplicity. For example, instead of simulating individual populations of ion channels or the geometry of the pyramidal neuron, the leaky integrate-and-fire model was used. Further, only the superficial cell's membrane potential is simulated, and the firing activity of other neurons in the network, the electroreceptors, granule cells and inhibitory cells, are estimated based on experimental data.

Nevertheless, previous experimental measurements on the parameters of the network guided the choice of parameter values and modelling techniques. The electroreceptor firing rate was known to linearly follow the AM frequency with a constant standard deviation, and so was replicated in the model with a sinusoidally varying input current and a fixed noise level. The membrane time constant was placed at a realistic value, as was the disynaptic inhibition, whose synapses were known to have a reversal potential close to the resting potential of the neuron. The contrast invariance of the model under

global stimulation was also investigated and compared to the properties of a realistic cancelling neural network. Although biological realism sometimes complicated analytical investigations, such as rectifying the input current or a multiplicative learning rule, it was nonetheless necessary so the model could be relevant.

7.3 Model Predictions

The reward of a realistic model is the ability to make experimentally verifiable predictions about the simulated system. In this thesis, a major prediction of the model was confirmed: the existence of independent frequency channels in the feedback pathway emanating from granule cells. This prediction was based upon the biophysical structure of the parallel fibre feedback pathways that suggested they act as delay lines. When the network needs to cancel multiple frequencies independently, delay lines must be segregated into frequency channels so as to preserve their phase relationships. This was a novel prediction in the neural circuitry of electric fish, and *in vivo* experiments support this hypothesis (Figure 5.4).

Another prediction is that the granule cell burst rate should be inversely proportional to AM frequency at AMs greater than 10 Hz, as the model required this relationship to optimally match the *in vivo* data (Figure 4.8). Decreasing granule cell bursting at high frequencies is also plausible from a combinatorial argument. Assuming D'Angelo's observations [50] are consistent with electric fish granule cells, then these cells have a constant firing rate as their driving frequency is varied beyond 10 Hz. If the spikes only occur on the crests of the input, this is akin to the problem of distributing k spikes into f crests during each second and the number of bursts per second is the expected number of crests that have more than one spike. For $k = 2$, the probability of both spikes at the same crest is $1/f$, and for $k > 2$ the burst rate should still scale by a similar $1/f$ factor. Although this inverse frequency dependence was inspired by activity in other systems, it has not been experimentally investigated in electric fish.

Some potentiating rule should also exist at the parallel fibre-superficial cell synapses. Otherwise, with only depressing burst rules in effect, the strength of any synapse would inevitably approach zero since there is always some non-zero probability of a burst at any time. This would make learning and cancellation impossible. However, no potentiating rule has been found at this synapse either *in vivo* or *in vitro*. In the model, a slow exponential potentiating rule was assumed because a potentiating rule is necessary for network function, and a rule that operates over a long time scale may be difficult to

detect experimentally. This is not the only possible potentiating rule at this synapse, and identification of another type of potentiating rule, such as one that is proportional to the firing activity of either the pre- or post-synaptic neuron, would not invalidate the results shown here.

2-spike burst pairs and 4-spike burst pairs were the only burst pairings assumed to depress the parallel fibre-superficial cell synapses in the model because experimental data existed only for these pairings. With these two burst rules, the model adequately mimicked the cancellation seen *in vivo*. However, at low AM frequencies, the firing rate tended to rise (Figure 5.6) and, as described above, this is a fundamental symptom of this neural network's properties. Higher order depressing burst rules such as 5-spike burst pairs or 6-spike burst pairs would help maintain a constant firing rate as the AM frequency decreased and not affect the results for high frequency cancellation since higher order bursts become exceedingly rare at high frequencies (Figure 6.1B). Moreover, mixed pair burst rules, such as a 4-spike superficial cell burst paired with a 5 spike parallel fibre burst, would likely also cause depression and aid in reducing the firing rate at low AM frequencies. These higher order burst rules should exist and should be depressing but have not yet been explored experimentally.

7.4 Theoretical Analysis

During local stimulation, the dynamics of the model are similar to those found previously for leaky integrate-and-fire models: it operates as a low-pass filter [65], so the firing rate decreases with AM frequency (Figure 6.1), and shunting inhibition leads to subtractive gain control (Figure 6.11) [94]. In addition, the firing rate of an LIF model with Gaussian white noise input can be calculated by solving a first passage time problem (Equation 6.2) [95]. However, the model in this thesis uses coloured noise, and the effect of this input to LIF models has only been investigated when the timescales of the noise filter and model neuron are sufficiently separated [96]. Since the cut-off frequency of the filtered noise and the membrane time constant are on the same order in this thesis' model, previous studies do not apply. Therefore, the corrections to the FPT equation outlined in this thesis are novel. Further, there have been no published studies on the rectification of the input to a stochastic LIF model, but its effects were successfully incorporated in the theoretical analysis. Rectification complicates the behaviour of the model because the Gaussian white noise input no longer converges numerically; in other words, the time step of integration begins to affect the results of the simulation.

Prior to rectification, solving for the membrane potential required numerically approximating a stochastic integral, which was accomplished using the Euler-Maruyama method. This method converges because the integration of a Gaussian white noise source is a Wiener process, $W(t)$, which is a continuous function with independent increments that are normally distributed: $W(t) - W(s) \sim N(0, t - s)$ for $0 \leq s < t$ [60]. Thus, when approximating the new membrane potential $V(t + \Delta_s)$ at each time-step, one must generate a Gaussian random variable at each step with zero mean and variance Δ_s , so $\xi(t)$, which has unit variance, is multiplied by $\sqrt{\Delta_s}$. However, when the input is rectified, the integral of the Gaussian noise source is no longer a Wiener process because the increments are not normally distributed, and they have a non-zero mean. Low-pass filtering the noise allows one to use the standard Euler method to numerically approximate the integral.

Negative feedback has already been investigated analytically in stochastic LIF models and solved by calculating the effective input current, I_{eff} , and identifying a system of equations that can be solved self-consistently to find the equilibrium firing rate [97]. A similar approach was used to study the depolarizing after-potential (DAP), a known source of positive feedback for superficial pyramidal neurons in the ELL. The DAP is a highly complex system that has been modelled previously [57] and was directly incorporated into the model in this thesis. Nevertheless, reducing the DAP to a single feedback parameter on the firing rate was a successful approximation in this model given the parameters that mimic experimental local stimulation. This was a novel method to analyze the dendritic after-potential.

During global stimulation, the feedback circuit between the synaptic weights and the depression and potentiation rules present in the model is unique to this thesis and as such has not previously been investigated. Nevertheless, a system of equations was derived that accurately predicted the equilibrium weight distribution, burst rates and firing rate modulation of the model. Furthermore, power law relationships between the mean weight and firing rate values and the global parameters η , L_w , τ_w , and f were identified. In other words, the system maintains the same behaviour despite changes in its parameters. These power laws hold for reasonable variations in the parameters, although the laws do breakdown when the fluctuations around the equilibrium weight values surpass a critical size. At this point, the system may exhibit new dynamics, for which these power laws are no longer valid. However, further investigation is required.

Note that if any of the global parameters η , L_w or τ_w , assumed to be constant in the 2-spike burst-rule-only model, were to vary inversely with frequency, then the firing

rate and average weight value would be identical for all AM frequencies. Experimental observations during global stimulation reveal that the average firing rate is approximately constant for AM frequencies less than 12 Hz (Figure 2.9). Although recruitment of higher order burst rules could conceivably maintain the firing rate, this may not be necessary if, for example, the time constant of weight potentiation or the learning rate varied at low AM frequencies.

The concept of contrast invariance was also studied in the global η model. While the model did not display true contrast invariance, a weaker form of it was still present that is putatively sufficient to accommodate natural spatial noise in the electric fish's environment. The lack of pure contrast invariance is due to the non-linearities of the system during global stimulation. This is best described by reducing the weight distribution to two parameters. When the feedback pathway is active, the parallel fibre input can be well approximated by a sinusoid in anti-phase to the feedforward input with amplitude w_{amp} and bias \bar{w} :

$$\Lambda(w_\phi - gV) \approx \Lambda\bar{w} - \Lambda w_{amp} \sin(2\pi ft) - \Lambda gV \quad (7.1)$$

Decreasing the contrast will lead to a smaller difference between the firing rates at the crest and the trough of the input and, therefore, to a smaller necessary modulation of the feedback (Λw_{amp}). However, Λ , the strength of the feedback, scales with the contrast and so w_{amp} should be approximately the same value despite contrast variations. Similarly, Λ affects both the mean weight and the disynaptic inhibition equally so that \bar{w} should be invariant to changes in contrast.

Unfortunately, this assumes that the system is linear and that a change in the input will produce a linear change in the equilibrium weight regardless of the phase. This assumption is false. For weights in the trough of the input, the equilibrium firing rate at most AM frequencies is low and both the R-I and Br-R curves (Figures 6.11 and 6.12) have a shallow slope. Therefore, changes in the input current due to different contrasts lead to small changes in the equilibrium weight value in the trough.

Conversely, the firing rate is high at the peak of the input, which places the R-I curve in its linear regime and the Br-R curve into a very steep section. This causes a large difference in equilibrium weight values at the crest of the input that were learned at different contrasts. Thus, in essence, there must be a unique Λ for each phase of the AM cycle, and using only one value creates the differences in the weight distributions between contrasts (Figure 6.20).

Like input cancellation, contrast invariance deteriorates at high AM frequencies. Both of these effects are due to the weakening of the feedback pathway as η decreases. However, since natural prey signals for electric fish are in the 0-16 Hz bandwidth [33], it is only necessary for the low-frequency region to be clear of redundant signals, which is accomplished in the ELL and in the model. Thus, in an economical manner, only the frequencies that are essential to survival are effectively cancelled.

7.5 Comparison to other Systems

The cellular network that produces cancellation in the ELL of electric fish has been investigated previously by John Lewis [98, 41]. Lewis' approach utilized short-term plasticity rules found *in vivo* at the parallel fibre-superficial cell synapse to create a negative image of the input, and it did not require a delay structure or independent frequency channels in the feedback. It does require, however, a potentiating synapse between parallel fibres and inhibitory cells, and parallel fibre activity to be continuous and modulated by the AM stimulus without a phase delay. In other words, the firing rate of the parallel fibre is sinusoidally modulated around a mean rate with the modulation at the same frequency as the input stimulus. In addition, the inhibitory cells' activity must be in phase with the parallel fibres' activity. Lewis' model assumed that long-term depression produced by bursts would decrease the conductance of the parallel fibre synapse to a stable level sufficient to create a negative image. It did not investigate how the burst-induced depression would behave online to create a stable synaptic strength value, nor how this may change as the stimulus frequency changed.

For simplicity, the model presented in this thesis assumes that a parallel fibre only bursts, and does so once per stimulus cycle. Since bursts are the critical feedback input, a single burst per period is approximately equivalent to a sinusoidally modulated input that always bursts at the crest of each cycle. The non-bursting input of the sinusoids at a multitude of delays would sum to a constant input and could be incorporated into the feedback strength, Λ . Thus, Lewis's parallel fibre simulation is analogous to the parallel fibre modelling approach in this thesis.

As this thesis employs a minimal model, no assumptions were made on the short term dynamics of any synapse. Including short term plasticity at the parallel fibre-superficial cell synapse would have a minor effect in changing the shape of parallel fibre input from a step function to something more complex. Furthermore, disynaptic inhibition in this thesis was modelled as a constant term. Since there are so few inhibitory cells compared

to parallel fibres, likely many parallel fibres synapse onto one inhibitory cell. When the input current of a multitude of parallel fibres, each with its own phase relationship, is added inside the inhibitory cell, the sum would not have any phase dependence (*i.e.* if all phases are represented equally, the input is constant over one period). Any plasticity at the inhibitory synapses is not required but could aid cancellation.

A common hypothesis for the function of cerebellar networks is to generate precisely timed inputs to process sensory signals [32, 31, 91]. Roberts and Bell have postulated the existence of such a network in the ELL of mormyrid electric fish that is used with synaptic plasticity for both the cancellation of expected stimuli and for the storage of temporal information [99, 40]. Another family of electric fish with a pulsatile EOD, mormyrids have a electrosensory processing structure comparable in physiology and function to *Apteronotus leptorhynchus*. Similar to pyramidal cells, medium ganglion cells in the mormyrid ELL receive both feedforward input from electrosensory afferents as well as feedback input provided by parallel fibres from a cerebellar-like network that putatively function as delay lines. Like *Apteronotus leptorhynchus*, the synapse between the parallel fibres and the MG cell is plastic, although an anti-Hebbian STDP rule governs the synaptic strength in mormyrids [20], as opposed to multiple correlative burst-induced depression rules in apteronotids.

The reason for different rules in these two similar networks likely arises from the origin of the parallel fibre feedback: parallel fibres in mormyrids are synchronized by an electric organ corollary discharge (EOCD) from the electric organ itself, which does not occur in apteronotids. Thus, the cancellation problem in mormyrids reduces to learning the time delays of inputs after the EOCD. On the other hand, apteronotid feedback must synchronize directly with the redundant global stimuli of arbitrary frequency by phase-locking granule cells. Because of this, the feedback is periodic with the predictable input, creating the frequency dependence of the firing rate described above and the need for multiple learning rules.

When the network needs to cancel multiple frequencies independently, delay lines must be segregated into frequency channels so as to preserve their phase relationships. Although assumed initially in this model, experimental evidence presented in this thesis supports the hypothesis of independent frequency channels in the apteronotid feedback pathway (Figure 5.7). Frequency channels have been suggested in other adaptive sensory processing networks, such as in the vestibulo-ocular reflex (VOR) [85], where the gain learned at one frequency is independent of the gains learned at other frequencies. Furthermore, studies have shown that the VOR's frequency-specific adaptation requires

long-term depression at parallel fibre synapses [90], which parallel the results shown here in electric fish. Experimental studies have investigated the cause of frequency tuning in the granular region [86, 50] but, unfortunately, the exact mechanism is still unknown.

There have been two separate hypotheses about the function of cerebellar structures: that they operate as an adaptive filter or as a generator of precise timing information. Results of the investigation in the ELL of *Apteronotus leptorhynchus* show that these two theories may operate concurrently. The ELL uses the delay structure of the diverse parallel fibre feedback pathway together with long-term synaptic plasticity to create an adaptive filter that can effectively cancel low frequency redundant input signals to maximize prey detection. In addition, the feedback is segregated into frequency channels so the network may learn to cancel individual frequencies independently. Unlike current theories of the mammalian cerebellum, the learning is unsupervised but nevertheless leads to stable storage of synaptic weights. Further, burst learning rules are employed instead of the standard spike-time dependent (STDP) rules. Yet, they are still able to effectively teach the appropriate synaptic strengths and may in fact be preferential to a single spike learning rule when multiple frequencies must be cancelled. The temporal shape of burst plasticity for each burst size has also been coordinated with the intrinsic dynamics of the burst rate in superficial cells to optimize cancellation across different AM frequencies. The matching of intrinsic dynamics with synaptic plasticity as well as the occurrence of independent frequency channels are important elements in the composition of the ELL and may be essential to cerebellar circuits in other systems.

Chapter 8

Conclusion

8.1 Summary of Results

The brain is perhaps the most fundamental structure for any living creature, and yet it is one of the least understood systems in science. Due to the highly connected nature of the brain, investigating how processing techniques are implemented at the microscopic level has proven difficult. Analysis of sensory processing structures can resolve some of these difficulties because the inputs to such regions can be experimentally observed and controlled. Further, studies of neural behaviour in model organisms, like weakly electric fish, can simplify the external anatomical and physiological conditions and make investigations more tractable.

Weakly electric fish emit an electric discharge to explore their environment. The region that analyzes the electrosensory input in their brain is called the ELL and has a cerebellar-like structure. Functional investigations have discovered that the ELL operates as an adaptive filter, which removes redundant global stimuli from the input, and physiological studies have suggested that a feedback loop with adjustable synapses produces this cancellation. In addition, electric signals are simple to detect and modulate and recording from neurons in the ELL is straightforward since the ELL has a relatively simple architecture. This is a system that can be easily manipulated and observed, and its analysis will help identify the dynamics of the cerebellar regions as well as learning neural networks, in general.

Recently, the exact synaptic plasticity rule that putatively induces learning in this structure was discovered to be based on burst events rather than single spikes or average firing rates. With this burst rule and the known anatomy of the ELL, a realistic minimal

model was built to gain insight into the dynamics of networks that employ burst-induced learning and what the necessary requirements for adaptive cancellation are. Due to the available data, two burst rules were incorporated: a 2-spike burst rule and a 4-spike burst rule. It was discovered that both rules operating simultaneously are necessary for optimal cancellation of redundant input frequencies, but the range of frequencies that can be cancelled *in vivo* depends crucially on the burst rates of the neurons that supply the feedback. Furthermore, the feedback must be separated into independent frequency channels if cancellation at any arbitrary frequency is required, a hypothesis that was later corroborated experimentally. These results show how an adaptive filter can be created with biological components: using a cerebellar structure with multiple learning rules and segregated frequency pathways.

This also illuminates the role of cerebellar structures, which have been thought to be either adaptive filters or timing devices. In the ELL, these two functions are not mutually exclusive and so harmonize the two disparate theories of cerebellar activity. Moreover, the use of a novel burst rule for unsupervised learning in the cerebellum both extends spike-time dependent plasticity to higher order events and shows that learning in the cerebellum can still be stable even without a teacher signal.

In addition, the dynamics of the model were successfully analyzed theoretically and generated some interesting questions about the behaviour of stochastic LIF models. Techniques to model the DAP and identify the equilibrium weight values during global stimulation were also discovered that were highly effective, at least in the neighbourhood of the parameters of the optimal model. As well, the model maintained its realism by displaying contrast invariance, a hypothetical requirement of any *in vivo* cancellation system.

8.2 Future Directions

Evidence of independent frequency channels in the feedback pathway raises some new questions for future research. Unlike the mammalian auditory network, where frequencies are segregated based on the mechanical properties of the input structure, there is only one input channel to the feedback pathway (deep cells), which transmits all AM frequencies equally. Although frequency-tuned cerebellar networks have been suggested before [86], there have been no prior investigations into the mechanism behind frequency identification and segregation of incoming AM stimuli in the electric fish granule cell network. Another interesting feature would be to ascertain the frequency channel band-

widths in the range of cancelled frequencies. For example, even though cancellation at 8 Hz is unaffected by cancellation at 2 Hz, it is affected by cancellation at 9 Hz (data not shown).

There is also preliminary evidence of different burst learning rules in different regions of the ELL. The data on 2-spike and 4-spike burst depression were gathered from observations of E type superficial cells in the centro-lateral segment of the ELL. However, there are two other regions of the ELL, each hypothesized to function at different frequency bandwidths, from low frequency AMs to search for prey to high frequency AMs for communication. It is possible that the burst-induced depression at the superficial cell-parallel fibre synapse in each region is specifically tailored to the function of that region. Further, it appears that superficial I cells may have different burst rules than the E cells in same segment. Although E and I cells do not synapse onto each other, both project to higher brain centres, and it is unknown what different information each transmits or why they should have different learning rules.

The property investigated in this thesis was the cancellation of single frequencies by the ELL in *Apeteronotus leptorhynchus*. This situation does appear in nature when two electric fish are in close proximity and their EODs begin to beat and removal of this signal does require an adaptive filter. However, these fish also adaptively cancel global signals such as tailbending that occurs over a range of frequencies. It is unclear how the ELL circuitry can cancel multiple frequencies simultaneously, even with independent frequency channels, as a single superficial cell burst will affect the weight distributions of all active channels even if the burst was induced at the crest of only one frequency. Therefore, the cancellation of multiple AMs will not be a linear superposition of cancellation at individual frequencies. Current studies are exploring the behaviour of the model when trying to cancel two discrete frequencies, and will hopefully expand their results to explain the cancellation of a continuous frequency range.

The analytical investigation of the model was fairly successful in predicting some of the outputs of the model under both local and global stimulation conditions. Nevertheless, there are still some areas that need to be explored. The mean input correction for low-pass filtered noise, though effective, is *ad hoc* and requires a theoretical foundation. The linearization of the burst rate to firing rate relationship with the addition of an AM frequency should also be investigated. In addition, the theoretical approach to global stimulation should be extended to include 4-spike bursts and the power law dependence on the average outputs with the input parameters requires further study. Experimental corroboration of the contrast invariance exhibited by the model is also necessary.

Nomenclature

AM	Amplitude modulation: the envelope of the EOD that carries information about the environment and stimulates electroreceptors
CLS	Central lateral segment: the region of the ELL where burst-induced depression of PF-SP cell synapses was discovered
CNS	Central nervous system: the neurons in the brain and spinal cord and all supporting cells
DAP	Depolarizing after-potential: a delayed excitatory input to an SP cell after it has fired
E cells	Excitatory pyramidal cells: a type of pyramidal cell that increases its firing rate to increased AM intensity in its RF
EGp	Eminentia granularis posterior: a cerebellum-like structure in the brain of weakly electric fish where granule cells reside
ELL	Electrosensory lateral line lobe: the first sensory processing region in the brain of weakly electric fish for their electric sense
EOCD	Electric organ corollary discharge: available in mormyrid fish to synchronize granule cells to their EOD
EOD	Electric organ discharge: high frequency electric signal used by apteronotids to sense their surroundings
EPSP	Excitatory post-synaptic potential: the increase in membrane potential of a post-synaptic neuron from activation of a synapse
GABA	Gamma-aminobutyric acid: a neurotransmitter for inhibitory synapses

I cells	Inhibitory pyramidal cells: a type of pyramidal cell that decreases its firing rate to increased AM intensity in its RF
LIF	Leaky integrate-and-fire: a simple non-linear model for neural firing activity
LTD	Long-term depression: a stable decrease in the strength of a synapse over many hours
LTP	Long-term potentiation: a stable increase in the strength of a synapse over many hours
NMDA	N-methyl-D-aspartic acid: receptors of this molecule are associated with synaptic plasticity and learning
nP	Nucleus praeminentialis: a region of the ELL that receives input from deep pyramidal cells and project to the EGp
PF	Parallel fibres: numerous thin fibres from granule cells in the EGp that putatively cancel redundant stimuli
PSTH	Post-stimulus time histogram: a histogram of neural firing activity synchronized to a characteristic of the stimulus, such as its phase
RF	Receptive field: the region on the skin of a particular neuron that, if stimulated, will alter the firing activity of that neuron
SP	Superficial pyramidal cell: the neuron in the CLS of the ELL that cancels low-frequency global stimuli. This thesis attempted to model its behaviour
STDP	Spike-timing dependent plasticity: the phenomenon that individual pre- and post-synaptic spike timing significantly effects the synaptic strength

Bibliography

- [1] E. C. Cherry, *On Human Communication: A Review, Survey, and a Criticism*, MIT press, Cambridge, Mass., 1957.
- [2] S. Haykin and Z. Chen, *Neural Comp* **17**, 1875 (2005).
- [3] P. D. Roberts and C. V. Portfors, *Biol Cyber* **98**, 491 (2008).
- [4] N. B. Sawtell and A. Williams, *J Neurosci* **28**, 1598 (2008).
- [5] P. Dean, J. Porrill, C.-F. Ekerot, and H. Jörntell, *Nat Rev Neurosci* **11**, 30 (2010).
- [6] W. Lytton, *From Computer to Brain*, Springer-Verlag, New York, 2002.
- [7] P. Dayan and L. F. Abbott, *Theoretical Neuroscience*, MIT press, Cambridge, Mass., 2001.
- [8] P. Nelson, *Biological Physics*, W H Freeman and Company, New York, 2002.
- [9] F. Rieke, D. Warland, R. de Ruyter van Steveninck, and W. Bialek, *Spikes: Exploring the Neural Code*, MIT press, Cambridge, Mass., 1997.
- [10] D. Gussin, J. Benda, and L. Maler, *J Neurophysiol* **97**, 2917 (2007).
- [11] P. D. Roberts and C. C. Bell, *Biol Cyber* **87**, 392 (2002).
- [12] D. Hebb, *The organization of behavior*, John Wiley and Sons, New York, 1949.
- [13] T. V. Bliss and T. Lomo, *J Physiol* **232**, 331 (1973).
- [14] J. H. Hopfield, *Proc Nat Acad Sci USA* **79**, 2554 (1982).
- [15] T. Kohonen, *Self-organization and Associative Memory*, Springer, Berlin-Heidelberg, 1989.

- [16] H. Markram, *Science* **275**, 213 (1997).
- [17] S. Song, K. D. Miller, and L. F. Abbott, *Nat Neurosci* **3**, 919 (2000).
- [18] X.-H. Xie and H. S. Seung, Spike-based learning rules and stabilization of persistent neural activity, in *Adv Neural Inf Process Syst*, edited by S. A. Solla, T. K. Leen, and K.-R. Muller, pages 164–170, MIT press, Cambridge, Mass., 2000.
- [19] G. Bi and M. Poo, *Nature* **401**, 792 (1999).
- [20] C. Bell, D. Bodznick, J. Montgomery, and B. J., *Brain, Behav and Evol* **50**, 17 (1997).
- [21] Y. Dan and M.-m. Poo, *Science* **256**, 1570 (1998).
- [22] V. Egger, D. Feldmeyer, and B. Sakmann, *Nat Neurosci* **2**, 1098 (1999).
- [23] J.-P. Pfister and W. Gerstner, *J Neurosci* **26**, 9673 (2006).
- [24] R. C. Froemke, I. A. Tsay, M. Raad, J. D. Long, and Y. Dan, *J Neurophysiol* , 1620 (2006).
- [25] A. Morrison, M. Diesmann, and W. Gerstner, *Biol Cyber* **98**, 459 (2008).
- [26] M. Ito, *The cerebellum and neural control*, Raven press, New York, 1984.
- [27] S. Albus, *Math Biosci* **10**, 25 (1971).
- [28] B. Y. D. Marr, *Journal of Physiology* **202**, 437 (1969).
- [29] D. Purves et al., *Neuroscience*, Sinauer Associates, Sunderland, Maine, second edition, 2001.
- [30] J. Porrill and P. Dean, *PLoS Comp Biol* **3** (2007).
- [31] J. F. Medina, K. S. Garcia, W. L. Nores, N. M. Taylor, and M. D. Mauk, *J Neurosci* **20**, 5516 (2000).
- [32] J. Meek, *Neuroscience* **48**, 249 (1992).
- [33] M. E. Nelson and M. A. Maciver, *J Exp Biol* **202**, 1195 (1999).

- [34] L. Chen, J. L. House, R. Krahe, and M. E. Nelson, *J Comp Physiol A* **191**, 331 (2005).
- [35] D. Babineau, A. Longtin, and J. E. Lewis, *J Exp Biol* **209**, 3636 (2006).
- [36] C. E. Carr and M. L., *Electroreception in gymnotiform fish: central anatomy and physiology*, in *Electroreception*, edited by T. H. Bullock and W. Heiligenberg, pages 319–374, Wiley, New York, 1986.
- [37] J. Bastian, *J. Comp. Physiol. [A]* **176**, 63 (1995).
- [38] J. Bastian, M. J. Chacron, and L. Maler, *Neuron* **41**, 767 (2004).
- [39] R. Krahe and F. Gabbiani, *Nat Rev: Neurosci* **5**, 13 (2004).
- [40] P. D. Roberts and C. C. Bell, *J Comp Neurosci* **9**, 67 (2000).
- [41] E. Harvey-Girard, J. Lewis, and L. Maler, *J Neurosci* **30**, 6152 (2010).
- [42] J. Saunders and J. Bastian, *J Comp Physiol A* **154**, 199 (1984).
- [43] N. Berman and L. Maler, *J Exp Biol* **202**, 1243 (1999).
- [44] M. J. Chacron, B. Doiron, L. Maler, A. Longtin, and J. Bastian, *Nature* **423**, 77 (2003).
- [45] A.-M. M. Oswald, M. J. Chacron, B. Doiron, J. Bastian, and L. Maler, *J Neurosci* **24**, 4351 (2004).
- [46] J. Bastian et al., *J Neurophysiol* , 10 (2001).
- [47] E. Sas and L. Maler, *Anat and Embryol* **177**, 55 (1987).
- [48] P. Chadderton, T. W. Margrie, and M. Hausser, *Nature* **428**, 856 (2004).
- [49] E. a. Rancz et al., *Nature* **450**, 1245 (2007).
- [50] E. D'Angelo et al., *J Neurosci* **21**, 759 (2001).
- [51] S. Sherman, *Trends in Neurosci* **24**, 122 (2001).
- [52] A. Kepecs and J. Lisman, *Network* **14**, 103 (2003).

- [53] G. Marsat, R. D. Proville, and L. Maler, *J Neurophysiol* , 714 (2009).
- [54] G. Marsat and L. Maler, *J Neurophysiol* (2010), doi: 10.1152/jn.00256.2010.
- [55] K. Wiesenfeld and F. Moss, *Nature* **373**, 33 (1995).
- [56] B. Lindner, L. Schimansky-Geier, and A. Longtin, *Phys Rev E* **66** (2004).
- [57] L. Noonan, B. Doiron, C. Laing, A. Longtin, and R. W. Turner, *Neurosci Res* **23**, 1524 (2003).
- [58] M. J. Chacron, a. Longtin, M. St-Hilaire, and L. Maler, *Phys Rev Lett* **85**, 1576 (2000).
- [59] J. Benda, A. Longtin, and L. Maler, *J Neurosci* **25**, 2312 (2005).
- [60] H. C. Tuckwell, *Introduction to Theoretical Neurobiology*, volume 2, Cambridge university press, Cambridge, Mass., 1988.
- [61] B. Doiron, a. Longtin, R. W. Turner, and L. Maler, *J Neurophysiol* **86**, 1523 (2001).
- [62] L. Maler and E. Mugnaini, *J Comp Neurol* **345**, 224 (1994).
- [63] N. J. Berman and L. Maler, *J Neurophysiol* **80**, 3173 (1998).
- [64] P. E. Kloeden and E. Platen, *Numerical Solution of Stochastic Differential Equations*, Springer, Berlin, 1999.
- [65] H. C. Tuckwell, *Introduction to Theoretical Neurobiology*, volume 1, Cambridge university press, Cambridge, Mass., 1988.
- [66] J. E. Lewis and L. Maler, *J Neurophysiol* **88**, 1695 (2002).
- [67] J. E. Lewis and L. Maler, *J Neurophysiol* **91**, 1064 (2004).
- [68] E. Harvey-girard, R. J. Dunn, and L. E. N. Maler, *J Comp Neurol* **505**, 644 (2007).
- [69] L. Maler, *J Comp Neurol* **183**, 323 (1979).
- [70] B. Doiron, C. Laing, A. Longtin, and L. Maler, *J Comp Neurosci* **12**, 5 (2002).
- [71] N. Lemon et al., *J Neurophysiol* , 1519 (2000).

- [72] R. W. Turner, L. Maler, T. Deerinck, S. R. Levinson, and M. H. Ellisman, *J Neurosci* **14**, 6453 (1994).
- [73] Z. F. Mainen and T. J. Sejnowski, *Nature* **382**, 363 (1996).
- [74] J. Bastian, *J Neurosci* **6**, 553 (1986).
- [75] F. Gabbiani, W. Metzner, R. Wessel, and C. Koch, *Nature* **384**, 564 (1996).
- [76] L. Maler, E. Sas, S. Johnston, and W. Ellis, *J. Chem Neuroanat* **4**, 1 (1991).
- [77] N. B. Sawtell, *Neuron* **66**, 573 (2010).
- [78] N. J. Berman and L. Maler, *J Neurophysiol* **80**, 3214 (1998).
- [79] W. H. Mehaffey, B. Doiron, L. Maler, and R. W. Turner, *J Neurosci* **25**, 9968 (2005).
- [80] C. a. Shumway and L. Maler, *J Comp Physiol A* **164**, 391 (1989).
- [81] L. Maler, *J Comp Neurol* **183**, 323 (2007).
- [82] N. Berman, R. J. Dunn, and L. Maler, *J Neurophysiol* **86**, 1612 (2001).
- [83] S. Jande, L. Maler, and E. M. Lawson, *Nature* **294**, 765 (1981).
- [84] L. Maler, S. Jande, and E. M. Lawson, *Brain Res* **301**, 166 (1984).
- [85] S. G. Lisberger, F. A. Miles, and L. M. Optican, *J Neurosci* **3**, 1234 (1983).
- [86] G. P. Dugué et al., *Neuron* **61**, 126 (2009).
- [87] J. Bastian and B. Bratton, *J Neurosci* **10**, 1226 (1990).
- [88] J. W. Middleton, A. Longtin, J. Benda, and L. Maler, *Proc Nat Acad Sci USA* **103**, 14596 (2006).
- [89] E. D'Angelo and C. I. De Zeeuw, *Trends Neurosc* **32**, 30 (2009).
- [90] C. I. De Zeeuw et al., *Neuron* **20**, 495 (1998).
- [91] W. M. Kistler, J. L. van Hemmen, and C. I. De Zeeuw, Time window control: a model for cerebellar function based on synchronization, reverberation, and time slicing, in *Prog Brain Res*, edited by N. M. Gerrits, T. J. H. Ruigrok, and C. I. De Zeeuw, volume 124, Wiley, New York, 1986.

- [92] D. Oertel and E. D. Young, *Trends Neurosci* **27**, 104 (2004).
- [93] T. Tzounopoulos and N. Kraus, *Neuron* **62**, 463 (2009).
- [94] G. R. Holt and C. Koch, *Neural Comp* **9**, 1001 (1997).
- [95] B. Doiron, B. Lindner, A. Longtin, L. Maler, and J. Bastian, *Phys Rev Lett* **96** (2004).
- [96] N. Brunel and S. Sergi, *J Theo Biol* **195**, 87 (1998).
- [97] C. Sutherland, B. Doiron, and A. Longtin, *Biol Cybern* **100**, 475 (2009).
- [98] J. E. Lewis, B. Lindner, B. Laliberté, and S. Groothuis, *J Exp Biol* **210**, 4437 (2007).
- [99] P. D. Roberts, *J Comp Neurosci* **7**, 235 (1999).