2011

Encoding strategies and mechanisms underpinning adaptation to stimulus statistics in the rat barrel cortex

Davies, Lucy Anne

ENCODING STRATEGIES AND MECHANISMS UNDERPINNING ADAPTATION TO STIMULUS STATISTICS IN THE RAT BARREL CORTEX

LUCY ANNE DAVIES

PhD 2010

Encoding strategies and mechanisms underpinning adaptation to stimulus statistics in the rat barrel cortex

Lucy Anne Davies

University of Plymouth

A thesis submitted to the University of Plymouth
in partial fulfilment for the degree of
DOCTOR OF PHILOSOPHY

School of Psychology,
Centre for Robotics and Neural Systems

December 2010

Abstract

It is well established that, following adaptation, cells adjust their sensitivity to reflect the global stimulus conditions. Two recent studies in guinea pig inferior colliculus (IC, Dean, Harper & McAlpine 2005) and rat barrel cortex (Garcia-Lazaro, Ho, Nair & Schnupp 2007) found that neural stimulus-response functions were displaced laterally in a manner that was dependent on the mean adapting stimulus. However, the direction of gain change, following adaptation to variance, was in contradiction to Information Theory, which predicts a decrease in gain with increased stimulus variance.

On further analysis of the experimental data, presented within this thesis, it was revealed that the adaptive gain changes to global stimulus variance were, in fact, in the direction predicted by Information Theory. However, following adaptation to global mean amplitude, neural threshold was displaced to centre the SRF on inputs that were located on the edge of the stimulus distribution. It was found that adaptation scaled neural output such that the relationship between firing rate and local, as opposed to global, differences in stimulus amplitude was maintained; with the majority of cells responding to large differences in stimulus amplitude, on the 40ms scale. A small majority of cells responded to step-size differences, in amplitude, of either direction and were classed as novelty preferring.

Adaptation to global mean was replicated in model neuron with spike-rate adaptation and tonic inhibition, which increased with stimulus mean. Adaptation to stimulus variance was replicated in three models 1: By increasing, in proportion to stimulus variance, background, excitatory and inhibitory firing rates in a balanced manner (Chance, Abbott & Reyes 2002), 2: A model of asymmetric synaptic depression (Chelaru & Dragoi 2008) and 3: a model combining non-linear input with synaptic depression.

The results presented, within this thesis, demonstrate that neurons change their coding strategies depending upon the global levels of mean and variance within the sensory input. Under low noise conditions, neurons act as deviation detectors, i.e. are primed to respond to large changes in the stimulus on the tens of millisecond; however, under conditions of increased noise switch their encoding strategy in order to compute the full range of the stimulus distribution through adjusting neural gain.

Contents

List of Figures

ACKNOWLEDMENTS

First of all, I would like to dedicate this thesis to the memory of Andrew Symes, fellow PhD-er, and sorely missed friend.

This PhD was funded , via the EPSRC, through the COLAMN project.

I would like to thank my supervisors, Dr Sue Denham and Dr Thomas Wennekers, for all their support, understanding, good advice and encouragement throughout the roller-coaster ride of this PhD. I would also like to thank Jan Schnupp and Jose Garcia-Lazaro for providing me with the experimental data. I need to thank Floortje Klijne, Pauline Campbell, Salvador Dura, Ulrich Bartsch, Beibei Jiang and everyone else who shared the red office space with me over the last four years - thanks for all the fun, laughter, coffee, sinterklass poems and general good times! Special mention also needs to go to my good friends Cathy McCabe, for being just lovely (and proof reading!) and Bea Rogers for all her friendship and support. In addition Martin Coath saved my life by teaching me \LaTeX.

Finally, I need to thank my family, who are everything to me, and have been so patient, supportive, loving, caring, as well as wonderful, fun, and a constant source of happiness and inspiration. So thank you, with all my heart, Seren Pascoe Davies, Sue Pascoe and Shaun Lewin. You have been and will always be truly amazing!

AUTHOR'S DECLARATION

At no time during the registration for the degree of Doctor of Philosophy has the author been registered for any other University award.

This study was financed with the aid of a studentship from the Engineering and Physical Sciences Research Council.

Relevant scientific seminars and conferences were regularly attended at which work was often presented.

Posters and conference presentations:

2008: **Federation of European Neuroscience Societies**; bi-annual meeting, Geneva. Poster presentation, detailing the initial analysis of the adaptation to stimulus statistics data.

2007: **Computational Neurosciences**; annual meeting, Berlin. Poster presentation. *Modelling the mechanisms underpinning sensory adaptation and gain control.*

Word count for the main body of this thesis: 68,973

Signed:

Date: 21/7/11

Chapter 1: Overview of thesis

The development of Information Theory (Shannon 1948; Shannon 1949) and, consequently the Efficient Coding Hypothesis (Barlow 1961) have had a profound influence on the understanding of neural dynamics over the years.

Information theory is mathematical system developed in the late 1940's as a tool for analyzing and understanding communication systems (Shannon 1948; Shannon & Weaver 1949). Its main purpose was to quantify compression, reliability and transmission of inputs to and within systems. This was achieved through formalising concepts such as information (e.g. signals are transmitted in a binary code, i.e. as a series of '0's and '1', the unit of information is the bit) and entropy (a measure of uncertainty about a random variable).

Nervous systems can also be viewed as biological communication systems, in that sensory inputs are transformed into binary, electrical signals, which are then transmitted across the network of neurons. As such, spiking activity can be quantified in terms of information, i.e. as the number of bits of information carried by each spike.

Due to the vast array of environmental information, with which organisms are constantly bombarded, signals must be highly compressed to allow for communication along finite neural pathways. Compression of any signal can result in loss of information; clearly this is an undesirable situation when sensory information is the key to survival. Horace Barlow (Barlow 1959; Barlow 1962; Barlow 1972) believed that nervous systems must have evolved to minimise both data loss due to compression and action potential generation (due to the high metabolic cost of spiking activity); however, these are obviously conflicting requirements. Using Information theory as its basis, the Efficient

Coding Hypothesis (ECH) states that nervous systems encode sensory input in as efficient a manner as possible through a process of reducing redundancy (e.g. ensuring any signals that do not provide additional information are discarded) and maximising the amount of information that single spikes can hold about a stimulus. Thus, the twin requirements of lossless compression (or rather low-loss compression!) and energy conservation can be achieved.

One of the main predictions of the ECH to be supported through experimental evidence is that neuronal response profiles reflect the prevailing stimulus levels within the environment (e.g. Laughlin 1981). Many neural cell types are not static, in terms of their sensitivity, and will adapt to the local stimulus conditions (e.g. Solomon, Peirce, Dhruv & Lennie 2004; Albrecht, Farrar & Hamilton 1984). The ability of cells to flexibly adjust the dynamic range of their output is especially important for luminance contrast, a stimulus attribute which can span a range of approximately 2 orders of magnitude within the visual scene.

In general, adaptation to low-level stimulus attributes, such as contrast or noise amplitude (which tend to have a Gaussian distribution), is believed to centre the firing threshold of a cell onto the global mean level of the predominant stimulus. The dynamic range of a cell will expand and contract to reflect the stimulus distribution. Changes that affect the overall response profile of a cell are generally classed as a form of gain control, the purpose of which is to ensure that the dynamic response range of a neuron encompasses the stimulus distribution. Much of the considerable body of work addressing adaptation has focused on contrast adaptation and gain control, in the visual system. However, a specific form of adaptation, to either the global mean or variance of a rapidly changing stimulus,

has been uncovered in the inferior colliculus of guinea pig auditory system (Dean, Harper & McAlpine 2005) and barrel cortex of the rat (Garcia-Lazaro, Ho, Nair & Schnupp 2007). The amplitude of the stimulus was refreshed every 40ms and drawn from a highly-skewed distribution. Such that a region of stimulus space (spanning a restricted range of stimuli) had an overall selection probability of 80% and the remainder of stimulus levels had 20% chance of selection. This gave the stimulus an overall global mean value that was determined by where in the stimulus space the high-probability region was located and an overall global amount of variance that was determined by the width of the high-probability region and increased with the width of the area. As such, mean and variance of the adapting stimulus could be manipulated independently and the adaptive responses to either mean or variance could be isolated.

Adaptation to stimulus variance was in the opposite direction to that expected by the ECH, i.e. the neural response profile was scaled upwards, thereby reducing the dynamic range of the cell, despite the presence of increased stimulus noise. Adaptation to global mean amplitude tended to affect neural threshold, thus on a superficial level, it was in line with Information theoretical principles. However, neural threshold was displaced, relative to the un-adapted response, to centre on inputs on the edge of the stimulus distribution, in contradiction to the expectations of the ECH. As such, the main inspiration for the work presented here was to examine the experimental data more thoroughly; with the aim of understanding the underlying strategies and mechanisms underpinning the changes in neural gain that were outside the predictions of the ECH.

Jan Schnupp generously provided the experimental data from the adaptation experiments in the barrel cortex. What was found was rather interesting, adaptation to

stimulus variance resulted in decreased neural gain, fully in line with the expectations of the ECH but nevertheless opposite to the findings of the published data. However, adaptation to global mean amplitude tended to prime neurons, not for detecting and discriminating the most common stimuli with the distribution, but for detecting large deviations in the input signal, occurring on the tens of millisecond scale.

In order to model the possible mechanisms underpinning adaptation to mean and variance, a computer simulation of a single-compartmental neuron was developed, using an implementation of Alain Destexhe's conductance-based integrate and fire computational neural model.

Adaptation to global mean was successfully replicated in the model cell through a combination of spike-rate adaptation and tonic inhibition. Adaptation to variance was modelled in several ways. Firstly, by implementing a model of gain control through the modulation of background firing rates (Chance, Abbott & Reyes 2002). This was achieved by increasing firing rates, of stochastic, excitatory and inhibitory background conductances, in proportion to the increase in stimulus variance. Secondly, via a model of asymmetric synaptic depression (Chelaru & Dragoi 2008) in which depression levels in both the inhibitory and excitatory populations that provided afferent excitation to the target cell were adjusted to reflect stimulus variance; such that as the amount of noise in the adapting stimulus increased depression levels became more balanced. This led to an increase in the amount of afferent excitation and inhibition, thereby reducing gain in the model cell. Finally, a model was produced which combined aspects of all the models reviewed in the thesis: the experimental data was fully replicated with a model of non-linear input in combination with recurrent synaptic depression.

The thesis is organised into six chapters, including this overview. Within the following chapter, the main concepts of Information theory and the Efficient Coding Hypothesis will be introduced. This will be followed by a review of contrast adaptation in the visual system. The purpose of this section is demonstrate the evidence for the ECH in early sensory processing as well as to introduce a few key concepts that underpin how gain control is described and quantified within neural systems, e.g. the parameters of the Naka-Rushton equation, which are used to fit cellular stimulus-response functions. Some extensions to the ECH are also introduced in the chapter, followed by some alternative coding strategies, namely novelty and deviation detection.

The third chapter will introduce the experimental data, as provided by Jan Schnupp. Here, the original and a related study will be reviewed, followed by the re-analysis of the data. It is in this chapter that the results regarding adaptation to stimulus variance will be presented. Namely, that in direct contradiction to the published results, adaptation, to increasing levels of stimulus variance, resulted in a decrease in gain, across the population. *There is also a brief section highlighting the difficulties in analysing temporal information* from the experimental data. Following on from the initial re-analysis, presented in Chapter 3, the fourth chapter of the thesis focuses on possible encoding strategies. Data is presented that suggests that neurons are adapted to detect large deviations in the ongoing stimulus.

The initial section presents the, equivocal, results of a data analysis that was based on the paradigm of stimulus-specific adaptation; with the aim of ascertaining whether adaptation to global mean levels improved novelty detection. Within the following sections the data is analysed in terms of stimulus history and context. Here, data relating neural output to difference in amplitude between the current and preceding stimulus (the step-size

function) is presented. Across the population, the adaptive response is scaled such that the step-size function, across all mean adapting conditions, is invariant.

Amongst both populations (adaptation to mean or variance) some cells show a preference for step-size irrespective of the direction of amplitude change and are classed as novelty-preferring cells. Novelty-preferring cells generally display longer time-to-peak and response-decay latencies than the remainder of the population.

The data was analysed in terms of sensitivity to local deviations from the global mean, at each time step. Across the population, neurons displayed a preference for positive, local deviations-from-the-mean; under conditions of adaptation to global-mean the responses were scaled as a function of the global standard deviation of the input signal.

The overall conclusion of this chapter is that adaptation to stimulus mean acts to rescale responses so as to maintain information transmission about local changes in stimulus amplitude, irrespective of global mean levels, and represents a distinct encoding strategy to that observed under conditions of adaptation to stimulus variance (in which neural responses adapt to reflect the expanded stimulus distribution). The fifth chapter presents results from computer simulation of adaptation to stimulus statistics. The introductory section to the chapter contains an in-depth review of some models of gain modulation as well as an introduction to spike-rate adaptation.

It will be demonstrated that the adaptive shift in the neural-stimulus response function and the invariance of the step size function can be reproduced by a model consisting of tonic inhibition and spike-rate adaptation.

Gain modulation is modelled in three ways:

1. by increasing the firing rates of background conductances in proportion to

stimulus variance (Chance et al. 2002) and

2. Adjusting levels of synaptic depression in recurrent excitatory and inhibitory units, such that the difference between afferent inhibition and excitation is reduced (Chelaru & Dragoi 2008).

3. Gain modulation was elicited by a combination of the models above, whereby non-linear afferent inputs were used to drive a recurrent population of excitatory and inhibitory I & F cells. The combination of recurrent depression acting within the network and feedforward depression, spike-rate adaptation and tonic inhibition acting on a read-out I & F cell elicited a decrease in neural gain as global variance increased, under various depression configurations.

All models replicate the adaptation to variance data, though the final model presents the most robust mechanism; gain modulation 'emerges' from the network and is independent of the balance of synaptic depression and other non-linearities acting on the membrane of the model read-out cell.

The final chapter is the discussion, which brings together the results from chapters 3-5. The main thrust of the argument within the discussion is that adaptation to amplitude of whisker deflection is an important aspect of sensory processing, even if it is not the most relevant feature of the stimulus, in terms of neuronal sensitivity. The results demonstrate that neurons change their coding strategies depending upon the overall levels of mean amplitude and variance in the sensory input. In conditions of low or constant stimulus variance, i.e. under the regime of adaptation to global mean, neurons are primed to respond to large and rapid changes in the stimulus; however, under conditions of increased noise, neural-gain functions broaden such that the cell is able to compute the full range of the

stimulus distribution.

Chapter 2: Efficient sensory coding and adaptation

2.1: Introduction

Neuroscience is a discipline that is fundamentally concerned with how nervous systems receive, analyze and respond to external stimuli. All afferent information, whether it be carried by light/sound waves, pressure or heat, is converted within the nervous system into electrical impulses. Thus, one of the most pertinent questions within neuroscience is how sensory information is encoded into electrical activity.

During the 1950-60's, Horace Barlow (Barlow 1959, Barlow 1961, Barlow 1972) and a number of other scientists (e.g. Attneave 1954; Craik 1948) turned to Information Theory as a possible mechanism for understanding neural encoding. Information Theory is a mathematical system for analyzing communication systems, developed by Claude Shannon (Shannon 1948; Shannon & Weaver 1949). According to Information Theory, messages travelling along any communication system can be broken down into basic units of information or 'bits', quantified by a binary code of 0s and 1s. Action potentials (spikes) can also be considered as binary signals, due to their all-or-nothing behaviour. Barlow proposed therefore that any information conveyed by spiking behaviour could also be quantified as 'bits' of information and that sensory input is encoded in such a way as to minimise the number of bits needed to transmit a signal through the neural system.

According to the Efficient Coding Hypothesis (Barlow 1961), neural response profiles should represent the inherent distribution of any incoming information in a maximally efficient manner. For example, if the predominant afferent stimulus has a Gaussian distribution, and incoming information is represented with a firing-rate code, the theory predicts that neural response profiles should be sigmoidal in nature, with the linear,

accelerating portion of the neural-response function exactly mapping the region under the curve of the stimulus distribution (please see Section 2.2.3, below). In general, most neurons have bio-physical limits that prevent the full encoding of external stimuli through a firing rate code alone, thus the Efficient Coding Hypothesis also predicts that cellular responsiveness is adaptive i.e. it is adjusted in concert with the distribution of incoming signals.

Adaptation, in response to contrast information within the visual stimulus, has been observed and widely studied in the visual system and found to be in line with the predictions of the efficient coding hypothesis.

An information theoretic account of sensory coding

As mentioned above, during the middle of the 20th century there was a move towards using information theory as a tool for explaining sensory encoding. In order to understand the context of this trend what follows is a description of the main principles of information theory, how it can be applied to nervous systems and an account of Barlow's efficient coding hypothesis (Barlow 1961).

Information theory

Information theory is a mathematical theory for analyzing communication systems (Shannon 1948; Shannon & Weaver 1949). It was developed to address the engineering problem of transmitting information along a noisy channel and provided a formalized framework to quantify the rate, reliability and capacity of message transmission within a communication system. A communication system can refer to anything that has a message or information source, a method of transmitting the message as a signal and a capacity to receive and output the message, see Figure 2.1, below.

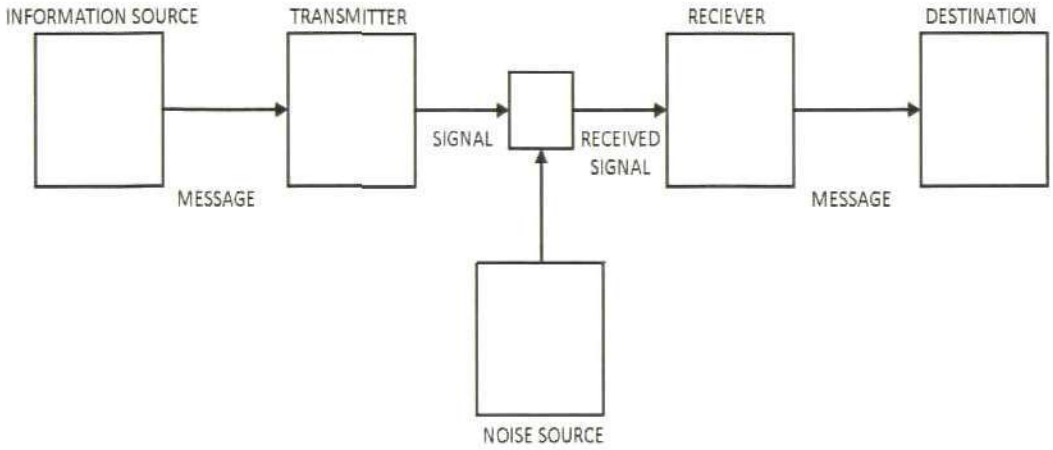

Figure 2.1: Schematic of a generalised communication system
The information source generates a message (e.g. text message to a friend), which is encoded into a signal by the transmitter (e.g. a mobile phone) and sent along the channel (e.g. a satellite or mobile phone transceivers), represented by the smallest box in the centre, where noise could be added. The receiver (e.g. the friend's mobile phone) receives the signal, decodes it and finally passes it on to the message destination (e.g. the receiving mobile phone's message inbox). Adapted from Shannon (1948).

In this sense, a communication system refers equally to arcane methods of signalling such as hilltop fire-beacons to modern day forms of communication, e.g. text messaging on mobile phones, as well as to natural languages in general. One of the motives behind developing information theory was to understand the limits of data compression and the need for robustness in a signal. Many forms of communication undergo data compression to a degree. For example, ancient beacon systems acted to compress messages, (such as: 'the Vikings have landed!') into the single symbol of a hilltop fire.

These forms of communication work because natural languages tend to be robust, i.e. can be understood even in the presence of a certain level of noise (such as spelling irregularities, abbreviations etc.), and are also highly redundant, in that messages can still be understood, even with parts of the signal missing (the missing parts are thus considered

redundant or unnecessary). For example, the English language has a redundancy of roughly 50% (when assuming a basic vocabulary of ≈ 850 words (Shannon 1948), such that the following text should still be legible, despite the fact that all vowels have been removed:

RD THS MSSG PLS, THNKY VRY MCH [1]

An excellent example of both the robustness and redundancy inherent in the English language is given in the SMS poem below. The poem was runner-up in a text-message poetry competition (organised by the company T-Mobile, 2007) to find a so-called `txt laureate':

> `O hart tht sorz, My luv adorz, He mAks me liv, He mAks me giv, Myslf 2 him, As my luv porz.' (Eileen Bridge, 2007)

In natural English, the poem reads: 'oh heart that soars, my love adores, he makes me live, he makes me give, myself to him, as my love pours'. Clearly, the English language lends itself well to the type of data compression needed when sending a message along a highly restrictive communication channel.

Obviously, the ideas outlined above are highly qualitative descriptions of some of the concepts underpinning Information Theory; Shannon developed strict mathematical methods (Shannon 1948) to quantify exactly the amount of data compression required to reliably and efficiently transmit a signal along a noisy channel.

Several of the key concepts of Information Theory (namely information, entropy, mutual information, channel capacity and redundancy), and their mathematical descriptions, are outlined below; with the assumption that one is discussing a noiseless communication system with a discrete set of inputs.

---

[1] READ THIS MESSAGE PLEASE, THANKYOU VERY MUCH

Information

The concept of information is generally taken to mean facts or knowledge. However, the compact Oxford English Dictionary also defines information as 'What is conveyed or represented by a particular sequence of symbols, impulses, etc'. For example, words within a given language represent semantic information using a specific sequence of symbols, i.e. the letters of the alphabet; or spiking behaviour within the nervous system is believed to represent sensory and motor information through either frequency or relative timing of action potentials.

In terms of Information Theory, it is useful to think of information in relation to uncertainty. If one knows exactly what response a system will generate then it follows that no new information is imparted through it; conversely if there is doubt over the output then the response of a system will clearly be informative. In any system, once the output has been received and understood, there is no more ambiguity, thus information can be considered as a reduction in uncertainty or an increase in certainty; the amount of randomness in a system determines how much information can be transmitted through it (see Section on Entropy, below).

The unit of information is the "bit". One bit measures the information contained in a single, binary (i.e. '0' or '1') event of equal probability, e.g. an unbiased coin toss. A binary code can transmit up to 1 bit/per symbol.

Finally, in discussing information it is important to stress that Shannon was not interested in quantifying the relevance or importance of information, but only in data compression and reliable transmission, thus, in information theory, the messages SOS (save our souls) and LOL (laugh out loud) are arguably equivalent.

Entropy

All communication systems are concerned with sending out and receiving messages. Each message is broken down into units of information, ready for transmission and reconstituted at the receiving end of the system. As described above, information is a reduction in uncertainty; the amount of potential information depends upon the degree of variability in the transmission and decoding of messages, which is quantified by the entropy of the system. For example, an arbitrary system in which the input and output are always the same (i.e. there is no uncertainty in the message) transmits no new information and, by extension, has zero entropy. In general, the inputs and outputs of a system are drawn from known distributions and the average or minimum number of bits needed to encode and decode messages is dependent upon the variability within these distributions. Entropy is determined by the number of symbols used to represent the input and output of a system and the probability of transmission.

Where an arbitrary system's signals are drawn from a distribution X with N possible values and all are equally likely to occur then the entropy (H) increases monotonically with N (see Figure 2.2, below) and is equal to:

$$H = -\log_2(1/N) \qquad \textbf{2.1}$$

As an example, one can consider a communication system in which the inputs are all drawn from a probability distribution X that has 4 possible values: A, B, C & D. If all possible values of X had an equal probability of being transmitted, H would equal 2 bits and any signal, drawn from X, could be represented with the following 2-bit binary code (see Table 2.1):

Table 2.1: Example of a 2-bit binary code

Symbol	Binary code	
A	0	0
B	0	1
C	1	0
D	1	1

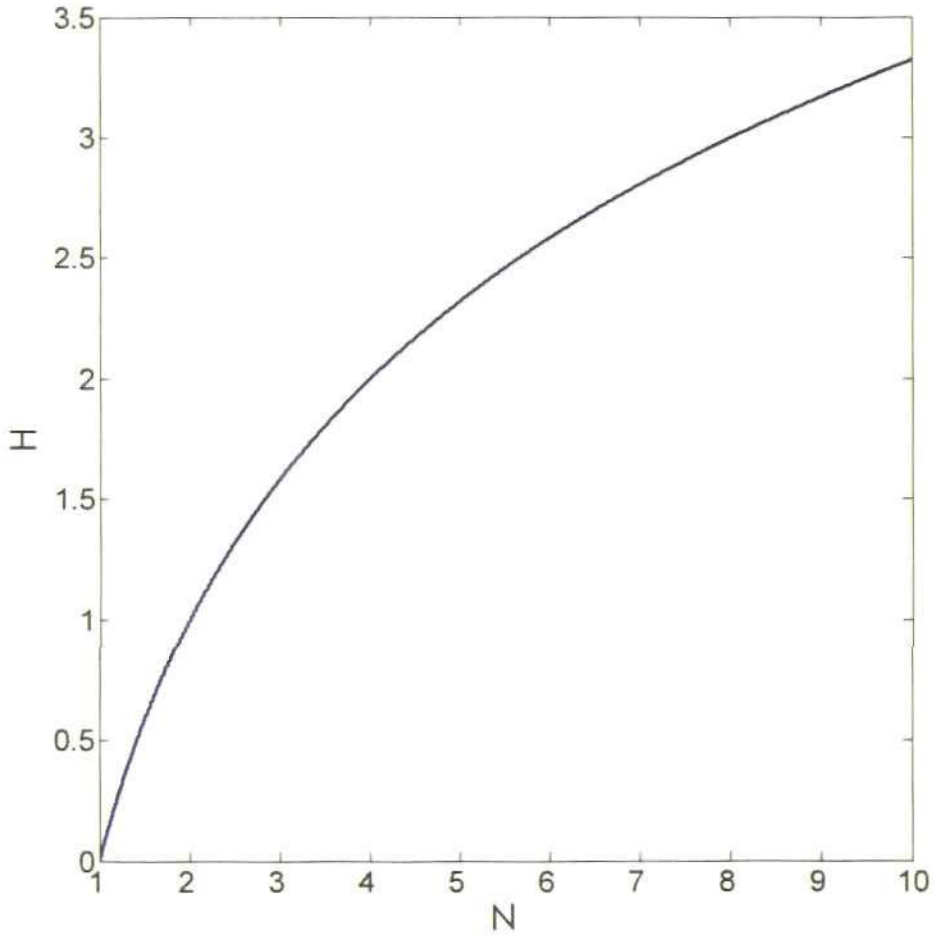

Figure 2.2: Entropy as function of the number of possible input symbols.
Entropy, H, is plotted as a function of N, the number of possible values that the signal distribution X can have and where all possible values of X ($X_{1 \ldots N}$), have equal probability of selection.

In the case where the 4 values of X occur with unequal probability (remembering, of course, that all probabilities add up to unity), H is given by:

$$H = -\sum_{i=1}^{n} p_i \log_b(p_i) \qquad \textbf{2.2}$$

Where H is the entropy of X, with a given number, n, of values $i = 1 \ldots n$ and p_i is the probability of observing the *ith* value of X; b is the base of the logarithm, which in Information Theory is always $\log_2$ as the information content for a binary event with $p = 0.5$ is equal to $\log_2(0.5) = 1$ bit.

For example, in the system outlined above, the p_i of the 4 symbols 'A, B, C, D' might be as follows:

$$p_A = \frac{2}{3},\ p_B = \frac{1}{6},\ p_C = \frac{1}{12},\ p_D = \frac{1}{12}$$

Which gives:

$$H = -\left[\frac{2}{3}log_2\left(\frac{2}{3}\right) + \frac{1}{6}log_2\left(\frac{1}{6}\right) + \frac{1}{12}log_2\left(\frac{1}{12}\right) + \frac{1}{12}log_2\left(\frac{1}{12}\right)\right]$$
$$= 1.4183$$

In this case, the average number of bits needed to encode each of the four symbols is 1.4183 bits; however some symbols will use more or less bits depending on how common they are, i.e.:

$$H_i = -\log_2(p_i) \qquad \textbf{2.3}$$

Therefore:

$$H_A = 0.5850,\ H_B = 2.585,\ H_C = H_D = 3.5850$$

$$H = \left[\frac{2}{3}H_A + \frac{1}{6}H_B + \frac{1}{12}H_C + \frac{1}{12}H_D\right]$$

$$H = 1.4183$$

As can be seen, the most common signal 'A' requires the least number of bits. In this system, no symbol is encoded using whole bits, thus it may be necessary to add bits to the system, for example:

$$H'_A = 1,\ H'_B = 3,\ H'_C = H'_D = 4$$

$$H = \left[\frac{2}{3}H'_A + \frac{1}{6}H'_B + \frac{1}{12}H'_C + \frac{1}{12}H'_D\right]$$

$$= 1.8333$$

Accordingly, H represents the minimum number of bits required for signal encoding, e.g. in the example given above, the signal could not be reliably encoded with less than an average of 1.4183 bits per symbol.

H has the following properties:

- $H = 0$, where there is no uncertainty, e.g. all probabilities, except one, are equal to zero, with the remaining probability equal to unity.
- H is maximal when all responses are equally probable (see Figure 2.3) and $H_{MAX} = log_2(n)$, where n is the number of output possibilities and $p_{1...n} = 1/n$.

Joint Entropy

One can define the joint entropy of a system in which values are drawn from the

joint distributions of two random variables, X & Y, with m & n possible values, respectively (e.g. a system in which the input is drawn from X and the output is drawn from Y). The joint entropy, $H(X,Y)$, is:

$$H(X,Y) = -\sum_{i,j=1}^{m,n} p(x_i, y_j) \log_2 p(x_i, y_i) \qquad \textbf{2.4}$$

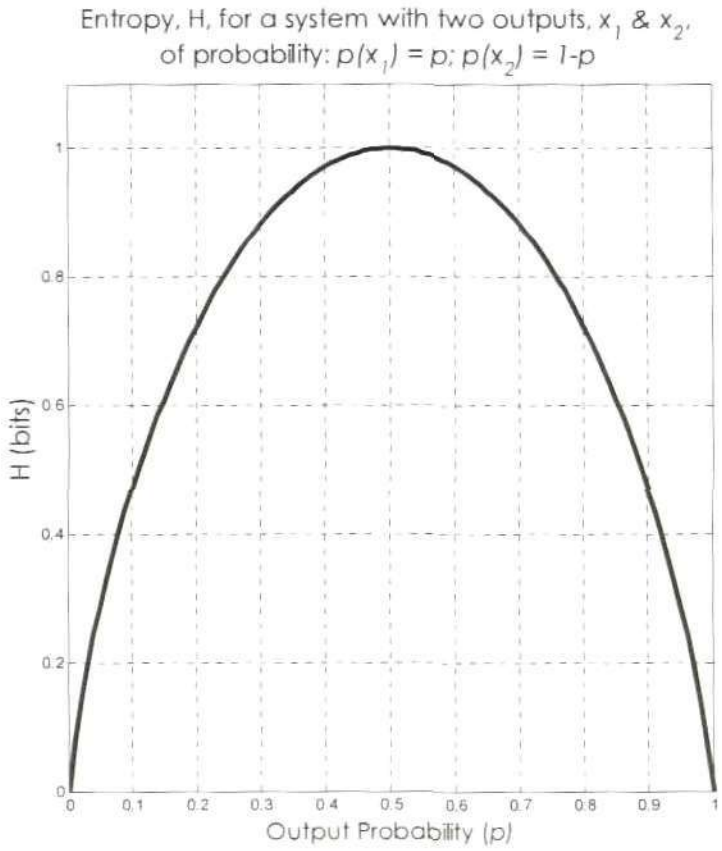

Figure 2.3: Entropy of an arbitrary system with two possible outputs
Entropy (H) of an arbitrary distribution p(X). X has two possible values only, with one having the probability p and the other being 1-p. H is maximal when the probability of selection for all values of X is equal (e.g. $p = 0.5$). Adapted from Shannon (1948).

So long as the individual entropies of X and Y are equal to:

$$H_X = -\sum_{i,j=1}^{m,n} p(x_i, y_j) \log_2 \sum_{j=1}^{n} p(x_i, y_j)$$

$$H_Y = -\sum_{i,j=1}^{m,n} p(x_i, y_j) \log_2 \sum_{i=1}^{m} p(x_i, y_j)$$

2.5

The joint entropy of X and Y will always be less-than or equal-to the sum of the entropy of the individual distributions, with equality only being observed when X & Y are independent:

$$H(X,Y) \leq H_X + H_Y$$

2.6

Conditional Entropy

If the distributions of X and Y are dependent on each other, it is useful to define the conditional entropy of Y given X, $H_X(Y)$:

$$H_X(Y) = -\sum_{i,j=1}^{m,n} p(x_i, y_j) \log_2 p_i(y_j)$$

2.7

The conditional probability, $p_i(y_j)$, that y is the *jth* value of Y when x is the *ith* value of X is given by:

$$p_i(y_j) = \frac{p(x_i, y_j)}{\sum_{j=1}^{n} p(x_i, y_j)}$$

2.8

The conditional probability has an alternative notation, which is widely used in the

literature:

$$p(y_j \mid x_i) = \frac{p(x_i, y_j)}{\sum_{j=1}^{n} p(x_i, y_j)}$$

2.9

The conditional entropy, $H_x(Y)$, can never be greater than the entropy of Y, $H(Y)$, i.e. knowledge of X can never increase the uncertainty of Y but only decrease it:

$$H(Y) \geq H_x(Y)$$

2.10

If the random variables X and Y are independent, X contains no information about Y, such that the conditional entropy must take its maximal value, $H(Y)$. If Y is completely dependent upon X then the uncertainty of Y disappears once X is known and all the entropy in Y is explained by $H(X)$; the conditional probability is equal to unity and $\log_2(1) = 0$. Equations 2.4 - 2.8 are derived from Shannon (1948).

Mutual information

Conditional and joint entropy could refer to two separate communication systems in which the signals are drawn from X *and* Y, respectively, or could refer to a single system with an input drawn from X and output derived from Y. In this case one is observing an input/output relationship. One can quantify the amount of mutual information, $I(X \mid Y)$, that exists in the relationship between X *and* Y, i.e. how much can Y tell us about X, when Y is known.

The mutual information of X and Y is given by:

$$I(X \mid Y) = \sum_{i=1}^{m} \sum_{j=1}^{n} p(x_i, y_j) \log_b \frac{p(x_i, y_j)}{p(x_i) p(y_j)}$$

2.11

Mutual information is also equivalent to the entropy of the input minus the remaining uncertainty of the input, once the output is observed, i.e. the conditional entropy of X given Y, $H_Y(X)$ (note this is opposite to the conditional entropy described by Equation 2.7, in the section above, which defined the entropy of Y, given X); the mutual dependence of X and Y quantifies the rate of information transmission (see Section on *Channel capacity*, below):

$$I(X \mid Y) = H_X - H_Y(X)$$

2.12

Mutual information is a useful measure of the reliability of information transmission when dealing with a noisy channel. Shannon (1948) gives the example of a noisy channel with an input x that is either 0 or 1 and transmitting at a rate of 1000 bits per second ($H(x) = 1 bit / symbol : p_0 = p_1 = 0.5$). Due to noise there is, on average, a 1% chance of the output, y, being in error. Thus if $y = 0$, there is a 99% chance that $x = 0$ and 1% chance that $x = 1$. One can calculate the conditional entropy of x given y, or $Hy(x)$, and the mutual information of x and y, $I(x \mid y)$, as:

$$H_y(x) = -[(0.99 \log_2(0.99)) + (0.01 \log_2(0.01))]$$
$$= 0.081$$

And:

$$I(x \mid y) = 1 - 0.081 = 0.919$$

The rate of error is then 81 bits/s and the rate of transmission is 1000-81 = 919 bits/s.

In the noiseless case there is no uncertainty, i.e. y is always equal to x, thus the *conditional entropy* $Hy(x)$ *is zero and* $I(x \mid y) = Hx$ *; in the case of a truly random system,* where there is a 50% chance of error, x and y are wholly independent and hold no mutual information about each other. For example, in the case above where $Hx = 1$ and $p0 = p1 = 0.5$, a 50% chance of error in y will give:

$$H_y(x) = -\left[(0.5 \cdot \log_2(0.5)) + (0.5 \cdot \log_2(0.5))\right] = 1$$

And:

$$I(x \mid y) = 1 - 1 = 0$$

Channel capacity

Channel capacity refers to the amount of information it is possible to transmit along a channel in a given unit of time and is equal to the maximum of the mutual information, where you have input X and output Y:

$$I_{MAX}(X \mid Y) \qquad \textbf{2.13}$$

Redundancy

In terms of information theory, redundancy either represents information that is either already known to the system or information that's unnecessary for the communication, e.g. vowels in natural English, which can often be inferred from the context. For example, in the hypothetical system outlined in the section on *Entropy*, above, it was stated that a useful code may have required the addition of extra bits of information.

These extra bits could be regarded as redundant information. Thus redundancy can be considered as the difference between the number of bits used to actually transmit a signal and the entropy of the message distribution, $p(X)$, or as unity minus the relative entropy.

Relative entropy, $D(p(X) \mid q(X))$, is the ratio of the true distribution of X, $p(X)$, to the actual distribution used to describe X, $q(X)$, and is often termed the Kullback-Leibler Distance (Cover & Thomas 1991).

The nervous system is a communication system

The nervous system of any organism acts to transform external information into an internal representation of the outside world and, as such, can be considered as a dense network of communication systems, acting in series and in parallel. For example, within the proprioceptive system, sensory receptors transduce external stimuli into graded potentials (an information source) that are rapidly converted (via synaptic interactions) at the soma of the somato-sensory relay neuron (the transmitter) into action potentials. These are propagated along the axon (the channel) to the spinal chord junction where post-synaptic neurons receive the sensory information, reconstituted from neurotransmitter release[2]; here the signal may be propagated further along into higher regions of the central nervous system using a communication system of pre-synaptic activity, somatic action potential generation, propagation and neurotransmitter release.

Barlow (1961) argued then that, as a communication system, the nervous system could be described using information theoretic principles, given three simplifying

---

[2] One could conjecture that the transmission of action potentials at the axon terminals, across the synaptic cleft, into post-synaptic dendritic activity is a separate/sub-system, nevertheless neural activity is essentially concerned with transmission of information and can thus be reduced into a communication system as defined by Shannon, see Figure 2.1, above

assumptions:

1. It is a noiseless system with discrete inputs

2. Action potentials are discrete signals and, due to their highly stereotyped behaviour, can be considered as binary signals

3. The capacity of a channel, *C* (channel here refers to a neural pathway, e.g. the optic nerve), depends on the number of fibres, *F*, within the channel, the number of discrete time intervals, per second, *R*, that an impulse can be sent along C and the average number of actual impulses, per second, I.

Nervous systems are subject to the same issues of compression as more traditional communication systems, in that a large range of possible messages (one could argue that the number of potential environmental inputs approaches infinity!) must be converted into a smaller number of signals for transmission. Compression is of vital importance for nervous systems as there can never be enough receptors to encode all incoming information; furthermore every spike that is transmitted costs metabolic energy. Barlow proposed that the nervous system acts to transmit information in a maximally efficient manner as possible. Reducing the redundancy in the encoding of sensory stimuli means that afferent input can be highly compressed, without loss of information, thus minimising the number of spikes needed to transmit a given signal through the nervous system. In terms of Information theory redundancy reduction is achieved through the maximisation of the entropy in the system.

Barlow's efficient coding hypothesis

The Efficient Coding Hypothesis posits that nervous systems have evolved to transmit sensory information with as few spikes as possible, i.e. to map stimulus to

response in a maximally efficient manner. To this end, it might be assumed that any redundant information must be discarded early on in the system. However, it is important to note that the Efficient Coding Hypothesis is not focused on removing extraneous information *per se*; rather it focuses on recoding incoming stimuli in such a way that the full range of sensory input can be encoded, but as economically as possible. An efficient coding strategy, as envisioned by Barlow, has the advantage of being reversible (as information isn't lost); thus redundant information can still be accessed during higher stages of cortical processing.

Redundancy reduction is achieved through maximising the relative entropy of the system, which is a function of the average entropy and channel capacity of the system, see below.

If the input to the nervous system is an ensemble of messages, m, that are of variable length (T_m) and independent, one can calculate the entropy, H_m, as for Equation 2.1:

$$H_m = -\log_b(p_m) \qquad \textbf{2.14}$$

Where p_m is the probability of m and b is the base of the logarithm and always equal to 2.

Therefore the average entropy across all messages (H_{AVE}) is (from Equation 2.2):

$$p_m \log_b(p_m) \qquad \textbf{2.15}$$

The rate of information flow is equal to H_{AVE}/T where T is the weighted sum of all message durations:

$$p_m T_m \qquad \textbf{2.16}$$

The channel capacity (C) is equal to the maximum amount of information, per unit of time, which travels along a channel, with constraints outlined in Assumption 3, given above.

As such C, for firing-rate codes, can be calculated as follows (Mackay & McCullough, 1952):

$$C = -FR\left(\frac{I}{R}\log_b \frac{I}{R} + \left(1-\frac{I}{R}\right)\log_b\left(1-\frac{I}{R}\right)\right) \qquad \textbf{2.17}$$

For F, R and I see Assumption 3, Section: 2.2.3, above.

The relative entropy (H_{REL}) of the message ensemble in relation to channel capacity is given by the ratio:

$$H_{REL} = H_{AVE}/C\,T \qquad \textbf{2.18}$$

The equations 2.14-2.18 are adapted from Barlow (1961).

Thus the redundancy is equal to 1-H_{REL}, i.e. the maximum rate of information transfer minus the actual rate of information flow across the channel. Clearly then, the ideal solution for redundancy reduction is to maximise H_{REL}.

H_{REL} is related to the rate of information flow, thus Barlow (1961) intuited that neurons encode external stimuli by firing rate over time, as the amount of information carried by average number of spikes per given unit of time can be quantified in bits due to the binary nature of action potentials. An example then of an efficient coding strategy would be a factorial code (Barlow 1961), whereby the most common stimuli are encoded using the lowest firing rate possible, with firing rate increasing proportionally with input rarity, given the limits of channel capacity. The type of encoding found in natural English, for example (Shannon 1948), can be considered factorial as there is a clear relationship between word length and rarity (i.e. commonly used words, such as "a", "the" etc., tend to be very short).

However, in a system in which stimuli are distinguished on the basis of firing rate codes a factorial code may not be practical, in terms of discriminability. For example, an

input with a Gaussian distribution has two tails of equal rarity; in a factorial code these stimuli would be represented with similar, low levels of activity despite the large distance between.

Nervous systems should ideally have encoding strategies that strive to equalise output probabilities as this will maximise the entropy of the system (Shannon 1948; Barlow 1959; Laughlin & Hardie 1978). For example, when relating the output of a neuron to a single stimulus parameter (e.g. levels of contrast in a visual scene, see Section: 2.3) with a Gaussian distribution, a maximally efficient mapping of stimulus to response should result in a neural rate function that is the integral of the stimulus distribution, i.e. the cumulative probability distribution function of the input (see Figure 2.4); the cumulative probability distribution function maps equal changes in response ranges with equal shifts in stimulus probability. Thus for rare stimuli, more of the stimulus space will be mapped onto a given unit of response relative to more common stimuli (compare highlighted areas in Figure 2.4, below: the step-change in response is equal for both areas, under the stimulus-probability curve, bounded by either the red or green lines but the amount of stimulus space is relatively greater for the area bounded by the red lines, with respect to the green lines), or:

"...economies could be effected by reducing the space in the sensory representation occupied by those familiar stimuli and allowing more space for the infrequent and unexpected stimuli" (Barlow 1959, p11).

This type of coding strategy has the advantage, over a purely factorial code, of utilising restricted resources efficiently by amplifying the responses to, and improving resolution of, more common stimuli (with respect to infrequent events) whilst minimising the number of possible responses states needed to encode rare stimuli. It is equivalent to histogram equalization (Gonzalez & Wintz 1977), a procedure used in digital image

processing, see Figure 2.4.

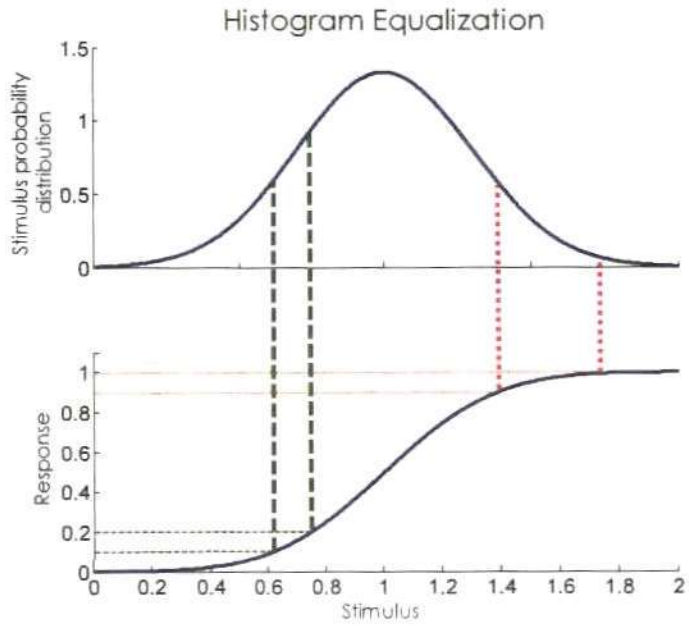

Figure 2.4: Histogram equalization and the neural code
Top panel is an arbitrary stimulus with a Gaussian distribution ($\mu=1$, $\sigma=0.3$), bottom panel is an idealized neural response function that is the cumulative probability distribution of the stimulus distribution. Thin lines define the upper/lower limits of a response range (equal to 1/10 of the maximum response); thick lines define the limits of stimulus representation for the associated response range (corresponding colour and line style). Response ranges delineated by red (dotted line) and green (dot-dash line) are equal; as are the areas under the curve bounded by the corresponding thick lines. However the representation of stimulus space is not equal, allowing for a greater resolution of common stimuli and minimising the number of responses that encode rare stimuli.

Accordingly, both information theory and the efficient coding hypothesis predict that the response profile of a sensory neuron will be highly dependent on the prevailing stimulus distribution (see Figure 2.5, below).

For example, given a stimulus with a Gaussian distribution, the cellular stimulus threshold (determined by the stimulus that elicits the half-maximal response or S_{50}) and linear region of the response function should be centred on the mean stimulus amplitude (Figure 2.5, left column); whilst the linear slope of the response profile is dependent on the

width of the stimulus distribution or stimulus variance (Figure 2.5, right column).

One would expect then that organisms have evolved with sensory neurons that are tuned to the prevailing environmental stimulus distributions and this does appear to be the case (see Section: 2.3); however organisms rarely exist in static environments.

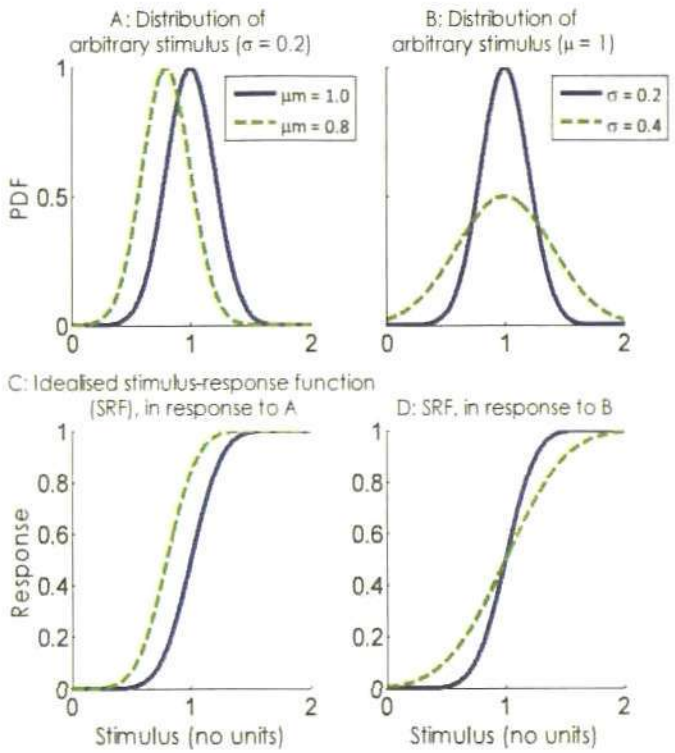

Figure 2.5: Predictions of the efficient coding hypothesis
Top row: Probability distribution functions (PDF) (solid lines) of an arbitrary stimulus with a Gaussian distribution. **Bottom row:** Response profiles of a hypothetical cell, assuming an optimal mapping of stimulus to response, given the corresponding (colour coded) PDF of the stimulus, each panel is an idealized neural response function (normalised response with no units) that is the cumulative probability distribution of the stimulus distribution. **Left column:** Dependency of response profile on the mean stimulus, blue (control): $\mu = 1$; red $\mu = 0.8$, σ is fixed at 0.2. The decrease in μ results in a lateral displacement of the stimulus-response function, (green, dashed line) relative to the control (blue, solid line). **Right column:** Dependency of response profile on the standard deviation of the stimulus, blue: $\sigma = 0.2$, green: $\sigma = 0.4$, $\mu = 1$. The increase in σ results in a decrease in the stimulus-response function gradient (green, dashed line) relative to the control (blue, solid line).

Thus the Efficient Coding Hypothesis also predicts that neural responses must be

flexible and that sensory systems will have mechanisms that allow for the adjusting of responsiveness to reflect transient changes in the prevailing stimulus statistics.

Before moving on to discussing experimental evidence for the Efficient Coding Hypothesis, one further point needs to be raised. The above discussion considers the case of single cells or channels of communication. However neurons do not exist in isolation, thus the Efficient Coding Hypothesis also holds that, within a population, neuronal responses should be independent of each other, such that the information carried by individual neurons is not redundant with respect to the information carried by the rest of the population.

2.2: Contrast adaptation in the visual system

The response of the visual system to contrast has been one of the most widely explored aspects of visual processing. Many visual cells (in both mammals and insects) have been shown to have highly non-linear responses to contrast, that are suggestive of an encoding strategy based on spiking efficiency and redundancy reduction, i.e. in line with the predictions of the efficient coding hypothesis.

2.2.1: Contrast-response function

Contrast is a ubiquitous property of the visual scene and describes the modulation of illumination, around mean intensity, of a given stimulus. It is often calculated using Michelson contrast:

$$c = \frac{L_{MAX} - L_{MIN}}{L_{MAX} + L_{MIN}} \qquad \textbf{2.19}$$

Where L_{MAX} and L_{MIN} describe the maximum and minimum luminance, respectively.

In order to cope with the omnipresent nature of contrast it might be expected that visual neurons encode contrast information in an efficient manner, i.e. with response profiles that reflect the inherent distribution of contrast within natural scenes. To test whether contrast encoding is in line with the efficient coding hypothesis, Laughlin (1981) compared the contrast-response of large mono-polar cells (LMC) in the compound eye of the blowfly (the arthropod equivalent of the retinal bipolar cell in vertebrates) with the distribution of contrasts found in the natural habitats of the blowfly, such as woodland and lakesides. LMC responses were recorded to stimuli drawn from the range of contrasts found in the wild and compared to the cumulative distribution function of the stimulus, see Figure 2.6, below.

One can see that the cellular responses closely matched the cumulative distribution function; thus there is a strong case to argue that contrast is encoded, in the blowfly, in the manner predicted by the efficient coding hypothesis.

Within the primary visual cortex (V1) of both cats (Ohzawa, Sclar & Freeman 1982; Ohzawa, Sclar & Freeman 1985; Albrecht & Hamilton 1982) and monkeys (Sclar, Lennie & DePriest 1989; Albrecht & Hamilton 1982), the majority of neurons have contrast-response functions (the average firing-rate response as a function of stimulus contrast) that are sigmoidal in nature. Figure 2.7, below, gives several examples of contrast-response function for simple and complex cells in cat or primate V1. Two-thirds of the cells in Figure 2.7 have contrast-response functions that increase for a limited subset of contrasts, saturate at high contrasts and are best fitted by a hyper-ratio function, when plotted on contrast-log coordinates.

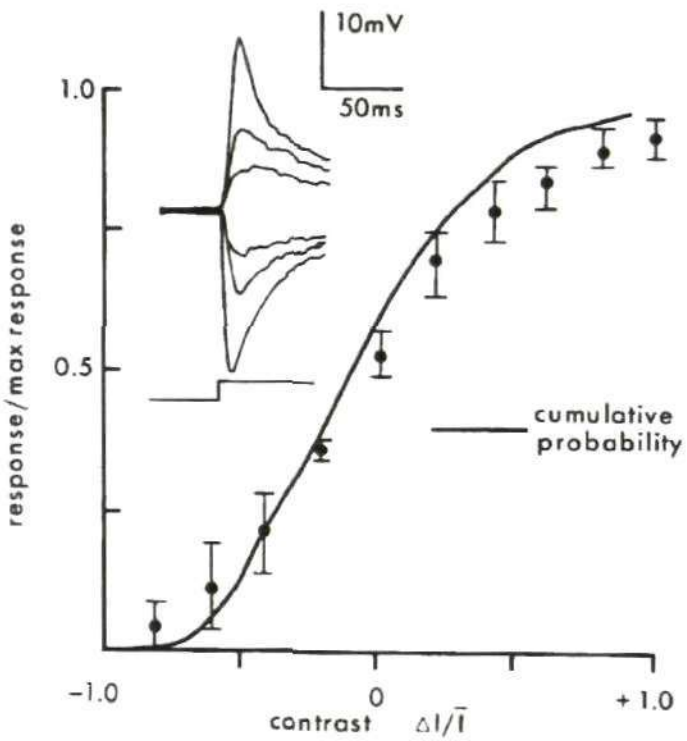

Figure 2.6: Contrast response in the compound eye of the blowfly
Comparison of distribution of natural contrasts and large mono-polar cells (LMC) contrast responses in compound eye of the blowfly. Solid line plots the cumulative probability distribution of contrast in the blowfly's natural environment, circles give the averaged contrast-response across LMC (n=6), with error bars giving total scatter of the responses across cells. The responses are normalized to set the maximum response, to an increase in contrast, at 1 and the maximum response, to a decrease in contrast, at 0. The inset plots the membrane response of a cell to a contrast step stimulus (hyperpolarisation to increments and depolarisation to decrements of stimulus contrast). Taken from Laughlin (1981).

The hyper-ratio function or Michaelis-Menton/Naka-Rushton equation (Albrecht & Hamilton 1982, Maffei & Fiorentini 1973) is given below:

$$r(c) = R_{MAX}\,[c^n / (c^n + c_{50}^n)] \qquad \textbf{2.20}$$

Where, R_{MAX} represents the maximum firing rate/response of the cell and is termed the saturation constant; n, the expansive response-exponent (Albrecht & Geisler 1994), governs the slope of the dynamic response at low contrasts (the average value, in the

cortex, of the power function is approximately 2.5 (Albrecht & Hamilton 1982) and C50, the half-maximal response or semi-saturation constant, describes the contrast sensitivity of the cell.

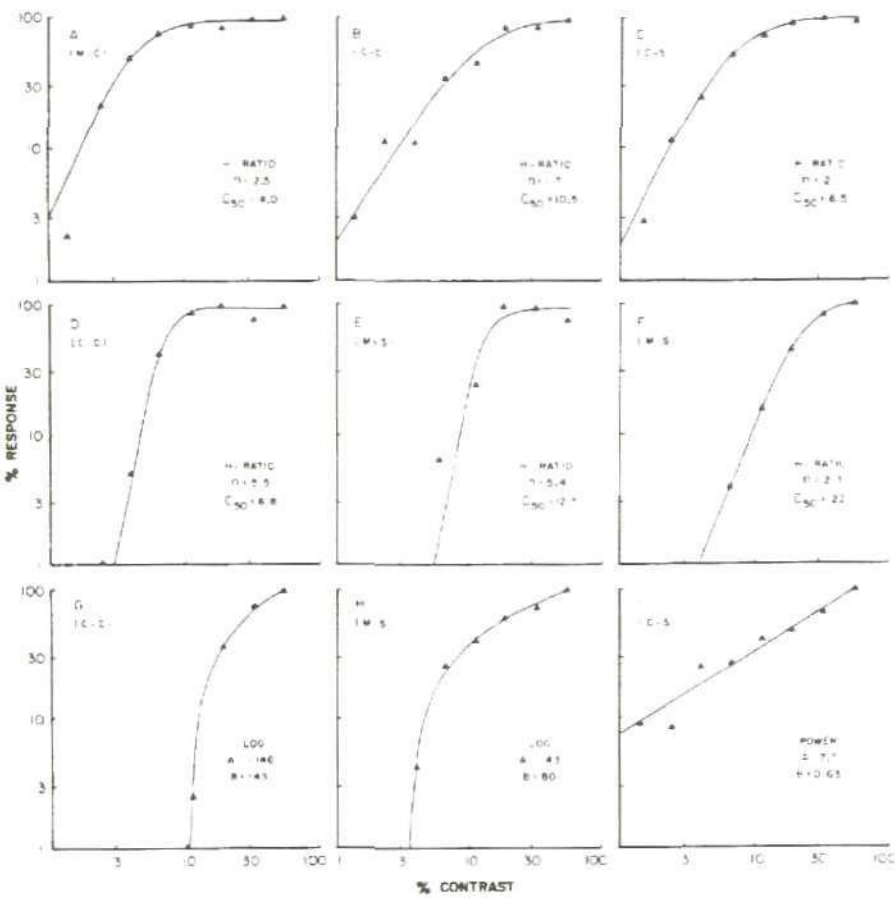

Figure 2.7: Contrast response functions in cat and monkey V1
Examples of contrast-response function in V1 of cat and monkey (C and M, respectively, see the 1st bracketed letter, on each contrast-response function plot, for species type) of both simple and complex cells (S and C, respectively, 2nd bracketed letter for cell type). Triangles give the actual percentage response; smooth lines are the best fit to the data points, with function type and parameters given on each plot (Taken from Albrecht & Hamilton 1982)

For primate V1 cells, the contrast-response function is fully developed within the first 10-20 ms after response onset (Albrecht, Geisler, Frazor & Crane 2002). Figure 2.9 A&D, below, shows the post-stimulus time histogram (PSTH) for a monkey V1 simple cell

in response to flashed (200ms presentation time) grating stimuli of fixed and optimal orientation, spatial frequency etc, at different contrasts (Figure 2.9A plots the actual PSTH whilst Figure 2.9C plots the shifted PSTH, such that time-to peak latency matched).

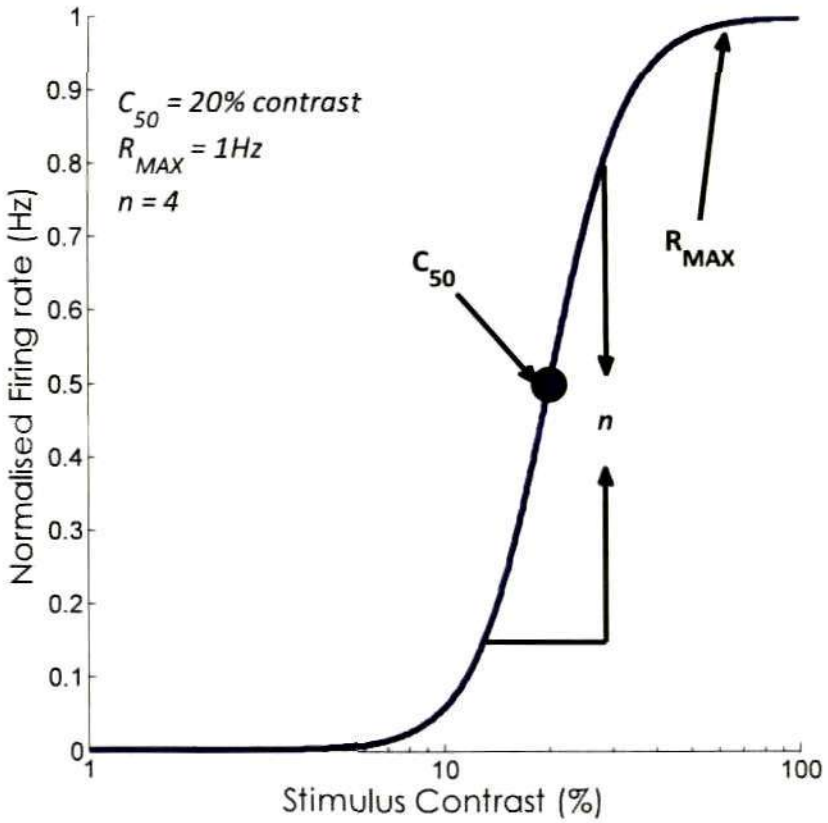

Figure 2.8: Idealized contrast-response function of a hypothetical cell
The idealized contrast-response function of a hypothetical cell showing response saturation for high contrasts and accelerating non-linear response at low-mid contrasts. R_{MAX}, C_{50} and n are the Naka-Rushton parameters, see Equation 2.20 and text for details.

As would be expected, activity increased and latency decreased for higher levels of contrast.

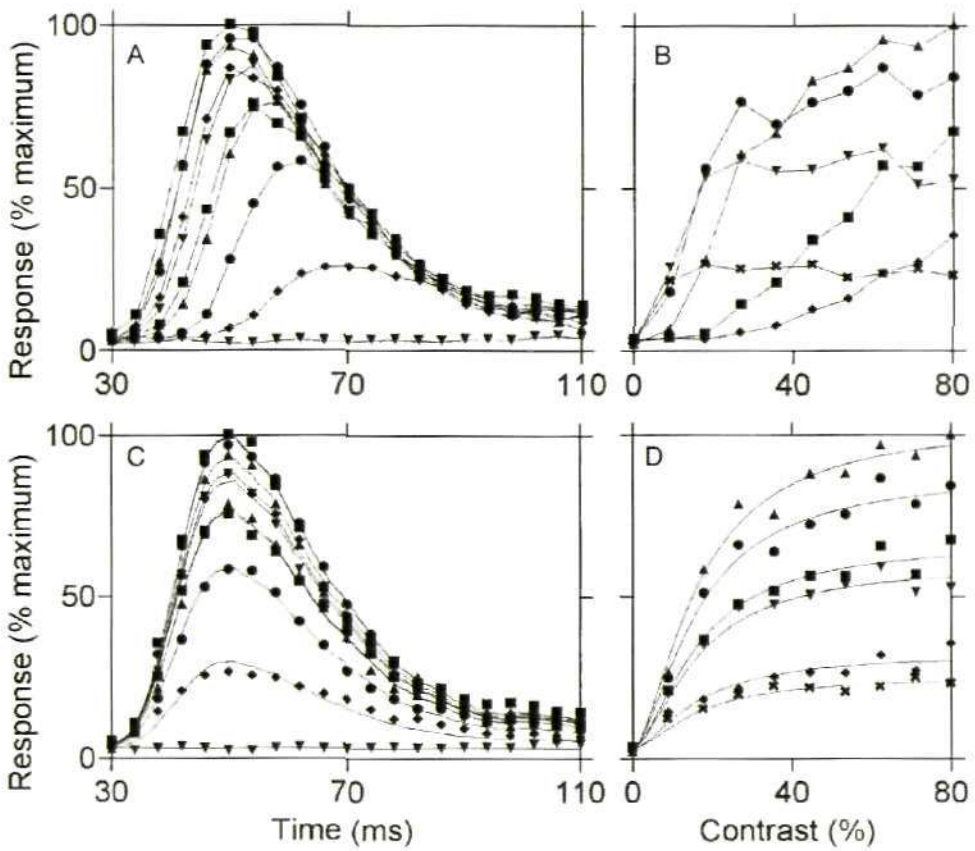

Figure 2.9: Contrast-response function in response to transient stimuli
A: Solid lines and symbols are the PSTH of response at different contrasts (diamonds-squares, contrasts ranging from 0.0-0.9 in incremental steps of 0.1) for a monkey V1 simple cell. **B**: contrast-response function from different time bins in A; 58ms (-diamonds), 62 ms (-square), 70 ms (-upward triangle), 78 ms (-circle), 86 ms (-downward triangle), 102 ms (-cross). **C**: PSTH, from A, but shifted laterally to remove latency differences and allow for comparison across all contrasts. Smooth lines are the average PSTH but scaled vertically and horizontally to account for contrast dependent amplitude and latency differences. **D**: contrast-response function at different time bins in C (symbols as in B). Smooth lines are the best fit of a Naka-Rushton equation in which only R_{MAX} is allowed to vary (C_{50} and n are 0.296 and 3.1, respectively). Taken from Albrecht et al. (2002).

The responses for a given time bin, as a function of contrast, are plotted in Figure 2.9 B&D (Figure 2.9B plots contrast-response functions for the 58,62,70,78,86, & 102ms time-bins from the PSTH plotted in Figure 2.9A, whilst Figure 2.9D plots the contrast-response functions, from the same time-bins, as derived from Figure 2.9C). It is clear from

Figure 2.9D that the contrast-response function is best fitted by a sigmoid function, even at the 58ms time-bin. The short latency of contrast-response function development indicates that this is an extremely rapid form of gain control, which can act within the space of a single saccadic fixation.

Are V1 contrast-response functions optimal, in terms of the distribution of contrasts within the environment? The levels of contrast within images of natural scenes can be measured on a pixel-by-pixel basis by determining the contrast-response of a filter, which resembles neuronal receptive field structure (e.g. a Gabor filter would mimic the receptive field of V1 simple cells, Brady & Field, 2000), to each pixel in the scene, thus quantifying the distribution of contrast in the natural world, as 'seen' by sensory neurons (Brady & Field, 2000, Tadmor & Tolhurst, 2000).

The contrast distribution of natural scenes was estimated in terms of V1 cell responses (Brady & Field, 2000). The average contrast of 46 grey-scale, logarithmically transformed, images of natural scenes (as calculated using Gabor filters of varying spatial frequency, orientation and phase to mimic cortical receptive fields) is plotted in Figure 2.10a in order of mean contrast (from low-high); the responses of the sensory filters were calibrated relative to their response to a sine-wave grating of equal spatial frequency, phase and orientation. There was considerable variety in average contrast levels across the images used for analysis, however the distribution of average contrasts across all scenes was approximately normal (Figure 2.10b).

The histogram of responses of the sensory filters, as a function of contrast (in 1% contrast bins) is plotted in Figure 2.11a, below.

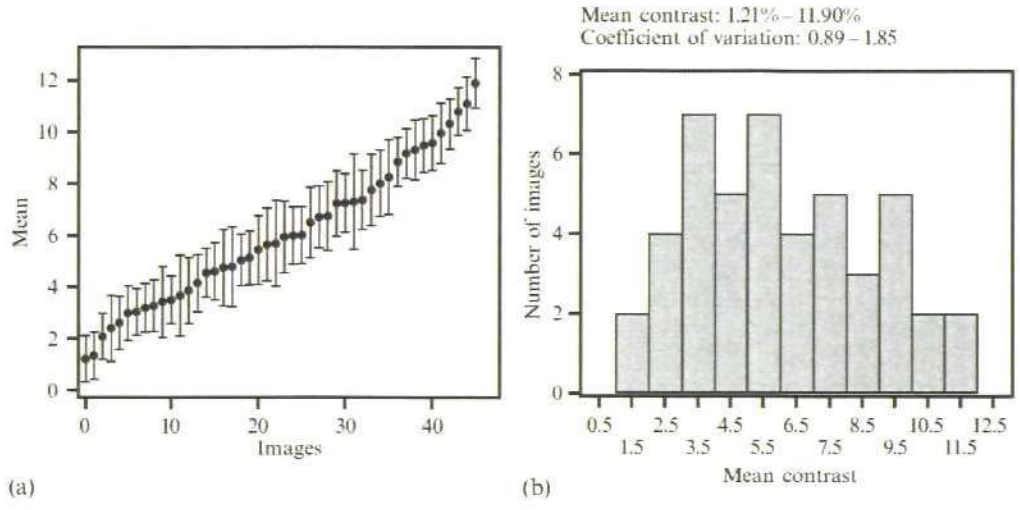

Figure 2.10: Levels and distribution of the average contrast observed in a series of 46 images of natural scenes
Plot (a): The average contrast level for each of the 46 images is plotted in order of increasing mean, ranging from 1.21-11.9%; the error bars show ±1 coefficient of variation. Plot (b): The distribution of mean contrasts is approximately normal. The figure is taken from Brady & Field (2000).

The distribution peaks for very low contrasts and has a high kurtosis (or 'peakyness'). The integral of the distribution of positive contrasts (see Figure 2.11b) is plotted in Figure 2.11c (solid line) and is best fit by a hyper-ratio function, with values similar to that observed in both cat and monkey V1 (for example see Table 2.2, in the Section: *Contrast dependent non-linearities in LGN and retina*, below). From this it would appear that the hyper-ratio function is optimal in terms of encoding environmental contrast. The distribution of semi-saturation contrasts observed in V1 (ranging from 1 to ≈ 40% contrast (Albrecht & Hamilton, 1982), see also Figure 2.7) could also allow for the efficient encoding of the wide range of average contrasts observed in the environment (Brady & Field, 2000), though this can also be achieved by shifting the dynamic ranges of response functions, depending on the prevailing levels of contrast in the input (see Section: 2.2.2:-2.2.3:, below).

A similar analysis was performed for retinal ganglion and LGN cells (Tadmor & Tolhurst, 2000); again the contrast-response functions of thalamic and retinal visual cells were found to match the distribution of contrasts encountered in natural scenes.

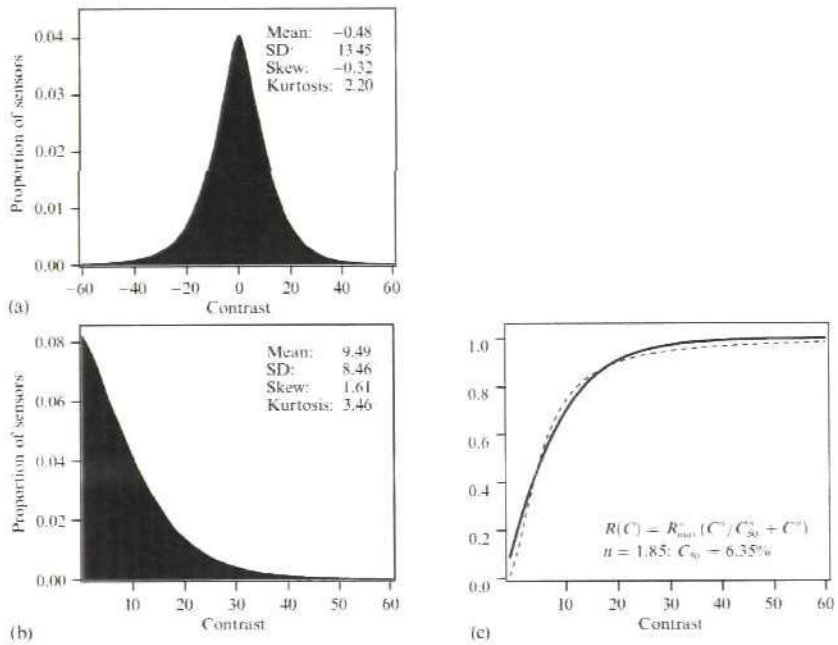

Figure 2.11: Distribution of contrast as 'seen' by V1 simple cells and the optimal response profile
Plot (a) gives the distribution of contrast in natural scenes, by binning the proportion of sensory filters that responded to each contrast; plot (b) is the distribution in (a) but for positive contrasts alone. The sensory filters were averaged across 5 spatial frequencies, and 4 orientations. Plot (c) is the integral (solid line) of the positive-contrast distribution in (b) and the dashed line is the best-fit to data: a hyperbolic-ratio function (R_{MAX} = 1:0). The figure is taken from Brady & Field (2000).

It should be noted that not all V1 cells display sigmoidal contrast-response functions. For example, in Figure 2.7g-h the relationship between contrast and response was a logarithmic function. The optimal output distribution, in terms of maximising entropy, depends on the constraints placed upon the system. When neural output in both feline V1 and macaque IT cortex is constrained by either the maintenance of long term,

average firing rates or minimising the average response, then the response distribution is an exponential function (Baddeley, Abbott, Booth, Sengpiel, Freeman, Wakeman & Rolls 1997; Rieke, Warland, de Ruyter van Steveninck & Bialek 1996), as opposed to a uniform output-distribution when the neuron is constrained by either response saturation or range (Laughlin 1981; see Figure 2.4 & Figure 2.6).

2.2.2: Gain control of the contrast-response function

V1 cells' contrast-response functions are not fixed, that is they are dynamic and are dependent on both stimulus attributes and history. Changes in the output, or gain, of a cell can lead to shifts in contrast sensitivity, response amplitude, and alter the dynamic signalling range of a cell.

Two forms of gain control, response acceleration and saturation, have already been outlined, above. However, the contrast-response function, as fitted with the Naka-Rushton equation, has three free parameters which are often used to quantify gain changes, relative to baseline levels (see Figure 2.12, below). For example, if the contrast-response function, CRF, is conserved but displaced laterally, compared to the control response (solid line in Figure 2.12, below), then only C_{50} is affected (dotted line; Figure 2.12, below); the resulting shift in contrast sensitivity is considered to be an example of contrast-gain control. Response-gain control is characterized by a vertical scaling of the contrast-response function (dashed line; Figure 2.12, below), in which the firing rate threshold that determines response saturation is increased or decreased (R_{MAX} in Naka-Rushton equation); C_{50} and n remain constant. Gain-adjustments that broaden or narrow the dynamic response range of the cell, a form of neural-gain control, are quantified by a shift in n (dot-dash line;

Figure 2.12, below). It should be noted that gain changes are often evoked that affect more than one of the parameters outlined above, however gain changes are usually classified by the predominant parameter shift.

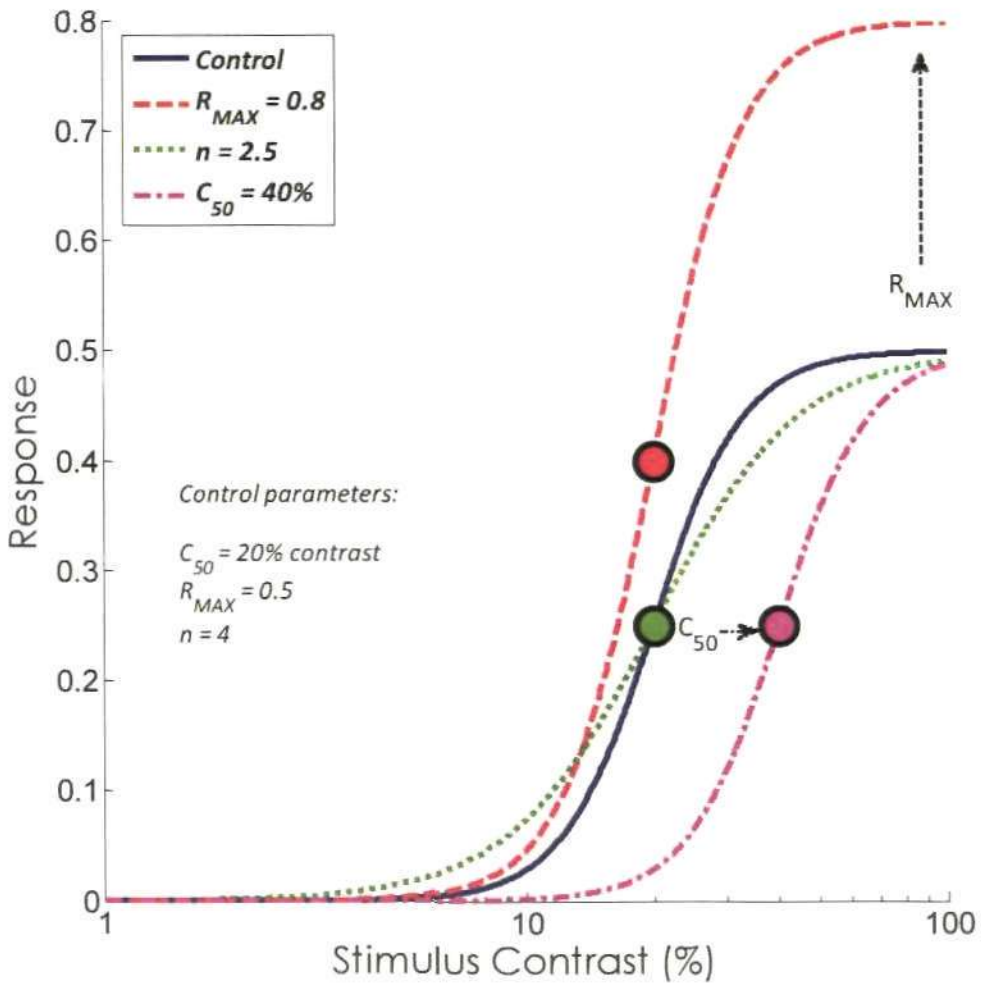

Figure 2.12: Gain control of the contrast-response function
Parameters of control response (smooth line) $n = 4$; $C_{50} = 20\%$ and $R_{MAX} = 0.5$. Parameter values following gain changes are given in the inset and are coordinated with line style. Dotted line shows effect of contrast gain control (only C_{50} is affected) on contrast-response function (CRF), dashed line plots response gain (R_{MAX}) and dot-dash line represents an example of a change in neural-gain (only the gradient, as determined by n, is affected)

2.2.3: Contrast adaptation

In general, many cells, within the early stages of the visual pathway, show response saturation in the presence of high contrasts; with the consequential effect of limiting their signalling capacity. However, even without contrast saturation, individual visual neurons, due to biophysical firing rate limits, do not generally have the capability to signal the full range of contrasts encountered during natural viewing (e.g. contrasts ranging from the faintest star in the night sky to dappled sunlight on the skin). In fact, many monkey and cat V1 neurons display dynamic firing ranges which are restricted to roughly 10 – 20% of the contrast range typically found within the environment (Albrecht & Hamilton, 1982; Frazor & Geisler, 2006).

A similar situation exists for the encoding of brightness levels. Retinal ganglion cells can encode around 2-log unit changes in illumination in their firing rate; however, cells must have the capacity to encode ≈10-log unit changes in brightness levels that are encountered within the visual environment. In terms of luminance, photoreceptors change the gain of their responsiveness, a process known as light adaptation, to match the prevailing levels of illumination; under dark conditions photoreceptors become preferentially sensitive to small changes in light levels, i.e. cells display increased gain to low levels of illumination, conversely, at high levels of overall illumination large changes in brightness are emphasised. Light adaptation is a slow process; for example, when walking into a darkened room from bright daylight; it takes several seconds for the eyes to adjust to the new light levels.

An analogous mechanism of adaptation to stimulus contrast exists within both cat and primate V1 (Maffei et al. 1973; Albrecht et al. 1984; Ohzawa et al. 1985; Sclar et al.

1989) and within primate lateral geniculate nucleus, or LGN, (magnocellular cells only, Solomon, Peirce, Dhruv & Lennie 2004). The dynamic response range of a cell shifts depending on the overall level of contrast. Figure 2.13a-b, below, plots the contrast-response function (solid lines and symbols) of a simple and complex V1 cell, respectively, after adaptation to stimuli ranging from low-high contrasts (left to right contrast-response function).

The dotted line gives the pre-adapted contrast-response function, whilst the dashed line plots the response to the adapting contrast; responses to the adapting contrast roughly equate to the C_{50}. As can be seen, adapting these cells to given contrast level displaces the *contrast-response function laterally, centring the C_{50} on the prevailing level of contrast*, as predicted by the Efficient Coding Hypothesis. Adaptation to low contrast levels centres the dynamic region of the contrast-response function onto the region of the log-contrast plot that emphasises small changes in contrast, whilst high contrast adaptation allows for the discrimination of large changes in contrast. For the cells in Figure 2.13c-d (complex V1 cell and Y-type LGN cell, respectively), the contrast-response function remains relatively unaffected by contrast adaptation.

Contrast adaptation can take up to 15 seconds to develop fully and the effects are long lasting (Ohzawa et al. 1982, 1985); this is in stark contrast to the time course of the contrast-response function itself which, in V1 at least, is fully developed within the first 10ms of response-onset (see Figure 2.9 above, Albrecht et al. 2002).

Contrast adaptation is generally considered a cortical mechanism as:

1. The earliest studies only found contrast adaptation in V1 (e.g. Movshon & Lennie 1979; Ohzawa et al. 1985)

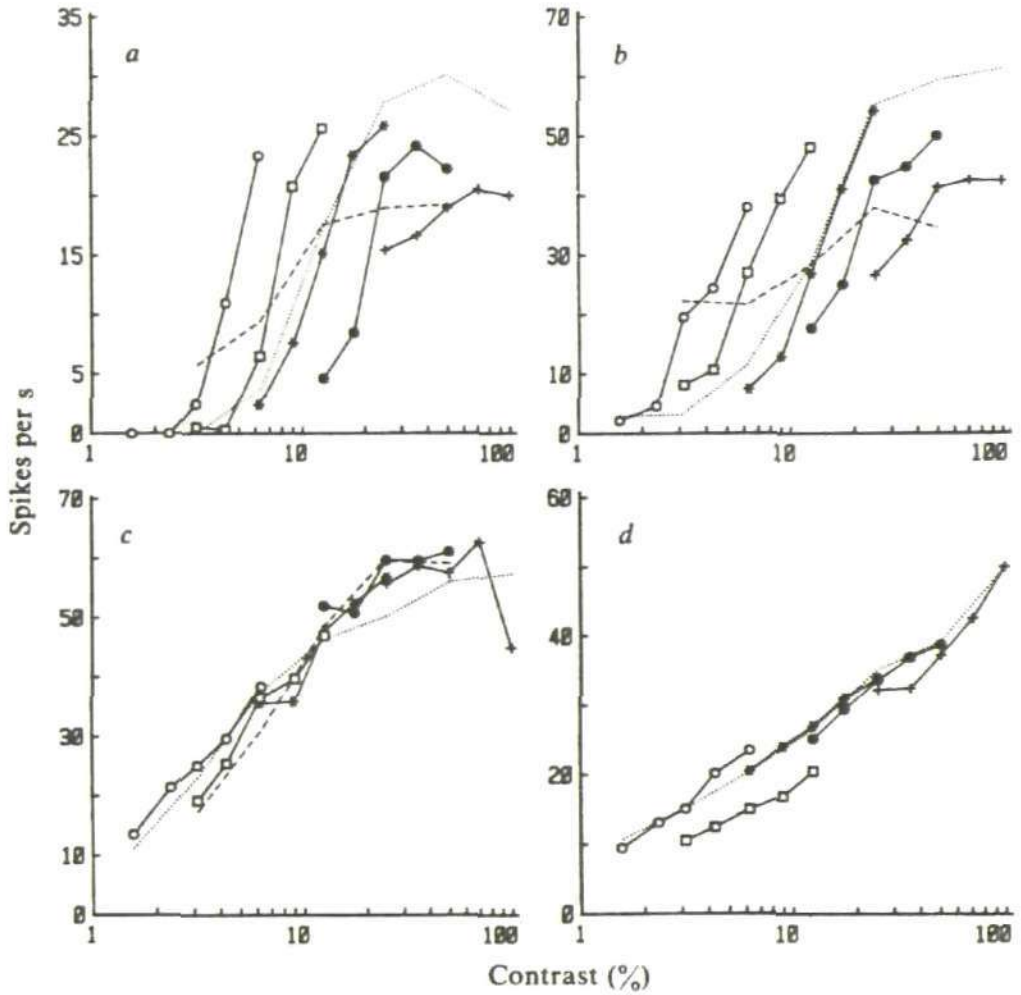

Figure 2.13: Contrast-response functions of cat V1 simple, complex and LGN Y cell, after contrast-adaptation

Contrast-response functions of cat V1 simple (a), complex (b-c) and LGN Y cell (d), after contrast-adaptation. Dotted line gives the pre-adapted, or control, response to a range of contrasts; solid lines & symbols plot post-adaptation contrast-response function, in response to increasing levels of adapting contrast (adapting contrast given by the intersection of the dashed line on the relevant contrast-response function). Top row, both cells display contrast adaptation. Bottom row, neither cells display contrast adaptation; note that whilst the contrast-response function of all V1 cells display sigmoidal contrast-response function, the LGN cell's response is best described by a linear function, thus response-profile is not an indicator of propensity for adaptation and not all cells have response profiles that exactly match the contrast-distribution (taken from Ohzawa et al. 1982).

2. Many cells within the early visual cortical areas respond strongly to grating stimuli and are preferentially sensitive to such stimulus features as orientation, spatial/temporal frequency etc. Contrast adaptation has been shown to reduce or alter responsiveness to, the pre-adaptation, preferred orientations and spatial frequencies (e.g. Maffei et al. 1973; Movshon & Lennie, 1979; Albrecht et al. 1984; Ohzawa et al. 1985; Sclar et al. 1989; Dragoi, Sharma & Sur 2000; Dragoi, Turcu & Sur 2001).

3. Some cells in V1 and V2 display contrast adaptation that is not orientation sensitive or have varying levels of adaptation dependent on the orientation of the stimulus (see Figure 2.14, below; Crowder, Price, Hietanen, Dreher, Clifford & Ibbotson 2006). These cells were most likely to found in areas where surrounding cells' orientations preferences changed rapidly (i.e. at the pinwheel centres of the orientation tuning map, Dragoi, Rivadulla & Sur 2001), thus their response properties are directly linked to the functional organisation of the cortex.

As neither contrast adaptation nor orientation tuning were observed in pre-cortical areas, such as the LGN, there was a general assumption that contrast adaptation must arise from purely cortical mechanisms.

However more recent studies have observed contrast adaptation in the retina, (Baccus & Meister 2002; Chander & Chichilnisky 2001; Kim & Rieke 2001; Rieke 2001) and LGN (Solomon et al. 2004), though it is cell type specific and species dependent.

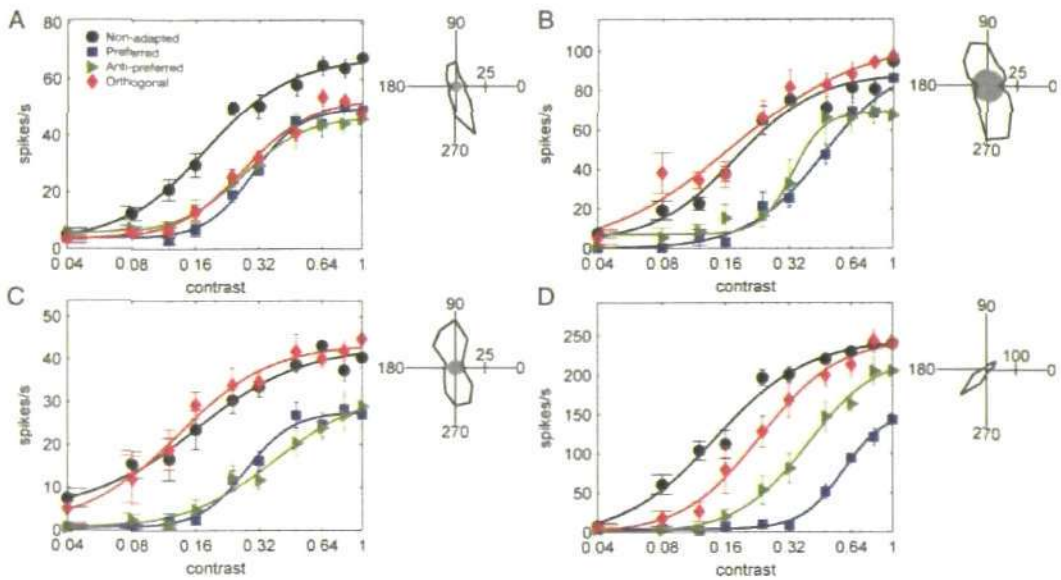

Figure 2.14: Example of 4 V1/V2 cells displaying different types of orientation-selective contrast adaptation

Contrast-adapting stimuli were either drifting gratings of the cell's preferred orientation and drift direction (preferred, blue line, solid squares), drifting gratings of preferred orientation but orthogonal drift (where orthogonal is always relative to the preferred stimulus parameter, anti-preferred, green line, solid triangles) or orthogonal orientation with preferred drift (orthogonal, red line, solid diagonals). Adapting contrast was 0.32; the black line and filled circles plot the un-adapted, control, contrast-response function. Symbols are the averaged firing rate (error bars give the SE); smooth lines give best-fit to the data. **A**: Example of non-orientation adaptation, the cell's contrast-response function to all adapting stimuli (preferred direction/orientation, anti-preferred direction/preferred orientation; orthogonal orientation) undergoes adaptation. **B-C**: Two cells displaying orientation selective adaptation, adaptation effects are only apparent for the preferred orientation stimuli. **D**: Example of a cell showing intermediate orientation selective adaptation. Taken from Crowder et al. (2006).

2.2.4: Possible mechanisms underpinning contrast gain effects such as contrast saturation and adaptation

Some of the possible mechanisms underpinning contrast gain are outlined below.

Divisive normalization model

The most widely accepted model of contrast gain is that of the divisive normalization model (Heeger 1992; Carandini & Heeger 1994; Carandini, Heeger & Movshon 1997; Carandini & Ringach 1997). Heeger (1992) proposed that the lateral shift in the contrast-response function observed (on log-coordinates) during contrast adaptation is equivalent to the contrast-response function being divided by the adapting contrast, thus cells cope with their limited firing capability by normalising their output with respect to overall stimulus energy. Contrast normalisation has also been shown to increase the information-carrying capacity of V1 cells (Brady & Field 2000). The normalization model extends the standard linear model of simple cell dynamics by introducing a divisive stage, post linear filtering: division arises from the pooled input of a large number of cortical cells, the normalization pool.

Figure 2.15 is a schematic diagram of the different stages of the divisive normalization model, applied to a V1, simple cell. Essentially, a cell's selectivity is determined by its linear spatiotemporal weighting function (Stage 1,

Figure **2.15**); the linear output is then scaled, or divided, (Stage 2) by the normalization pool (Stage 3) and then rectified (Stage 4) before finally generating a response (Stage 5); as can be seen the cell's own output is fed back into the normalizing pool (Stage 6).

Cells within the normalization pool are tuned across the whole spectrum of stimulus parameters, thus the scaling factor is the overall energy of the pooled response, determined by stimulus contrast alone. The normalizing pool is believed to contribute to a gamma-aminobutyric acid (GABA) mediated shunting inhibitory conductance which holds the cell membrane close to its resting potential, thereby preventing the cell from firing. Only

excitatory inputs that are potent enough to counteract the influence of divisive inhibition on membrane potential can drive the cell to fire.

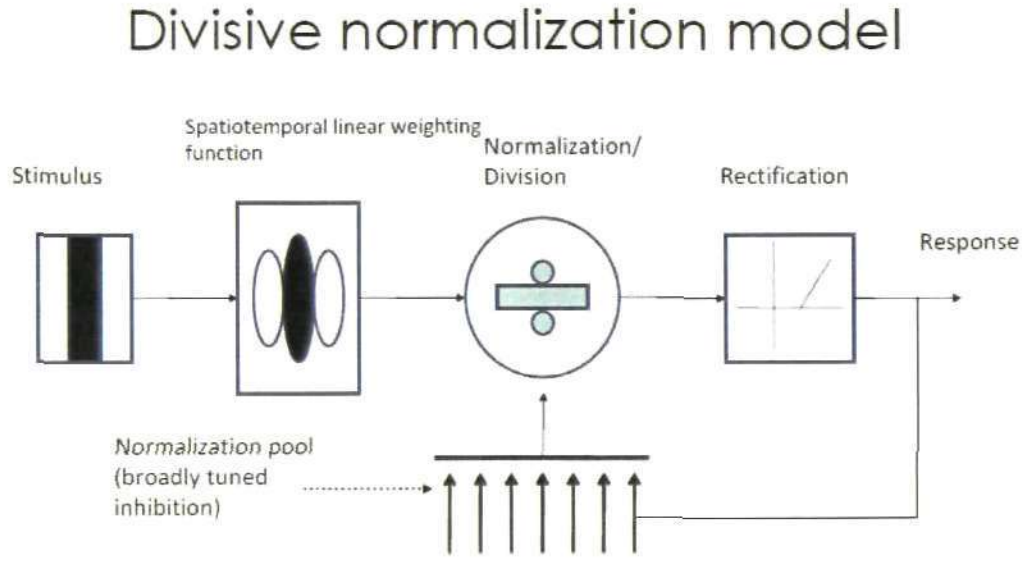

Figure 2.15: Cartoon of divisive normalisation model
Divisive normalisation model, as proposed by: Heeger (1992); Carandini & Heeger (1994); Carandini et al. (1997) and Carandini & Ringach (1997). See text for details.

How does shunting inhibition work and why is it considered divisive? GABA opens up chlorine ion channels which allow negatively charged chlorine ions to enter the cell; thus negating the depolarising influence of any voltage dependent inward flow of sodium ions. The chlorine-current reversal potential is -70mV and as such is close to most cells resting potential. Shunting inhibition increases overall conductance rather than the synaptic current, I_{SYN}; this changes the slope of the input-output relationship (in this case, membrane potential as a function of synaptic current, which is linear) and is thus considered divisive (see Figure 2.16A, below).

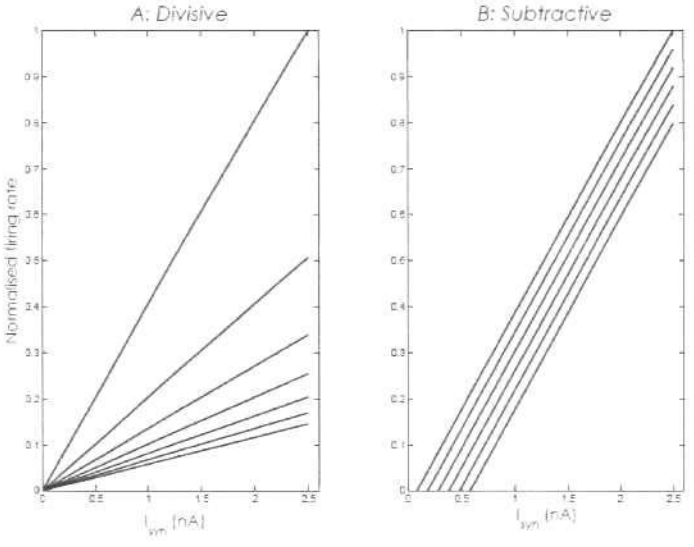

Figure 2.16: A comparison of the effect of divisive and subtractive inhibition on the f-I curve
A: Divisive inhibition changes the slope of the *f*-I curve. Membrane voltage (V) is calculated as $V = I_{SYN}/G$, where I_{SYN} is the synaptic current and G is the conductance, the firing rate (*f*) was calculated from the linear function $f = g(V)$ and normalised; *g* can be any monotonic function and in this case was *tanh*(V) (as in Hopfield, 1984). G ranged from 10 to 70 nS in equal steps. **B:** Subtractive inhibition shifts the curve by subtracting an inhibitory current I_{INH}. I_{INH} ranged from 0.08 to 0.58 in equal steps and G was equal to 10nS. Both plots are adapted from Holt & Koch (1997).

Conversely, inhibition with a reversal potential far from resting potential (i.e. potassium mediated inhibition, reversal potential = -100mV) acts as a current source, I_{INJ}, and is therefore subtractive: $I_{SYN} = I_{SYN} - I_{INJ}$. The input-output function is shifted along the abscissa (see Figure 2.16B, below).

For low stimulus contrasts, pooled inhibition is negligible and there is a fairly linear transduction of afferent input into post-synaptic firing rates, however, as stimulus contrasts increases, so does the strength of the shunting inhibition, thereby making it harder for feedforward excitation to drive the cell across threshold. A threshold mechanism gives rise to the response-exponent (or half-squaring, within the normalization model), with *n*

determined by how close threshold is to resting potential. A similar model, the contrast-gain exponent model, has been developed (Albrecht & Geisler 1991). This is a purely descriptive model in which input is filtered through three stages prior to rectification: a contrast gain control mechanism, which generates response saturation to high contrasts, a linear spatial summation stage that sets up stimulus selectivity and finally the output is filtered with a response-exponent.

Issues with the divisive normalization model

There are serious issues over whether cortical inhibition can be the source of divisive normalization. The evidence of the role and tuning of cortical inhibition is highly contradictory. For example, whilst several studies have found diversity in the orientation tuning profiles of inhibitory inputs (Martinez, Wang, Reid, Pillai, Alonso, Sommer & Hirsch 2005; Monier, Chavane, Baudot, Graham & Frégnac 2003), others indicate that tuning preference of both excitatory and inhibitory conductances share the selectivity of the target cell (e.g. Priebe & Ferster 2006; Anderson, Carandini & Ferster 2000). Also, according to Holt & Koch (1997), when translated to a spiking neuron model, shunting inhibition becomes subtractive rather than divisive, as spiking effectively clamps the membrane potential significantly above resting potential.

Finally, the divisive-normalization model assumes that at low-contrasts the transduction of stimulus to response is linear, however as outlined above, the majority of V1 cells have contrast-response function functions that are best fit with an expansive nonlinearity for low contrasts (e.g. the mean value of the Naka-Rushton parameter n, in V1,

is ≈2.5) and this has also recently been found to be the case in feline LGN (Duong & Freeman 2007).

Spike-rate adaptation can account for contrast adaptation

Whilst the divisive-normalisation provides a possible model of contrast-gain control, it does not necessarily provide a bio-physical mechanism to account for it. One possible membrane mechanism that may underpin contrast adaptation is an after-hyperpolarisation in the membrane potential, post spike generation. The after-hyperpolarisation is believed to be essential for spike-rate adaptation (or spike-frequency adaptation: e.g. Madison & Nicoll 1984; Connors & Kriegstein 1986), a property of regular-spiking cells (either pyramidal or spiny stellate cells, Connors & Gutnick 1990) whereby, under conditions of constant stimulation, the inter-spike interval, ISI, increases (see Figure 2.17, below). Within regular spiking cells, after-hyperpolarisation and spike frequency adaptation (Sanchez-Vives, Nowak & McCormick 2000) has been linked to both sodium and calcium mediated potassium current ($I_{K(Na)}/I_{K(Ca)}$).

Synaptic depression

Synaptic depression has also been implicated as a cortical gain mechanism (Abbott, Varela, Sen & Nelson 1997; Tsodyks & Markram 1997; Varela, Sen, Gibson, Fost, Abbott & Nelson 1997). Synaptic transmission is a dynamic process, that is, synapses can often display use-dependent short-term plasticity, expressed by enhancement or reduction in synaptic efficacy, the effects of which can last for tens of milliseconds to seconds. The essential premise is that for each presynaptic spike, the synaptic resource is depleted by a given factor (a utilisation parameter, u) and then recovers back to its base level within a given time, set by the synaptic recovery time constant (τ_R).

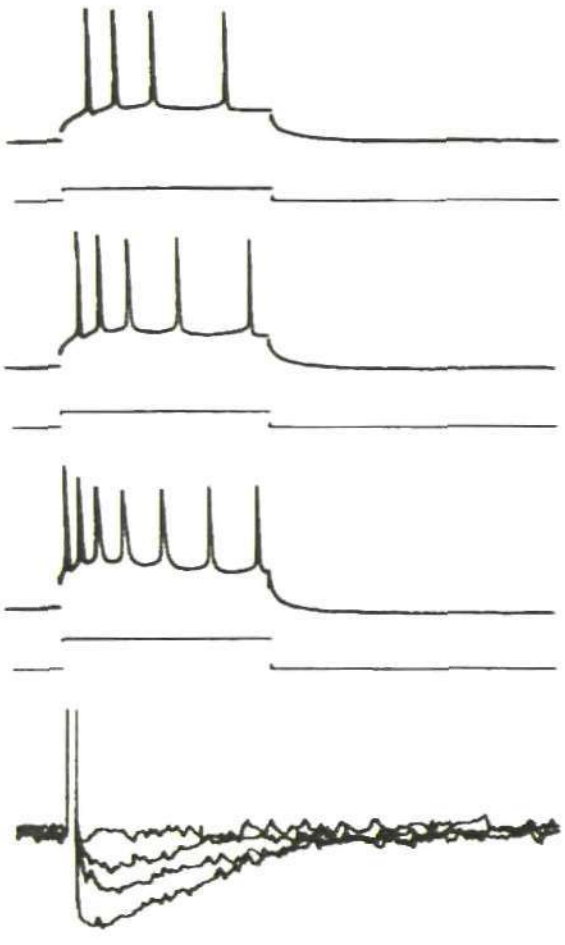

Figure 2.17: Spike-rate adaptation in a hippocampus pyramidal cell
The top three traces show the membrane response to an injection of stepped, depolarising, current (the current profile is below each trace, the level of injected current increased from top to bottom trace). The bottommost trace is the after-hyperpolarisation, following current offset, for each of the three traces above (the smallest after-hyperpolarisation followed the offset of a current that elicited one action potential and isn't shown in the three traces above), the depth of the after-hyperpolarisation increased with injected current. Taken from Madison & Nicoll 1984.

During the time period between depletion and recovery the probability of synaptic transmission (p) is reduced and thus the synapse can be said to be depressed. Experimental (e.g. Varela et al. 1997; Varela, Song, Turrigiano & Nelson 1999) and modelling studies

(Abbott et al. 1997; Tsodyks & Markram 1997; Markram & Tsodyks 1996) indicate that *cortical synaptic depression enhances the signalling of transient changes in firing rate*, rather than average firing rate, and may therefore be a rapid and automatic form of cortical gain control that is important for temporal filtering of afferent input.

In terms of thalamocortical synapses and contrast-gain control, depression is strongest when afferent firing rates are high, thereby leading to response saturation in the post-synaptic response. Under low contrast conditions firing rates are slow, therefore the amount of synaptic depletion is minimal and synapses have time to recover fully between presynaptic spikes. For high contrasts, the presynaptic firing rate is too rapid to allow for complete synaptic recovery between spikes. Thus the probability of synaptic transmission is chronically low, leading to response saturation in the afferent input and consequently the target cell itself.

Geniculocortical synaptic depression has been observed *in vitro* across sensory cortices (e.g. rat somatosensory cortex [Gil, Connors & Amitai 1999]; cat V1, Layer 4 excitatory spiny stellate cells [Stratford, Tarczy-Hornoch, Martin, Bannister & Jack 1996]) and *in vivo* within the somatosensory cortex (Chung Li & Nelson 2002); several modelling studies have used both fast and slow forms of thalamocortical depression to model contrast adaptation and saturation (Chance, Nelson & Abbott 1998; Adorján, Piepenbrock & Obermayer 1999). In conclusion, synaptic depression provides a powerful mechanism for describing many of the non-linear, contrast-response dynamics of V1 simple cells that can't be fully accounted for by the long-standing model of intracortical divisive inhibition. Another advantage is that synaptic depression has also been used to successfully model

contrast adaptation (Adorján et al. 1999), therefore synaptic depression may be a 'unifying' mechanism in that it can account for both long and short term contrast gain dynamics.

However, one key issue with thalamocortical synaptic depression, as a model for contrast gain, is that depression levels are highly saturated *in vivo*. Boudreau & Ferster (2005) measured the effect of thalamocortical synaptic depression on post-synaptic potentials (PSP), *in vivo*, to small trains of spikes elicited in the LGN by electrical stimulation. Monosynaptic PSP (i.e. thalamocortical PSP) displayed only small levels of depression, especially when compared to those of polysynaptic PSPs and much less than expected given comparable *in vitro* data (e.g. Stratford et al. 1996). Reducing spontaneous activity in the LGN through increased intraocular pressure led to an increase in amplitude of the 1st PSP followed by reduction of subsequent PSPs to similar amplitudes relative to the normal condition, thus suggesting that spontaneous thalamic activity saturates thalamocortical depression *in vivo* (see also Reig, Nowak & Sanchez-Vives, 2006).

This has also been observed in the somatosensory cortex (Castro-Alamancos 2004b) in relation to changes in thalamic activity during quiescent and active states. Increased thalamic firing, during arousal states, resulted in sensory suppression, in the rat barrel cortex, due to thalamocortical synaptic saturation (Castro-Alamancos & Oldford 2002). This suppression is manifested as a reduction in the RF size of cells in the input layer (layer 4) of the barrel cortex, such that cells only respond to principle whisker stimulation (Castro-Alamancos 2004b). The saturation of synaptic depression during active states is proposed to focus sensory information transmission onto the strongest, or most salient, stimulus feature.

This suggests that the thalamocortical synaptic depression model of contrast gain

control may not be readily applied to awake behaving systems, however, synaptic depression is not fully saturated even in the *in vivo* condition (Boudreau & Ferster 2005) and could thus still have an impact on the transmission of information from thalamus to cortex. Intracortical synaptic depression could be an alternative source of synaptic plasticity-mediated gain control (e.g. Chelaru & Dragoi 2008). However, it should be noted that intracortical depression levels in the V1 cortex have also been shown to be highly dependent upon ongoing levels of background activity, with post-synaptic potentials (from both intracortical and thalamocortical connections) displaying less depression for cells with large amounts of spontaneous synaptic activity (Reig, Gallego, Nowak & Sanchez-Vives 2006)

Contrast dependent non-linearities in LGN and retina

To what degree can contrast-gain effects in the cortex be inherited from feedforward inputs? Certainly it has been known for several decades that retinal ganglion cells are subject to a rapid type of gain control that leads to response saturation at high contrasts (Shapley & Victor 1979). Within the cat, this phenomenon is seen in both classes of retinal ganglion cells (X- and Y-type, Victor & Shapley 1979). However, within primate retina the situation is somewhat different. Primate retinal ganglion cells are classified as either parvocellular (P) or magnocellular (M) type cells (depending on morphology and response preference to such stimulus features as temporal and spatial frequency). The general observation is that only M-retinal ganglion cells undergo contrast saturation whilst P-retinal ganglion cells display a linear contrast-response function (For a review see Shapley 1986).

Retinal ganglion cells provide afferents that are selectively connected to corresponding type thalamic cells, which inherit the response dynamics of the feedforward retinal input. It should be noted, that within cat retina and thalamus, response saturation is not a universal property, and is also not specific to a cell class, as in primate (Ohzawa et al. 1985).Within rabbit thalamus (Cano, Bezdudnaya, Swadlow & Alonso 2006), contrast saturation is only observed with transient-type cells, whereas sustained cells display highly linear contrast-response functions.

It would be useful to quantify to what degree cells in V1 share characteristics of LGN responses. Whilst there isn't a wealth of data on the subject, C_{50} values in V1 and LGN appear to be relatively similar (average C_{50} for cat LGN = 20.3 (±12.1) and for cat V1 =21.3 (±13.5); Sadakane, Ozeki, Naito, Akasaki, Kasamatsu & Sato 2006). However, values of n appear to be almost doubled in V1, relative to LGN (Sclar, Maunsell & Lennie 1990; Contreras & Palmer 2003; Albrecht, Farrar & Hamilton 1984, see Table 2.2).

Table 2.2: Comparison between the average values of the response exponent, n, in cat and monkey V1 and LGN.
Where n describes the slope of the linear region of the contrast-response function, see Equation 2.20

Species	**Value of n:**	
	LGN	*V1*
Cat	1.4 (1)	2.9 (3) ; 2.2 (1)
Monkey	1.6 (2)	2.9 (3) ; 2.4 (2)

(1): Contreras & Palmer 2000; (2): Sclar et al. 1990; (3): Albrecht & Hamilton 1982

Thus, to what extent are the contrast-dependent dynamics of LGN cells inherited by

V1 cells? Both complex and simple V1 cells display a variety of contrast-response functions (Albrecht et al. 1984). This could arise from the strict segregation of V1 afferents, along the lines of contrast-response function properties, such that V1 cells with saturating contrast-response function receive inputs only from LGN cells which show similar levels of saturation in their contrast-response function. However, whilst there is evidence that that C_{50} values are conserved during thalamocortical transmission, the exponent value (n) governing the slope of the dynamic contrast response is increased, see Table 2.2.

In conclusion, the above discussion indicates that response saturation is a feature of many retinal cell response dynamics; and is inherited by connecting cells in the LGN and presumably then passed on to V1. However what is not clear is what relationship there is between levels of contrast saturation in the cortex and LGN.

2.3: Other evidence for Efficient Coding Hypothesis

The above discussion has focused on the response to visual contrast, as evidence of efficient sensory coding. However, the experimental paradigms used to test the adaptive processes outlined above tend to be static and un-naturalistic (i.e. the 'adapt and probe' technique, whereby any shifts in responsiveness, following adaptation with sustained stimuli, is tested using probe stimuli, post-adaptation). For the Efficient Coding Hypothesis to be considered a universal encoding strategy one would expect to find evidence for efficient coding following adaptation to dynamic, naturalistic stimuli and across a range of parameters and modalities.

As discussed, Laughlin (Laughlin 1981) studied the contrast response in the early fly visual system by presenting contrast stimuli drawn from a distribution of contrast levels found in the wild; however presentation techniques were static and bore little relationship

to the temporal distribution of contrast one would expect a blow fly to encounter in the wild. More recent studies have explored adaptation to dynamic, motion stimuli in the visual system of the blow fly. The fly H1 cell is a direction-sensitive visual cell that conveys velocity information about horizontal motion across its receptive field. H1 firing rates (Brenner, Strong, Koberle, Bialek & de Ruyter van Steveninck 2000) were shown to track slowly-varying (time scale of 1s) Gaussian velocity stimuli (see Figure 2.18A&B) and also underwent adaptive rescaling; in that when the adapting stimulus was scaled by doubling the standard deviation, firing rates to both the scaled and non-scaled stimuli displayed considerable overlap (Figure 2.18C, below). Adaptation to the higher standard-deviation stimulus resulted in a slight lateral displacement of the velocity-response function (Figure 2.18D), however when neural rate functions were normalised, with respect to average firing rate and standard deviation, there was a high degree of overlap (Figure 2.18E, below) between functions, with divergence only observed for the highest velocity stimuli. A similar form of rescaling was observed when the stimulus varied rapidly (on the time scale of ≈10ms), thus H1 cells normalise their output by adapting to the local statistics of the stimulus. The measured rate functions were also shown to maximise information transfer, relative to rate-functions that were derived through either stretching or contracting the measured stimulus-response functions.

Adaptation to velocity in fly H1 neurons occurs on many time scales, ranging from the 100ms to tens of seconds scale (Fairhall, Lewen, Bialek & de Ruyter Van Steveninck 2001). Using a switching technique, whereby the standard deviation of a white-noise velocity signal was switched between two values every T seconds, it was observed that, on the longest time scale, firing rates adapted to the standard deviation of the stimulus (over

the time period *T/2*), with the length of the adaptation time constant (a measure of the time for firing rates to relax to the steady-state level, from the transient changes in firing rate observed following the stimulus switch) dependent on the length of *T*. In terms of rapid adaptation, the rescaling of the neural-response function occurred within the first second of the stimulus switch. This type of rapid rescaling, to a switching, adapting stimulus, is not restricted to the fly visual system and has also been observed in the inferior colliculus (Dean et al. 2008); adaptive responses, following a stimulus switch, were elicited to noise bursts stimuli drawn from one of two highly skewed distributions (see also Dean et al. 2005 and Chapter 3 for a description of the stimulus distribution); the rescaling of neural-response functions, towards the steady-state level, was observed within 300ms of the stimulus switch.

The stimulus-information transfer rate was found, in the fly H1 cell (Fairhall et al. 2001), to be relatively constant, at around 1.5 bits per spike, except for a transient dip (with a time constant of ≈40ms) in information rate following a switch to the lower standard deviation stimulus. Thus the rescaling of firing rate and neural response profiles serve to maintain a constant flow of information about the external world, in the presence of transient changes in distribution.

The above experiments highlight the fact that adaptation acts on many time scales; in fact adaptation can be considered as a continuous, life-time process, even an evolutionary one. Organisms, and sensory systems, evolve within and adapt to the confines of their ecological niche; thus if efficiency is the driving force behind adaptation, as predicted by the Efficient Coding Hypothesis, one would expect that neural codes have evolved to represent environmental signals in as economical manner as possible. Laughlin

(Laughlin 1981) demonstrated that the response (in the large mono-polar cells of the blowfly, at least) to contrast distributions, drawn from a natural environment, follows the optimum code expected by the Efficient Coding Hypothesis and Information Theory. However, visual contrast is a relatively simple parameter (i.e. the fluctuation of illumination about its mean), especially when compared to more complex sensory input, such as auditory stimuli. Auditory stimuli are not spatially discrete and have a variety of parameters that need to be resolved, e.g. frequency, amplitude and direction. A key functional attribute of auditory stimuli is communication, e.g. speech, mating calls or warning sounds.

Thus it could be argued that the complex nature and communicative function of auditory information provides the ultimate test of the Efficient Coding Hypothesis.

As such, theoretically derived optimal codes (Lewicki 2002; Smith & Lewicki 2006) for the efficient encoding of human speech, animal vocalizations and environmental sounds (e.g. leaves rustling, twigs snapping etc) have been found to closely resemble experimentally defined features of auditory nerve and cochlear processing. Thus, it would appear that human speech evolved to have spectral features that closely match the optimum, in terms of auditory processing along information theoretical lines.

Accordingly, the evidence above reveals that adaptation is a dynamic process, acting on many time scales and serving to reflect the local and global statistics of incoming stimuli through the rapid rescaling of the input/output function as well as transient and long-term changes in firing rates.

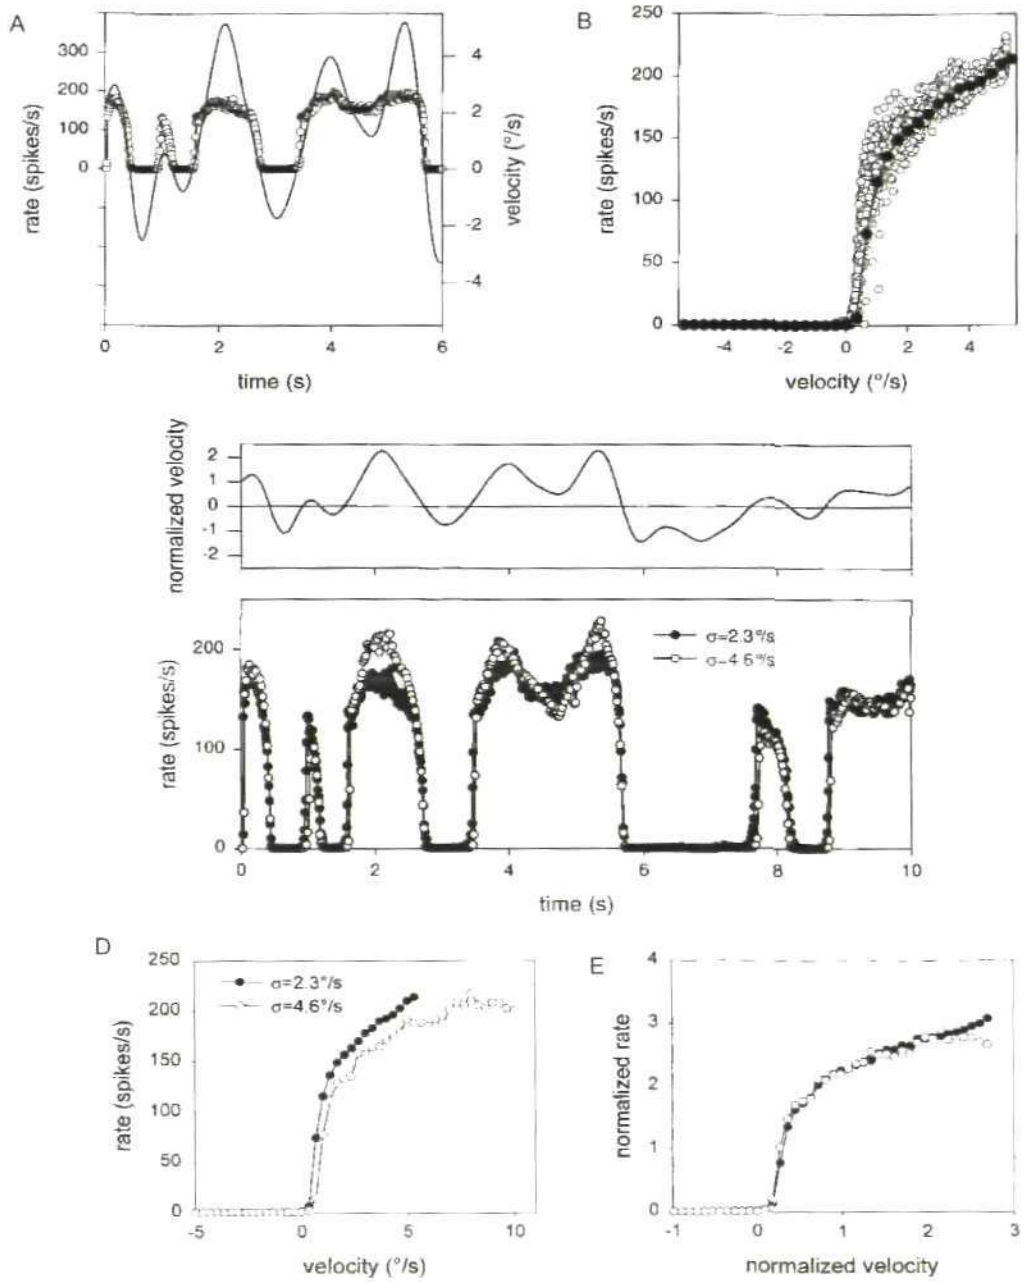

Figure 2.18: Adaptive rescaling in fly H1 neuron to slowly-varying velocity stimuli Adapted from Brenner et al. (2000). **A:** Solid line: stimulus (right axis), circles; rectified response (left axis). **B:** Neural-response function, empty circles: firing rate verses velocity, filled circles: average firing rate as function of stimulus velocity. **C:** Bottom trace: firing rate to two stimuli (see trace above) with $\sigma = 2.3$ (filled circles) & $\sigma = 4.6$ (empty circles). Top trace plots normalized stimulus. **D:** Velocity-response functions from stimuli in B. **E:** Velocity-response functions normalised by mean firing rate and σ of the stimulus.

2.3.1: Extensions to the Efficient Coding Hypothesis

The above discussion reveals that there is strong evidence that sensory systems encode signals in a manner consistent with the Efficient Coding Hypothesis. However, a main weakness of the Efficient Coding Hypothesis is its assumption of no noise (Barlow 1961). This has been addressed by the principle of information maximisation (or Infomax, see below).

Noise

Whilst the assumption that sensory systems are essentially noiseless was useful for hypothesising about the efficiency of the sensory code, Barlow himself acknowledged (Barlow 1961) that it was not a biologically realistic representation of actual neural communication. Neuronal interactions are inherently noisy, with neural stochasticity being intrinsic to the neuron (e.g. fluctuations in membrane potential), extrinsic (i.e. background or input noise) or synaptic in origin (e.g. due to the stochastic nature of transmitter release). In the presence of noise, some redundancy is necessary for resolving ambiguity; as such incorporating noise into the hypothesis leads to changes in the expected profile of a redundancy reducing code.

As outlined above, where the input to a noiseless system is Gaussian, the optimal, redundancy-reducing, mapping of stimulus to response is the cumulative probability function (cumulative probability distribution) of the input, as output entropy is maximised due to histogram equalization (i.e. all response states have equal probability). If noise is introduced into the system then the optimal strategy is to have a response probability distribution that is maximal about the signal (as opposed to a uniform distribution in the

noiseless case), such that response resources are focused on the signal itself rather than the noise (Zhaoping 2006). However, in this case, due to the unequal response-probability, output entropy is not maximal, thus redundancy in the system is increased, relative to the high signal-to-noise ratio case.

For example, in the case of a low signal-to-noise, stimulus ensemble with a Gaussian distribution, the signal is essentially the average stimulus of the distribution and the optimal transfer function (from stimulus to response) is one which ensures that responses elicited to the signal occur most frequently. An example hypothetical stimulus-response function and corresponding response probability function is plotted in the bottom row of Figure 2.19, below; the slope, or gain, of the stimulus-response function is reduced for those stimuli that occur most frequently. A noise-dependent adjustment in coding strategy has been observed experimentally in the compound-eye of the blowfly (Laughlin, Howard & Blakeslee1987).

Understanding and defining the optimal code is a major problem when testing the *validity of the Efficient Coding Hypothesis, in that one must know the full range of signals* to which a sensory system is exposed, their distributions and the levels of intrinsic and extrinsic noise within a system.

One method of quantifying the effect of noise on encoding of sensory input is to measure the mutual information between the stimulus and neural response. If one can recall from the earlier introduction to information theory (see the Section 2.2.2), mutual information essentially quantifies the amount of information an input/output relation, $(x|\ y)$, hold about each other, after taking into account transmission noise. Thus it has been proposed (Linsker 1990) that noisy systems encode input efficiently by maximising mutual

information.

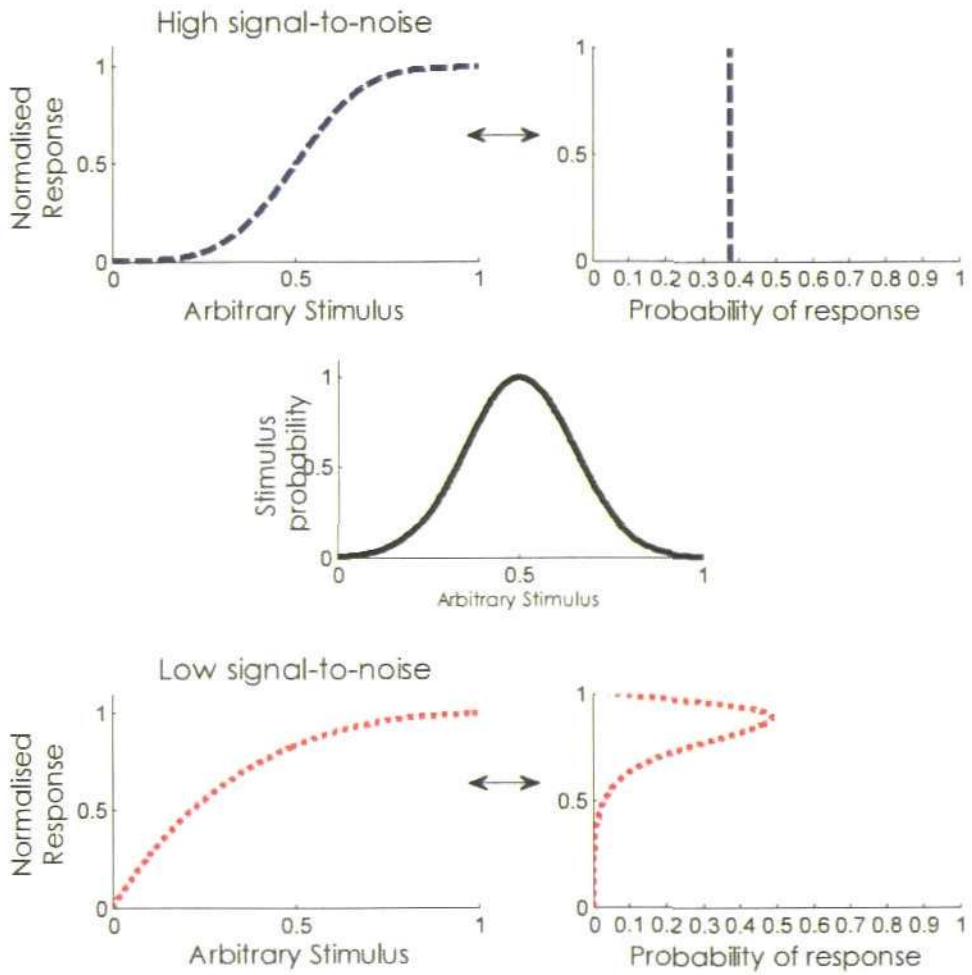

Figure 2.19: Hypothetical example of optimal response profiles under conditions of low and high noise

The central plot is an arbitrary stimulus with a Gaussian probability distribution. Top row plots the optimal response profile (left) when the signal-to-noise ratio is high (i.e. low noise) and the response probability (right). The response profile is the cumulative distribution function of the input, where all possible responses have equal frequency and the response probability has a uniform distribution. Bottom row plots the optimal response profile (left) and the response probability (right) when the signal-to-noise ratio is low. The hypothetical stimulus-response function (left) does not follow the cumulative distribution function; the response probability function (right) is maximal for responses elicited to inputs located close to the mean of the stimulus distribution.

Information maximisation principle

The information maximization (or Infomax) principle is essentially an optimization method for neural networks such that:

"...network connections develop in such a way as to maximize the amount of information that is preserved when signals are transformed at each processing stage, subject to certain constraints" Linsker (1990)

Here, the information that is preserved is the mutual information between the input and output of each layer of in the network and the constraints on the processing units/stages are biological in origin, e.g. lateral connectivity, latency or synaptic transmission probabilities.

The Infomax principle can be considered as an extension to the efficient coding hypothesis as, in the absence of noise, maximising the mutual information between the input and output is equivalent to maximising the input entropy (see Section: 2.2.2, above); thus Infomax, in the zero-noise case, is equivalent to redundancy reduction.

Linsker (Linsker 1990) found that neural networks of units, modelled on sensory cells and with Hebbian-type synapses (Hebb 1949; Hebbian here refers to the learning rules for synaptic weight adjustment and is easily summarized by the phrase: 'units that fire together, wire together'), self-organized into feature-detecting units with similar properties to that observed in V1, e.g. orientation selectivity, when optimized using an Infomax learning algorithm.

Information maximization is generally taken to be the underlying principle of many adaptive processes (e.g. Adorján et al. 1999; Laughlin et al. 1987; Yu, Potetz & Lee 2005) as adaptive shifts in neural responsiveness tends to improve information transmission for sensory inputs within the prevailing stimulus distribution. It has also proved to be a useful,

and successful, concept for analysing sensory systems and developing self-organising neural-network models (e.g. Okajima 2001).

Sparse coding

According to Barlow, sensory systems aim to represent input in as economical and efficient manner possible. This is achieved through exploitation of the inherent redundancy in sensory input (Barlow 1959; Barlow 1972). For example, contrast adaptation (see Section: 2.3) is a cellular mechanism that exploits the fact that contrast levels, over a given spatial and temporal scale, rarely fluctuate across the full range of possible contrasts encountered in the visual environment. Thus efficiency can be achieved through encoding relative, over absolute, levels of visual contrast.

Another major source of sensory redundancy is in the form of local temporal and spatial correlations. For example, in a natural-scene image, neighbouring pixels are more likely to be similar than not (Barlow 1972). As such, sensory systems can exploit these correlations by ignoring them and only transmitting what is not redundant.

Thus it is possible for each processing stage to encode more specific features of the visual scene without increasing the numbers of cells needed to code for each feature. This is termed sparse coding and Barlow believed it was crucial to neural encoding (Barlow 1972).

Sparse coding (Olshausen & Field 2004; Field 1994) represents a trade off between the robustness of a high density code and redundancy reduction of a local code. Sparse coding is defined by the ratio of active-to-passive cells for a given input; if significantly less than half the cells in a population are active then the input has a sparse representation, within the cortex. A schematic comparison between a compressive-code versus a sparse

code is given in Figure 2.20. Both types of encoding elicit the same number of spikes; however a sparse code is able to represent the input with the minimum amount of spiking activity, without the need to compress the input.

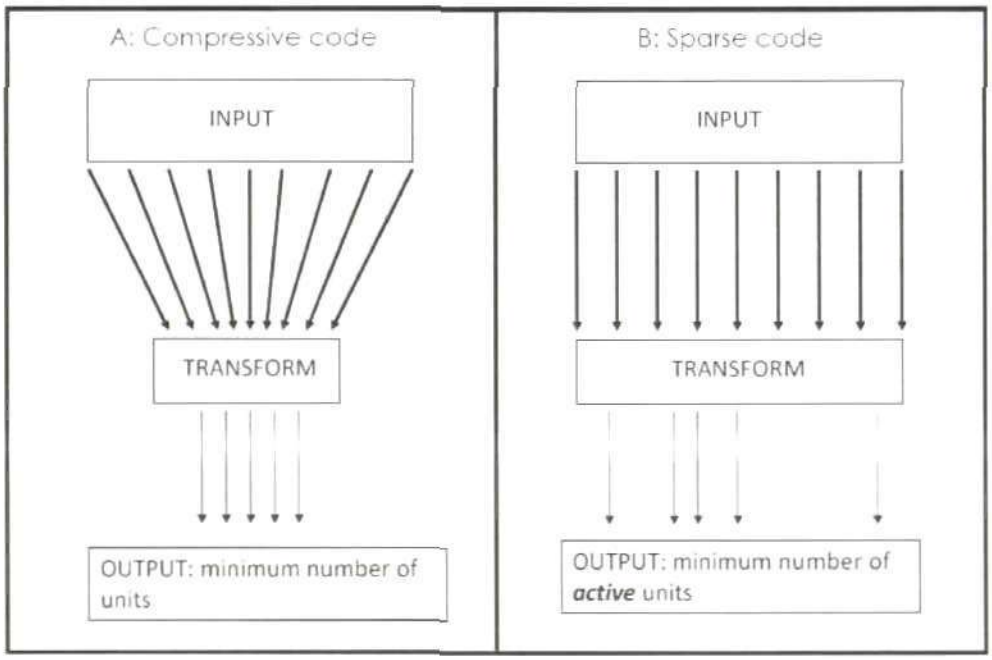

Figure 2.20: Schematic comparison of redundancy reduction by either compressive or sparse encoding of sensory input
See text for details. Adapted from Field (1994).

2.4: Criticisms of the efficient coding hypothesis

According to the efficient coding hypothesis sensory systems compress incoming information using a strategy of redundancy reduction; compression is a feature of biological structure in the early stages of sensory processing (e.g. in the transmission of retinal information along the bottle neck of the optic nerve). However, as one traverses the processing hierarchy neurons become much more numerous. As such, it can be argued that

redundancy may actually increase, as more neurons are able to respond to the compressed signal. However, as we have seen in the previous section the representation of sensory inputs becomes progressively sparser at each successive cortical stage (e.g. Olshausen & Field 2004) and firing rates are generally very low. Nevertheless, Barlow recently acknowledged (Barlow 2001) that redundancy may be important, or even necessary, for the processing of sensory information in the cortex.

However, it should be remembered that the fact that redundant information may be utilised, post transmission, is implicit in the efficient coding hypothesis. The purpose of information maximization/redundancy reduction is to package information as efficiently as possible **whilst** minimising loss of information, as redundant information is not necessarily the same as useless information.

As such the efficient coding hypothesis is not a universal theory and depending on the constraints placed upon a system and the levels of information processing/transmission taking place within a given cortical area information maximisation may not always *represent the optimal strategy for the encoding of sensory information. Some examples of* alternative strategies are outlined below.

A perfect redundancy reducing code will be fully reversible, i.e. it should be possible to reconstruct the stimulus from the neural response. However in reality this is not often the case as, firstly, transmission is rarely noiseless and, secondly, many neurons respond selectively to stimulation, thus not all stimulus information is encoded. A good example is orientation selectivity in V1, where complex cells respond preferentially to the orientation of a grating stimulus and the direction of drift.

As has been discussed above, the efficient coding hypothesis assumes that neural

responses adapt to the distribution of inputs, conserving/maximising information. Alternatively, neurons that selectively encode sensory inputs, adapt to reflect a particular feature/or subset of the stimulus distribution and ignore or suppress the remainder. An example of selective coding is background-suppression in the auditory system of insects, whereby responses to a loud stimulus suppress responses to subsequent stimuli, over a given time window (Sobel & Tank 1994; Baden & Hedwig 2007)

A comparison of how both strategies (effective and selective coding) would affect stimulus-response curves (stimulus-response function), following adaptation to stimuli with multi-modal distributions, is given in Figure 2.21A&B below. Adaptation that is underpinned by a background-suppression mechanism (see Figure 2.21B, below) would result in stimulus-response function that only covers the highest-intensity mode of the stimulus distribution; rather than covering the whole distribution (as per adaptation along information theoretical principles, see Figure 2.21A, below).

A recent study (Wimmer, Hildebrant, Hennig & Obermayer 2008) attempted to elucidate the principle mechanism underpinning adaptation to sound stimuli in the auditory interneuron (AN2) of crickets, through a stimulus protocol that involved comparing adaptive stimulus-response functions to stimuli with bi- and tri-modal distributions and calculating the amount of mutual information the adaptive neural-response functions conveyed about the input distribution. Example results from three AN2 cells are given in Figure 2.21C, below.

The general observation was that, following adaptation to the tri-modal stimulus, the slope of the stimulus-response function tended to become shallower (relative to the bi-modal condition), as per the expectations of the efficient coding hypothesis, however the

amount of gain change in slope was not as large as expected.

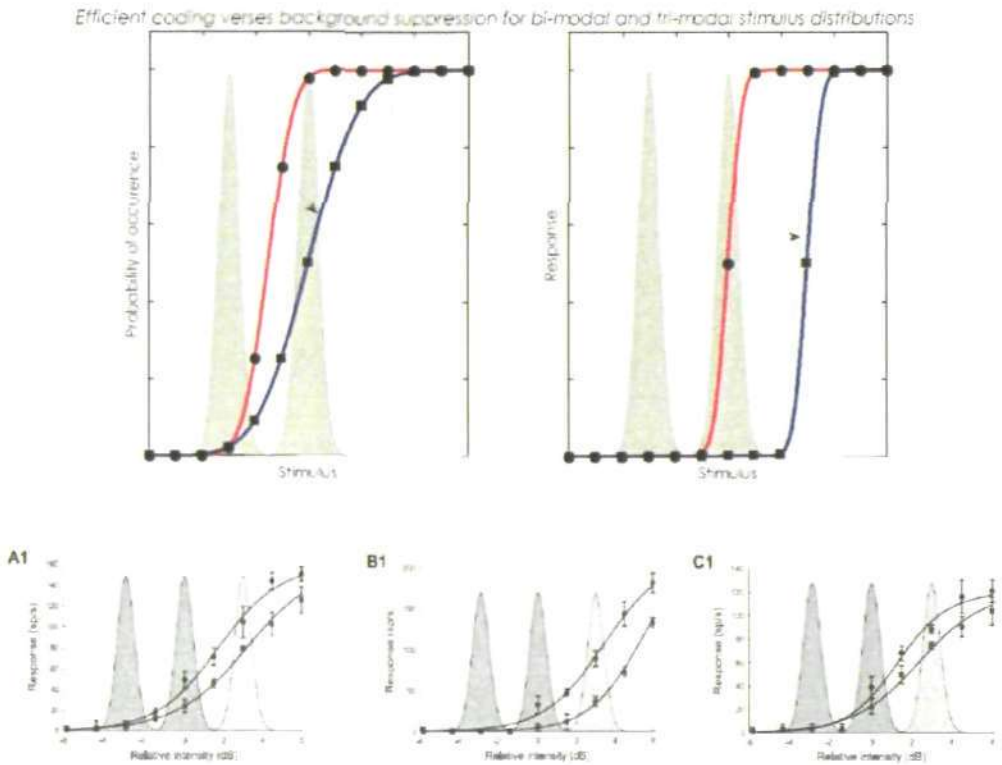

Figure 2.21: Efficient coding verses selective coding
Dark- and light grey regions represent adapting stimuli with bi- and tri-modal distributions, respectively. **A:** Expected influence of adaptation, mediated by an *efficient coding strategy*, on stimulus-response curves of a hypothetical neuron. Red line and circles gives expected response following adaptation to an arbitrary stimulus with a bi-modal distribution; blue line and squares give expected adaptive response to a tri-modal distribution. Probability of the distribution is plotted on the sinistral axis and hypothetical response on the dextral axis. **B:** As for A, but for *selective coding strategy*. A & B adapted from Wimmer et al. (2008). **C:** Actual experimental responses from three cricket AN2 neurons showing adaptation under the different stimulus distributions (circles: bi-modal; squares: tri-modal). Taken from Wimmer et al. (2008).

Lateral displacements of the stimulus-response function were larger and shifted thresholds (the half-maximal point or S_{50}) to stimulus intensities that were higher than expected according to information theory and resulted in stimulus-response functions that didn't fully cover the stimulus space; this resulted in a loss of information following

adaptation. The results were not compatible with a selective coding strategy either; whilst some background suppression was observed there was no change in information transmission for the higher intensity sounds, under both conditions.

One possible strategy underpinning adaptation in cricket AN2 neurons is to prime the cells for novelty detection. In general, the stimulus-response functions were shifted such that cells were still able to respond dynamically to stimuli with intensities higher than that found within the input distribution.

Novelty detection, as a result of stimulus-specific adaptation (SSA), has been observed (Ulanovsky, Las & Nelken 2003) in the feline auditory cortex (A1). SSA is also present in cat visual cortex (e.g. orientation adaptation; Dragoi et al. (2000), see below).

A standard mechanism for testing novelty detection is by using an oddball stimulus design (see Figure 2.22A, below), whereby rare/deviant stimuli are randomly inserted into a stimulus presentation of a common/standard stimulus. The target cell or cells are thus adapted to the standard stimulus.

Within cat A1, neurons undergo SSA that primes the cell to respond strongly to the rare stimulus, relative to the common stimulus, irrespective of the frequency values of the standard/deviant stimuli (Ulanovsky et al. (2003) see Figure 2.22C above). Here adaptation was to one of two possible frequencies of equal amplitude (the two frequency values, $f1$ and $f2$, were located well within the boundaries of the RF of the target A1 cell, as was the amplitude of the stimuli). The strength of the adaptive effect was dependent on both the difference in value ($\Delta f = f2 - f1$) between the standard/deviant stimuli (compare the three columns under the 90/10% condition, the effect decreases with the size of Δf) and on the ratio of presentation (e.g. if one compares the first and second columns in Figure 2.22C, Δf

was equal under both conditions but the effect was more pronounced under the 90/10% condition). Novelty detection, in this case, is similar in principle to the phenomenon of background suppression, outlined above, with the difference being that it is novelty rather than intensity that suppresses the response to the background (i.e. the standard stimulus).

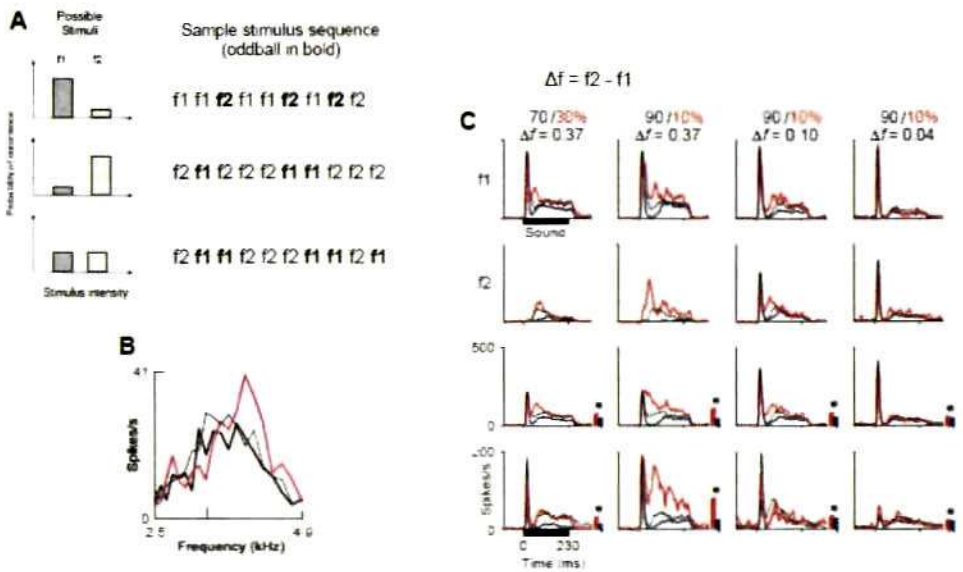

Figure 2.22: Stimulus-specific adaption (SSA) results in enhanced novelty detection **A:** Schematic illustration of odd-ball experimental paradigm (adapted from Ulanovsky et al. 2003). **B:** Frequency response curve of an A1 neuron before (black line), during (magenta) and after SSA (thin black line, recovery time of 30s). Adaptive frequency was 3.33 KHz. **C:** Responses to deviant (red), standard (blue) and control response (black, 50/50% presentation) for four stimulus configurations (columns, ratio of presentation and frequency difference (Δf) given above column). Top row plots responses for a single A1 neuron when *f1* was standard, second row plots responses for *f2*. Third row plots the average response across *f1* and *f2*; bars indicate spike counts and asterisk indicates significant difference in firing rate of deviant verses standard frequency. Bottom row plots averaged responses for another neuron.

The effect of frequency adaptation on the frequency tuning-curve of an A1 neuron is plotted in Figure 2.22B above. As can be seen adaptation results in a repulsive shift in

preferred frequency, away from the control response, and facilitation of the far-flank of the tuning curve (Condon &Weinberger 1991; Malone & Semple 2001; Ulanovsky et al. 2003). This is similar to the effect of adaptation on orientation tuning curves in V1 (Dragoi et al. 2000; Jin, Dragoi, Sur & Seung 2005, see Figure 2.23).

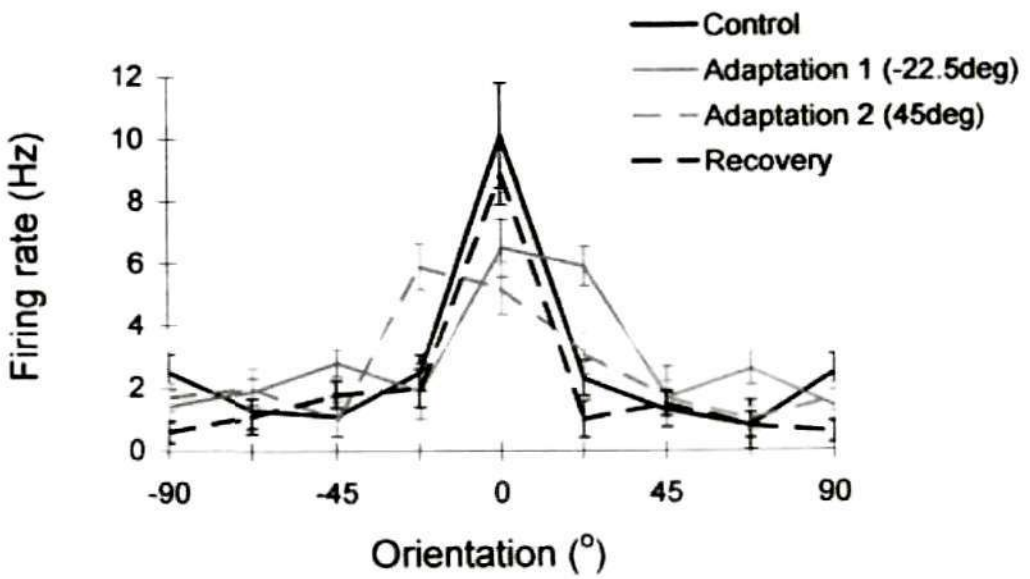

Figure 2.23: Orientation Adaptation in feline V1
Adapting orientations are marked with red arrows. Orientation tuning curves of a representative V1 neuron. Solid black line is the un-adapted response, with preferred orientation of 0o; recovery response is the dashed black line. Light- and dark-grey lines plot the adaptive orientation tuning curves. Adaptation results in a repulsive shift in preferred orientation and facilitation of the far-flank of the tuning curve, relative to the control response. Taken from Dragoi et al. (2000).

A recent study (Ringach & Malone 2007) indicates that novelty detection may not be limited to the action of a subset of neurons undergoing SSA but may be the basis for setting the operating point of the cortex as a whole. The responses of monkey V1 neurons to grating stimuli was modelled using a linear/non-linear model, whereby the linear input to the model was the standard deviation (SD) of the generator potential:

"a variable representing a combination of the history of intracellular variables yielding an overall measure of neuronal drive" (Ringach & Malone 2007).

The non-linearity was a half-rectifier with parameters: threshold, gain and input noise, see Figure 2.24 below for a description of model parameterisation.

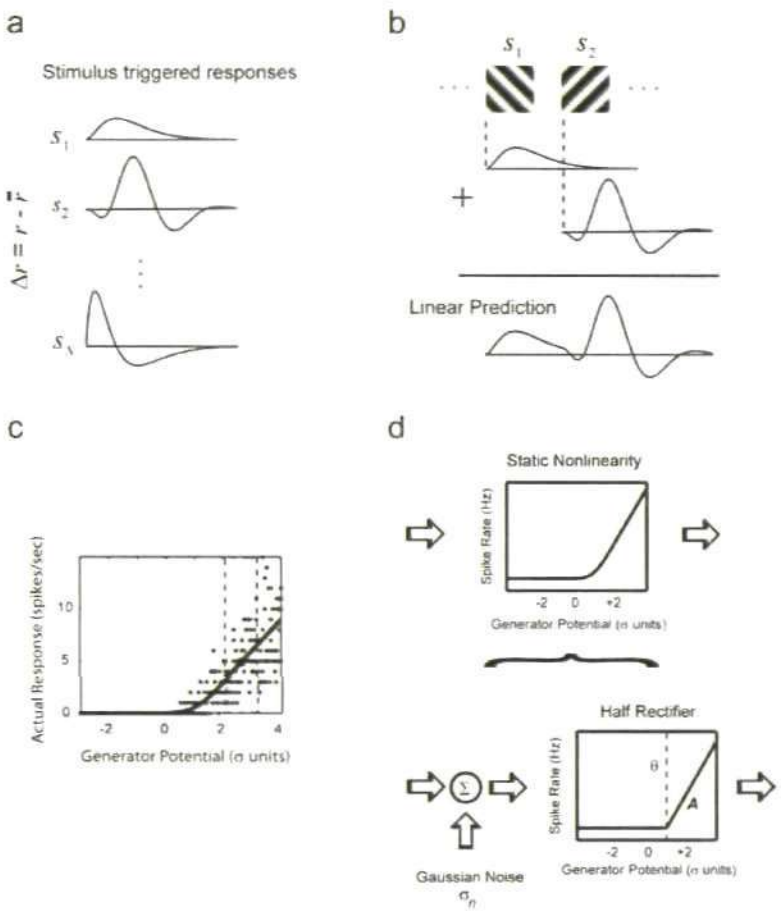

Figure 2.24: Estimation of the generator potential and non-linearity in a linear/non-linear model of V1

Plot **a:** The stimulus-triggered responses (r_i) about the mean response ($\bar{r}$) was calculated for each stimulus (s_i) in the set (S). Plot **b:** A linear prediction for a stimulus sequence was determined from the responses (r_i) and then normalized to have zero mean and unit variance. Plot **c:** A non-parametric estimate of the non-linearity (solid line) was found by smoothing the output of the generator potential and the actual responses within a given window (dashed lines). Plot **d:** The parameters: threshold (θ), gain (A) and input noise (σ_n) were determined by fitting a half-rectifier to the non-parametric estimate in plot c. Taken from Ringach & Malone (2007).

The general observation across the cortex was that the non-linearity threshold was roughly 1 SD from the mean of the generator potential and threshold was invariant to levels of stimulus contrast (e.g. with contrast levels of 25, 50 and 99%).

Due to the shape of the non-linearity, once the signal crosses spiking threshold it is amplified in a linear manner and there is no response saturation. Overall, spikes were generated when the generator potential fluctuations were ≈2 SD away from the mean activity of the generator potential. Thus, in monkey V1 at least, sensory cells appear to act as deviation detectors that amplify large signals at the expense of background noise.

As mentioned in the beginning of this section, information maximisation may only be optimal during early stages of visual processing when the full transmission of the signal is the priority. In general, V1 was always considered to be early part of the visual processing hierarchy, where only simple transformations of the signal occurred. As such, one might have expected that sensory encoding within V1 would be in line with the principles of information theory and the early evidence, from contrast adaptation studies for example, were in support of this hypothesis. However, recently a growing consensus has emerged that V1 is more than just a relay stage; the large number of feedback connections from higher cortices hint at a role in the complex processing of visual information (Angelucci, Levitt, Walton, Hupe, Bullier & Lund 2002; Angelucci & Bullier 2003; Bullier 2001), V1 receptive fields are not fixed but exhibit plasticity (e.g. Dragoi et al. 2000; Gilbert, Li & Piech 2009; Li, Piëch & Gilbert 2004) and even reward signals have been observed in rodent V1 (Shou, Li, Zhou & Hu 1996). Therefore V1 should be distinguished from 'lower' areas in the visual processing hierarchy, for which redundancy reduction is the

optimal strategy, as an area that implements a number of strategies depending on the types of processing mechanisms active at any given time (Schwabe & Obermayer 2005).

2.5: Fisher Information

The above discussion has focused entirely on Shannon's concept of information, i.e. that information is a reduction of uncertainty and is quantified by its entropy.

Essentially, mutual information is a measure of how certain one can be of the input to a system given the observed output, i.e. one is making an estimate of the input distribution; the accuracy of the estimate can be quantified by Fisher Information (named after the statistician and geneticist, R.A. Fisher, who formulated the concept of maximum likelihood estimation, Fisher 1922).

Formally, Fisher Information, $I(\theta)$, is the variance of the score, where the score is the first derivative of the negative-log of the likelihood function, $L(\theta \mid X)$, which gives the likelihood of a parameter, θ, being a certain value, given the set of observations, X (where X is dependent on θ). The likelihood function is so called as it represents the likelihood, not the probability, of θ, given X (as $L(\theta \mid X)$ will not always integrate to unity).

2.5.1: Likelihood Function

The conditional probability, $p(X \mid \theta)$, describes how a set of observations, X, may vary, for a given parameter, θ. However, the likelihood function, $L(\theta \mid X)$, describes how θ may vary, given a fixed set of observations, X; in essence, $L(\theta \mid X)$ is the converse of $p(X \mid \theta)$.

As an example one can consider a coin throw; the coin can only land on 'heads',

Hd, or 'tails', T, and one is interested in discovering whether the coin is biased towards landing on *Hd*; in this case the probability of the coin landing on *Hd* is equal to $p(Hd)$ and the probability of observing *T*, is equal to $p(T) = 1 - p(Hd)$. If one assumes that the coin is unbiased, i.e. $p(Hd) = 0.5$, then the conditional probability of the coin landing on *Hd* for a single throw is (see equation 2.8):

$$p(Hd \mid p(Hd) = 0.5) = 0.5.$$

And the likelihood of $p(Hd) = 0.5$ is:

$$L(p(Hd) = 0.5 \mid Hd) = 0.5.$$

Consequently, the likelihood of $p(Hd) = 0.5$ following two coin throws which land heads-up is:

$$L(p(Hd) = 0.5 \mid Hd, Hd) = p(Hd, Hd \mid p(Hd) = 0.5) = 0.25.$$

2.5.2: Maximum-Likelihood

In the above example, the likelihood for a single parameter value was discussed. If one wishes to estimate which parameter (from a set of possible parameter values) is most likely to generate a given set of observations, one can take the maximum, L_{MAX}, of the likelihood (as a function of all the possible parameters values within the set θ, see Figure 2.25), to be the best (or most likely) estimate, θ, of the true parameter. In the case of $p(Hd)$, θ is a probability and would thus range from 0 (the coin will never land on heads) to 1 (the coin will always land on heads).

The likelihoods as a function of $p(Hd)$, for single trials of different sequences of

coin throws (X) are plotted in Figure 2.25, below.

Each hypothetical sequence X comprises n number of coin throws, such that $X = \{x_1 \ldots x_n\}$ and $x_i \in \{Hd, T\}$. The assumption is that x_i is independent of all other coin throws in the sequence and that the probability of scoring either Hd or T is constant, thus X is an independent and identically distributed (i.i.d) Bernoulli trial. The likelihood can therefore be calculated using individual conditional probabilities for each coin throw:

$$L(\theta \mid X) = \prod_{i}^{n} p(x_i \mid \theta)$$
$$\theta^r (1-\theta)^{n-r} \qquad \textbf{2.21}$$

Where $p(X \mid \theta)$ has a binomial distribution, r is the number of Hd per sequence of n trials and $i = 1 \ldots n$.

In the top plot of Figure 2.25, $X_{1 \ldots n} = Hd$ and the maximum-likelihood estimate, $\hat{\theta}_{MLE}$, is always unity; as the sample number, n, increases the slope of $L(p(Hd) \mid X)$, about L_{MAX}, becomes steeper. In the middle plot of Figure 2.25, X contains both Hd and T; $\hat{\theta}_{MLE}$ tends towards $p(Hd) = 0.5$. Finally, in the bottom plot of Figure 2.25, the sample size is increased; $\hat{\theta}_{MLE}$ tends towards $p(Hd) = 0.5$ and the gradient of the slope $L(p(Hd) \mid X)$, about L_{MAX} increases, relative to the middle plot of Figure 2.25 .

The behaviour of the likelihood function about L_{MAX} is an indicator of how close $\hat{\theta}_{MLE}$ is to the true value of θ. A flat curve about L_{MAX} indicates that a change in θ would not significantly affect the probabilities of observing X (as the same observations would have a similar likelihood of being observed for a range of θ) and therefore $\hat{\theta}_{MLE}$ could be highly inaccurate; one would expect the $\hat{\theta}$ to be biased, i.e. the difference between $\hat{\theta}$ and

the true value of θ would not equal 0 (when $\hat{\theta} - \theta = 0$ the estimate is said to be unbiased) and one would also expect that $\hat{\theta}_{MLE}$, across trials, would be highly variable.

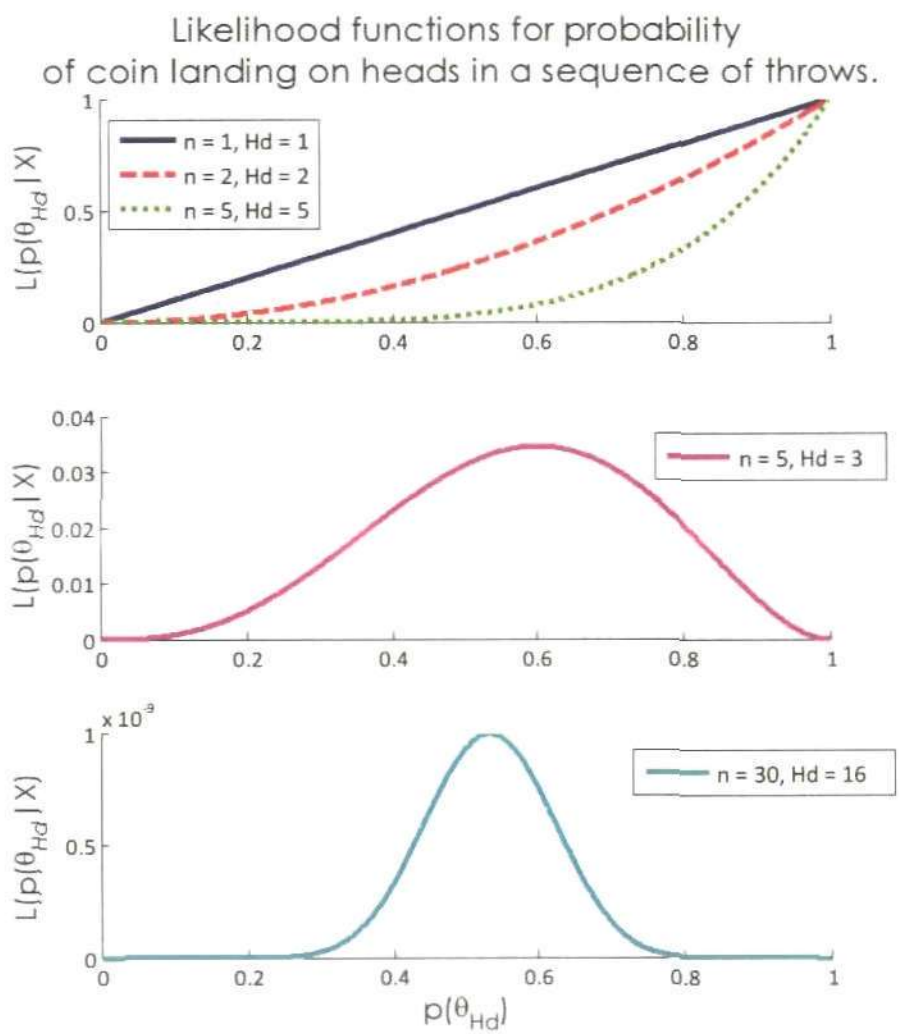

Figure 2.25: Likelihoods as a function of *p(Hd)*
Likelihoods as a function of the probability of a coin landing on 'heads', *p(Hd)*, given a hypothetical data set, *X*. For each *X*, the number of trials, *n*, and number of 'heads', *#Hd*, can be found in the corresponding legends. The number of 'tails', *#T*, in *X* is equivalent to *n*-*#Hd* and $p(T) = 1 - p(Hd)$. In the top plot, all values of *X* were set to *Hd*. In the middle plot, $n = 5$ (same number of trials as the green, dotted line in the top plot) but the *#Hd* was set to 3. For the bottom plot, $n = 30$ and $\#Hd = 16$.

The variance of $\hat{\theta}$ depends upon the number of trials and/or sample size of X. For example, the top and middle plots of Figure 2.25 display the likelihood functions of two hypothetical cases of X, in which $n = 5$ but the number of Hd ($\#Hd$) was either 5 (Figure 2.25: top plot, green dotted line only) or 3 (Figure 2.25: middle plot). Intuitively, one could argue that both cases of X could be easily elicited from the throws of an unbiased coin; however, both samples gave rise to different estimates of θ ($\#Hd = 5$, $\hat{\theta}_{MLE} = 1$; $\#Hd = 3$, $\hat{\theta}_{MLE} = 0.6$).

As trial and/or sample size increases one would expect $\hat{\theta}$ to converge upon its true value and that the variance of $\hat{\theta}$ would decrease, i.e. the estimate would become more accurate. A low variance estimator is said to be efficient and an optimal estimator is one in which the bias and variance of the estimate is minimised; Fisher Information provides a measure of quantifying how reliable and efficient the likelihood function is, as an estimate of the parameter set.

2.5.3: Fisher Information from the log-likelihood

Fisher information, $I(\theta)$, is defined as the variance of the score. Under mild conditions, .it also equals the 2nd derivative of the negative, log-likelihood function:

$$I(\theta) = -E\left[\frac{d^2\ell(\theta \mid x)}{d\theta^2}\right]$$

$$\ell(\theta \mid x) = \ln L(\theta \mid x) \qquad \textbf{2.22}$$

E is the expectation operator over the observations of x with respect to the likelihood distribution $L(\theta \mid x)$, as a function of x.

As equation 2.22 shows, the amount of Fisher information is directly related to the curvature of the averaged log-likelihood function about L_{MAX}. A sharp maximum in the log-likelihood has a large second order derivative and corresponds with high Fisher information. A flat maximum, conversely has low information. As discussed above, a steep curvature of the likelihood function about L_{MAX} indicates that the bias and variance of the estimate of θ are low; thus, there exists an inverse relationship between the variance of the score (the Fisher information) and the variance of the estimate of θ.

2.5.4: Cramer-Rao Bound

The inverse of the Fisher information is called the Cramer-Rao bound. The Cramer-Rao inequality (named after the mathematicians who derived it: Harald Cramér and Calyampudi Radhakrishna Rao) defines the lower bound of the variance of the maximum-likelihood estimate, MLE, of θ, $var(\theta_{MLE})$, such that $var(\theta_{MLE})$ can never be less than the inverse of the Fisher Information:

$$\mathrm{var}(\theta_{MLE}) \geq \frac{1}{I(\theta)} \qquad \textbf{2.23}$$

And the standard error (SE) of the estimator is:

$$SE(\theta_{MLE}) \geq \frac{1}{\sqrt{I(\theta)}} \qquad \textbf{2.24}$$

An optimal estimator is one which saturates this bound and is unbiased.

2.5.5: Fisher Information and the neural code

Generally, neuronal responses to single, simple features of environmental input (i.e.

contrast, orientation) can be quantified by plotting the average response as a function of the stimulus (i.e. can be quantified by the rate-function). However, neural responses tend to vary from trial-to trial. Therefore the ability of other neurons to decode or estimate the input from the average firing rate will be highly dependent upon the levels of variance in the response. Fisher information provides a measure of how well one can estimate the stimulus feature from the average firing-rate response by means of the Cramer-Rao bound.

As stated above, there are two main factors that determine how accurate and efficient an estimator is, namely the bias (the difference between the estimated and true value of the parameter) and the variance. The tuning curve is the mean response over many trials. One can assume that the tuning curve is, on average, an unbiased estimate (i.e. $E[\hat{\theta} - \theta] = 0$, where E is the expectation, across trials) of the input; the quality of the estimate is therefore dependent only upon the amount of output variability, observed in response to a given stimulus.

Neuronal noise is often assumed to be independent and normally distributed; with the experimental evidence (e.g. Tolhurst, Moshvon, Dean 1983; Gershon, Weiner, Latham, Richmond 1998) indicating that response variance is either equal to or proportional to mean levels of activity.

The probability, $p(r \mid \theta)$, of observing a response, r, for a stimulus parameter, θ, can be captured by the Gaussian distribution given in equation 2.25, where σ^2 is either fixed or equal to the mean firing rate, $f(\theta)$:

$$p(r \mid \theta) = p(r = r \mid \theta) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(\frac{(r - f(\theta))^2}{2\pi\sigma^2} \right) \qquad \textbf{2.25}$$

However, if one is counting spikes, then $p(r \mid \theta)$ is often best described by the

Poisson distribution:

$$p(r \mid \theta) = p(r = k \mid \theta) = \frac{f(\theta)^k e^{-f(\theta)}}{k!} \qquad \textbf{2.26}$$

where r is the response of a neuron, θ is the stimulus, k is the number of spikes per time window and $f(\theta)$ is the average response over all trials.

The Fisher information can be calculated, in terms of the noise distribution as:

$$I = -\int \partial r P(r \mid \theta) \left(\frac{\partial^2 \ln P(r \mid \theta)}{\partial^2 \theta} \right) \qquad \textbf{2.27}$$

The Fisher Information can be also approximated if one knows the noise distribution, such that for Gaussian noise, the Fisher Information is:

$$I = \frac{f'(\theta)^2}{\sigma^2} \qquad \textbf{2.28}$$

Where $f'(\theta)$ is the derivative of the stimulus-response function, $f(\theta)$, and σ^2 is the variance (which is either fixed or proportional/equal to the mean response). For Poisson noise, the Fisher Information can be approximated by (Seung & Sompolinsky 1993):

$$I = \frac{f'(\theta)^2}{f(\theta)} \qquad \textbf{2.29}$$

For the Poisson distribution, the variance is equivalent to the mean and the denominator, in equation 2.29, is the average response, $f(\theta)$. The population Fisher Information, $I_{P(\theta)}$, is calculated by summing $I(\theta)$, for each neuron in the population, under the assumption of independence:

$$I_{P(\theta)}(\theta) = \sum_{i}^{N} I_i(\theta) \qquad \textbf{2.30}$$

Where N is the total number of neurons in the population and $I(\theta)$ is the Fisher Information

for the *ith* neuron ($i = 1...N$).

From equations 2.28-2.29, it is clear that Fisher Information is greatest when the gradient of the tuning function is high and variance of the response, as a function of θ, is low. For those neurons that display bell-shaped tuning curves in response to stimulus features (e.g. orientation sensitive cells in V1), Fisher Information will peak on the slopes of the tuning curve (assuming that variance is equal or proportional to mean activity, see Figure 2.26 below).

Responses are therefore most informative (in terms of estimating the stimulus), for those stimuli which elicit intermediate responses, either side of the peak response, and not (as may be intuitively expected) for those stimuli which elicit the strongest response, e.g. the preferred orientation.

Fisher Information, then, is a measure of how well a neuron is able to discriminate between stimuli. As already outlined within this chapter, small changes in stimuli located on the steepest regions of the rate-function slope will result in relatively large changes in the response, with respect to stimuli located on the saturated region of the function. The ability of a neuron to discriminate between two stimuli, θ and $\theta+\Delta\theta$, can be quantified by a discriminability index, d', which is related to the Fisher Information by the equation (see Seung & Sompolinsky, 1993):

$$d' = \Delta\theta\sqrt{I(\theta)} \qquad 2.31$$

Equation 2.31 demonstrates that the ability of a neuron to resolve a stimulus increases with increasing Fisher Information.

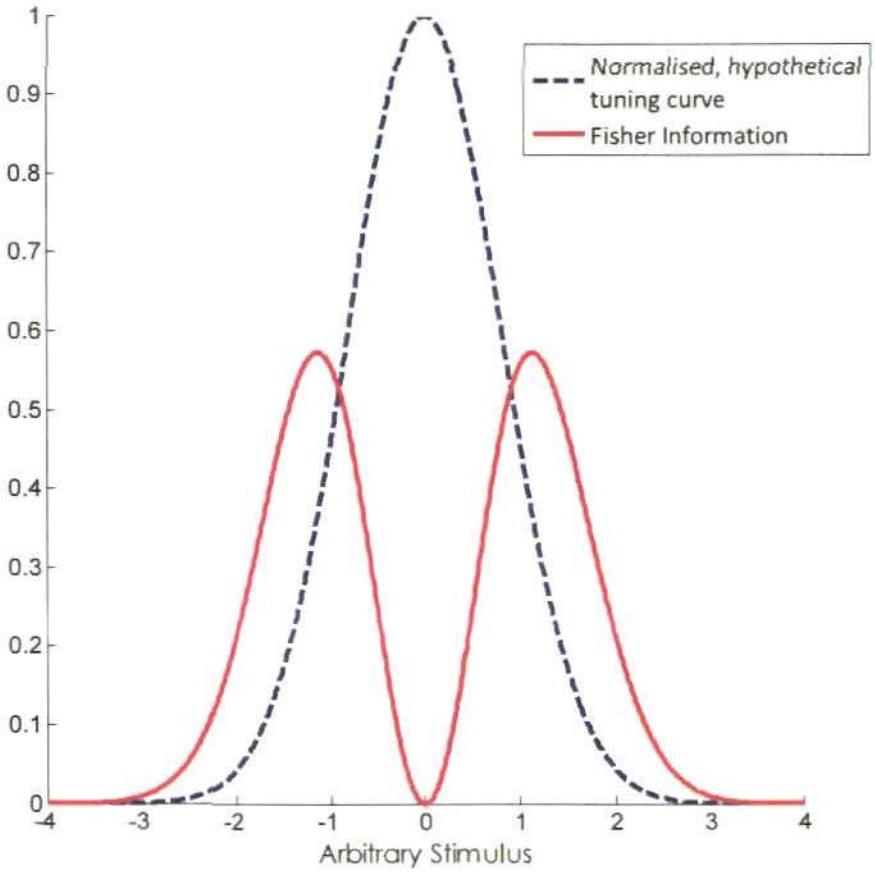

Figure 2.26: Fisher Information for a hypothetical single neuron with a Gaussian tuning profile
Blue, dashed lines plot the normalised response of a hypothetical neuron to an arbitrary stimulus (no units). Red solid lines plot the Fisher Information calculated as for Equation 2.28 and assuming that variance is equivalent to the mean firing rate. Fisher Information is peaked on the slopes of the tuning curve, not on the preferred stimulus, and does not reflect the direction of the slope of the function but its gradient.

The inverse of the square-root of the Fisher Information quantifies the discrimination threshold, or just-noticeable-difference in response, for the stimuli, θ and $\theta + \Delta\theta$; which if one refers back to equation 2.24 is the lower-bound for the standard error of θ.

For sigmoidal tuning-curves (e.g. in response to stimulus contrast), the Fisher Information would be expected to peak on the linear, accelerating region of the curve (see Figure 2.27); if the slope is centred on the mean of the stimulus distribution, as predicted by the efficient coding hypothesis, then neural output will be most informative, in terms of estimating and discriminating stimuli, for inputs that lie within the higher probability regions of the stimulus space.

2.5.6: Correlated Activity

The population Fisher Information, that is the sum of all individual Fisher functions across the population, is calculated under the assumption of independence (see Equations 2.28 - 2.30, above). Independence implies additivity, i.e. the joint noise distribution of two arbitrary neurons, $p_{1,2}(r \mid \theta)$, will be equal to the sum of the individual distributions:

$$p_{1,2}(r \mid \theta) = p_1(r \mid \theta) + p_2(r \mid \theta)$$

However, if the noise distributions $p_1(r \mid \theta)$ & $p_2(r \mid \theta)$ are correlated then the joint noise distribution will be less than the sum of the individual noise distributions:

$$p_{1,2}(r \mid \theta) < p_1(r \mid \theta) + p_2(r \mid \theta)$$

Clearly, if correlations between neural outputs are present, then calculating the population Fisher Information under an assumption of independence can lead to its over-estimation.

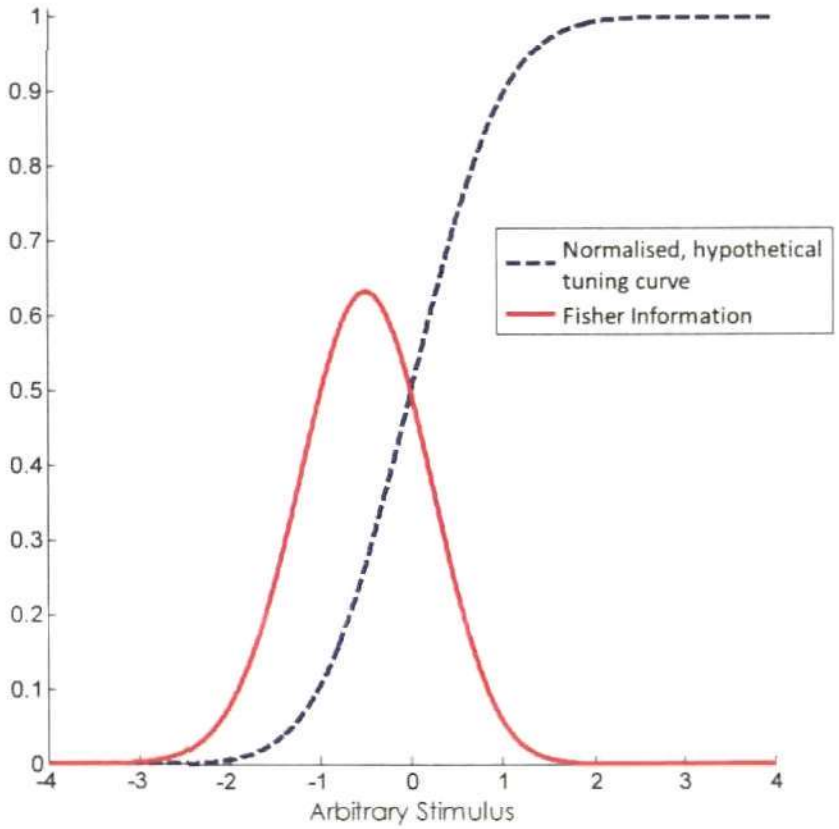

Figure 2.27: Fisher Information for a hypothetical neuron with a sigmoidal tuning curve
Blue, dashed lines plot the normalised response of a hypothetical neuron to an arbitrary stimulus (no units). Red solid lines plot the Fisher Information calculated as for Equation 2.28 and assuming that variance is equivalent to the mean firing rate.

In the case of correlations, Fisher Information can be calculated by taking into account the covariance matrix, $Q(\theta)$, of the neural noise; for example (Abbott & Dayan, 1999):

$$I = f'(\theta)^T Q^{-1}(\theta) f'(\theta) + \frac{1}{2}\mathrm{Tr}\left[Q'(\theta)Q^{-1}(\theta)Q'(\theta)Q^{-1}(\theta)\right] \qquad \textbf{2.32}$$

Where Tr is the trace operation and $Q'(\theta)$ is the derivative of the covariance matrix.

The effects of noise correlations on Fisher Information appear to be varied. For example, in a population with a large number of neurons, positive correlations can have the effect of decreasing Fisher Information whilst negative correlations have the opposite effect (Sompolinsky, Yoon, Kang & Shamir 2001).

The sharpening of tuning curves, for single stimulus parameters only, has been shown to increase the amount of Fisher Information contained in the population code (for a review, see Dayan & Abbott 2002); for example, sharpening of stimulus-response functions, within V1, can be achieved by narrowing the width (essentially increasing the slope) of the bell-shaped tuning-curves typically observed in response to stimulus orientation. However, a recent modelling study (Series, Latham & Pouget, 2004) has demonstrated that the introduction of local correlations (through lateral connectivity) into a V1 model of tuning curve sharpening actually decreased the amount of Fisher Information contained with the population.

There is certainly experimental evidence to suggest that neural response variability is strongly correlated (e.g. Zohary, Shadlen & Newsome, 1994), thus it would appear to be important to consider the correlations inherent in the population activity when calculating the amount of Fisher Information contained within it.

2.5.7: Comparison to Shannon Information

Both Shannon (in the form of Mutual Information) and Fisher Information provide quantitative measures of the encoding of stimulus into response. However, these measures belong to distinct theoretical fields and, as such, tell us different things about the relationship between input and output. Fisher Information is a form of statistical estimation analysis that describes the accuracy and efficiency with which an estimate of a parameter,

θ, describes the true value of θ, where the estimate is derived from a set of observations that are dependent upon θ. Mutual Information is derived from Information Theory and is an entropy, or uncertainty, measure that describes the mutual dependency between the stimulus and response distributions.

Mutual Information is bounded by the entropy of the stimulus ensemble, see Equation 2.11. It is often considered too difficult to assess for a large number of neurons (Kang & Sompolinsky, 2001). In general, within studies looking at populations of neurons, Fisher Information is often calculated as an alternative to Shannon Information (despite the differences in what each information measure says about the relationship between input and output) as it is relatively easier to quantify, especially under the assumption of independence, for example see Equations 2.28 to 2.30 .

However, a quantitative relationship between Mutual and Fisher Information has now been derived for certain neural systems (Brunel & Nadal, 1998; Kang & Sompolinsky, 2001). Building on earlier work (Clarke & Barron, 1990 & Rissanen, 1996), Brunel & Nadal's theoretical study (Brunel & Nadal, 1998) found that for low-dimensional stimuli in large populations the Mutual Information between the population activity and the stimulus becomes equal to the mutual information of an efficient Gaussian estimator and the input. As described above, an efficient estimator is one which reaches or saturates the Cramer-Rao bound, thus the Mutual Information can be calculated from the inverse of the Fisher Information; Mutual Information for continuous-valued inputs is also related to the log of the Fisher Information for large populations of neurons (Kang & Sompolinsky, 2001).

A relationship between Fisher Information and stimulus-specific information (Butts & Goldman 2006) has also been found for populations containing at least 50 neurons

(depending on the time window of integration used to calculate the information measures, (Challis, Yarrow & Series 2008). The Stimulus-Specific Information is an entropy measure (derived from Information theory) of a particular stimulus value, averaged over all responses; whereas Mutual Information is essentially the uncertainty of the whole stimulus ensemble minus its uncertainty once the responses are known, averaged over all stimuli and responses.

The above theoretical work implies that Shannon and Fisher Information are related, for large groups of neurons at least, and it is legitimate to calculate the information that any observed neural activity holds about the stimulus using Fisher Information.

2.6: Conclusion

In conclusion, Barlow's efficient coding hypothesis suggests that neural systems have evolved to encode sensory information in such a way as to encode incoming information with the minimum amount of resources (e.g. spiking activity) and redundancy; where reducing redundancy refers to the 'packaging' of information along neural pathways in such a manner that sensory stimuli can be encoded efficiently without necessarily losing information. This is achieved through sensory neurons having response profiles that reflect the inherent distributions of the parameters which they encode; with the optimal example being response profiles following the integral of the input distribution. However, sensory systems must be flexible and as such the Efficient Coding Hypothesis predicts that rather than responses being hard-wired there must be a mechanism for responding dynamically to changes in stimulus distribution. Adaptation, at the very least adaptation to visual contrast, has been shown to be one of these mechanisms. As such, adaptation is not simply a case of

neural fatigue but reflects the dynamic nature of sensory systems that must be flexible in order to cope with an ever changing world.

Whilst there is a wealth of evidence to support the Efficient Coding Hypothesis, it is clear that other coding strategies may be more 'efficient' depending on the processing stage within the cortex and the overall stimulating regime (c.f. the brief review of stimulus-specific adaptation in Section 2.6.8, above). The rest of this thesis examines an example of adaptation to stimulus statistics that, superficially at least, provides further evidence for the Efficient Coding Hypothesis; however, following a more robust analysis of the data it is clear that the adaptation serves to encode local changes in the stimulus, irrespective of the global mean.

Chapter 3: Adaptation to stimulus statistics

The previous chapter focused mainly on adaptation to contrast, within the visual system, as to date this remains one of the most widely studied aspects of gain control in the cortex. However, recent evidence has emerged that indicates that adaptation to intensity is present in other sensory systems (see below). The adaptive responses are in line with the principles of the Efficient Coding Hypothesis - but only to certain degree. In fact adaptation tends to centre the responsiveness of neurons onto features within the stimulus space that have a higher intensity than global mean levels.

3.1: Adaptation to stimulus statistics: Introduction

In a recent study, Dean and colleagues (Dean et al. 2005) developed a novel method for investigating adaptation using a dynamic stimulus (to mimic natural stimuli more accurately) that allowed for rate-functions to be calculated during the adaptation process. The main advantage of the stimulus configuration was that levels of variance, within the adapting stimulus, could be manipulated independently of the mean (see below), thus the purpose of the study was to measure the effects of adaptation to stimulus statistics on mid-brain neurons within guinea-pig inferior colliculus (IC).

3.1.1: Adaptation of Inferior Colliculus neurons to noise burst stimuli

The IC is an important nucleus in the sub-cortical auditory system. Cells were adapted using noise burst stimuli, in which only the amplitude (in decibels) of the noise burst was varied. Decibel levels were refreshed every 50ms and drawn randomly from a pre-determined set of possible values (along the interval 21-96dB in 1dB steps) with a highly skewed selection-probability distribution, such that for a restricted range of noise

amplitudes the probability of selection was 0.8 (the high probability region), and 0.2 for the rest (the low probability region). This gave the stimulus both a dynamic profile (in that decibel levels fluctuated rapidly) and a global-mean amplitude that was both centred- and dependent-on the location of the high-probability region in the stimulus space (see Figure 3.1a-c). Furthermore, by increasing the width of the high-probability region, while holding its location constant, it was possible to increase stimulus variance independently of the stimulus mean (please compare Figure 3.1c and Figure 3.2a).

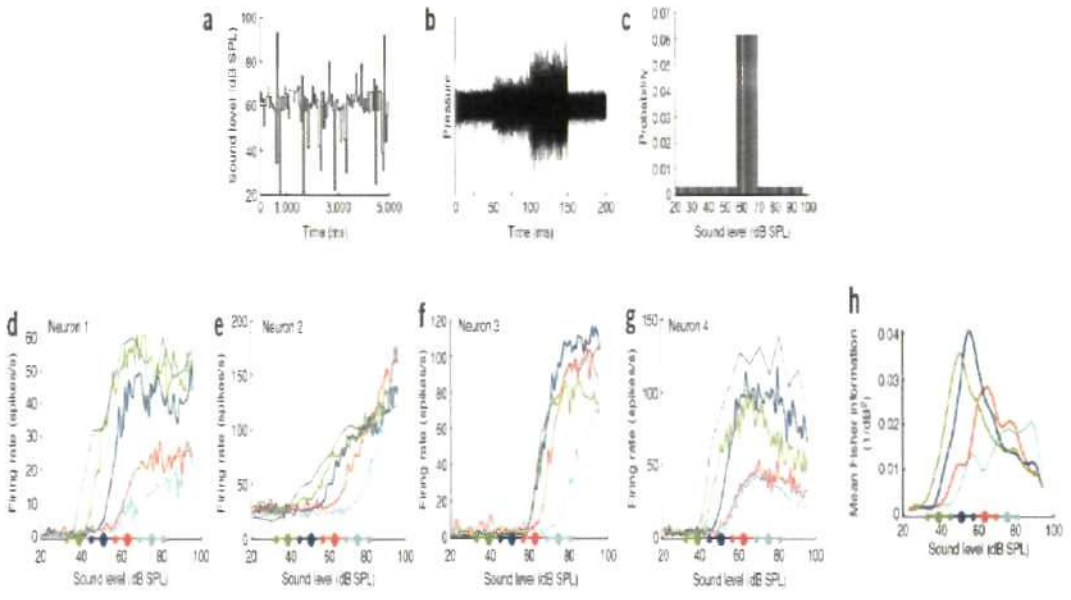

Figure 3.1: Adaptation to the global mean amplitude of noise bursts in inferior colliculus.
Plot (**a**): a 5s sample showing the variation in an adapting stimulus with an overall mean value of ≈60 decibels. (**b**): an actual, 200ms, example of the noise burst stimuli. (**c**): uneven stimulus distribution with the high-probability region centred on 63 decibels. (**d-g**): examples of rate-level functions in 4 neurons following adaptation to stimuli with a mean values of 39 (green), 51 (blue),63 (red) and 75 (cyan) decibels. (**h**): population Fisher functions. Filled circles and lines on the abscissa are the corresponding high-probability region regions. Taken from Figures 1 & 2, Dean et al. (2005).

Stimulus-response functions were calculated by taking the mean number of spikes fired (recorded extracellularly) during each 50ms epoch (with a 8ms response latency), that corresponded to a given decibel level. Cells were adapted at four different global mean amplitudes or high-probability regions (with width, and thus stimulus variance, held constant) and, for those cells that displayed adaptation, the stimulus-response functions were displaced laterally along the abscissa (see Figure 3.1d-g). This was in line with earlier single-cell contrast-adaptation studies in the visual system (e.g. Ohzawa et al. (1982); for further information see Chapter 2), which also observed a contrast-dependent lateral shift in stimulus response functions, following adaptation.

However, following adaptation in the IC, the slopes of stimulus-response functions were often centred on stimulus levels that were slightly higher in amplitude than the global mean of the adapting stimulus; this effect was most pronounced following adaptation to the lower amplitude, high-probability region conditions. Further, adaptation only occurred when the adapting high-probability regions were higher in amplitude than the region of the stimulus space covered by the pre-adapted response (c.f. Figure 3.1f).

In order to explore more accurately what the rate-level functions were encoding, Fisher information curves were calculated from each individual neuron and then for the population as a whole (by taking the sum of the individual Fisher-functions).

As discussed in the previous chapter, Fisher information is essentially a measure of how much information a set of observations, X, holds about the parameter set, θ. In terms of neural rate-level functions, the Fisher Information can be used to quantify how accurately one can estimate the value of the stimulus, given the average neural response. In general terms, one would expect Fisher Information to be highest when variance in the

spike count data, used to generate the mean response, is low (i.e. the response is reliable) and the curve of the stimulus response function is steepest, as small changes in the stimulus will elicit notably different responses. Fisher Information functions from individual neurons can be calculated when the distribution of the response noise, $P(r \mid s)$ is known (see Equation 3.1); the population response can then be estimated by summing the individual Fisher Information functions, across the whole population, but only under an assumption of independence, i.e. neuronal output, across the population, is uncorrelated.

For the IC adaptation data (Dean et al. 2005), the Fisher Information function, $f_a(s)$, of an individual neuron was calculated by taking the sum of the probability for each response elicited to a stimulus (1st term, equation 3.1), weighted by the curvature of the log-probability of the response, r, as function of the stimulus, s (2nd term, equation 3.1).

$$f_a(s) = \sum_r P_a(r \mid s)\left(\frac{d \ln P_a(r \mid s)}{ds}\right)^2 \qquad 3.1$$

Where $P_a(r \mid s)$ is the probability (P) of neuron a firing r spikes for stimulus s (see Chapter 2, also).

Correlations in spiking activity between pairs of neighbouring neurons were low (with an average correlation coefficient of 0.056 [STD of ±0.13] across all neuron pairs and stimulus conditions), thus independence in neural output was assumed and the population Fisher Information was estimated by taking the sum of the individual Fisher functions:

$$F(s) = \sum_a f_a(s) \qquad 3.2$$

The population Fisher Information was greatest (peak Fisher Information) for stimulus levels that were slightly higher than the global mean (see Figure 3.1h) of the

adapting stimulus.

The most intriguing finding of the study was a novel form of adaptation, observed in response to increasing stimulus variance only (the width of the high-probability region was widened but the stimulus level on which it was centred was held constant); the rate-level function was scaled upwards or increased in gain, i.e. the slope of the linear region became steeper and maximum responses increased (Figure 3.2a-b).

Why was this result intriguing? Following on from the discussion in the preceding chapter, it is clear from Information Theory that one would expect to see the stimulus response function slope become shallower, following adaptation to stimuli of increased variance, and not steeper as was actually observed.

However, it should be noted that this increase in gain was a rare response, the majority of cells showed either no adaptation (see Figure 3.2c) or a decrease in gain (resulting in a population Fisher curve (Figure 3.2d) that was slightly wider at the tails of the function, relative to the control).

Nevertheless, it can be argued that these results do present a challenge, albeit a small one, to the orthodoxy of Barlow's Efficient Coding Hypothesis. The authors of the study (Dean et al. 2005) conjectured themselves that the adaptive response to changes in the global mean of the stimulus may serve to prime the cell to act as a novelty detector, over faithfully encoding stimulus information.

It has been suggested (Dean et al. 2005) that the observed adaptive responses to stimulus statistics might be due in part to intrinsic mechanisms within the auditory pathway; therefore a more recent study Garcia-Lazaro et al. (2007) used a similar experimental design, this time in rat barrel cortex, to test whether adaptation to stimulus

statistics was a universal phenomenon or restricted to the auditory modality only. Here, the adapting stimulus was the deflection amplitude of the primary whisker; all other stimulus variables (e.g. direction, deflection frequency) were held constant.

3.1.2: Adaptation of rodent, cortical somatosensory cells to whisker deflection amplitude

A full description of the stimulus conditions is given in Garcia-Lazaro et al. (2007).

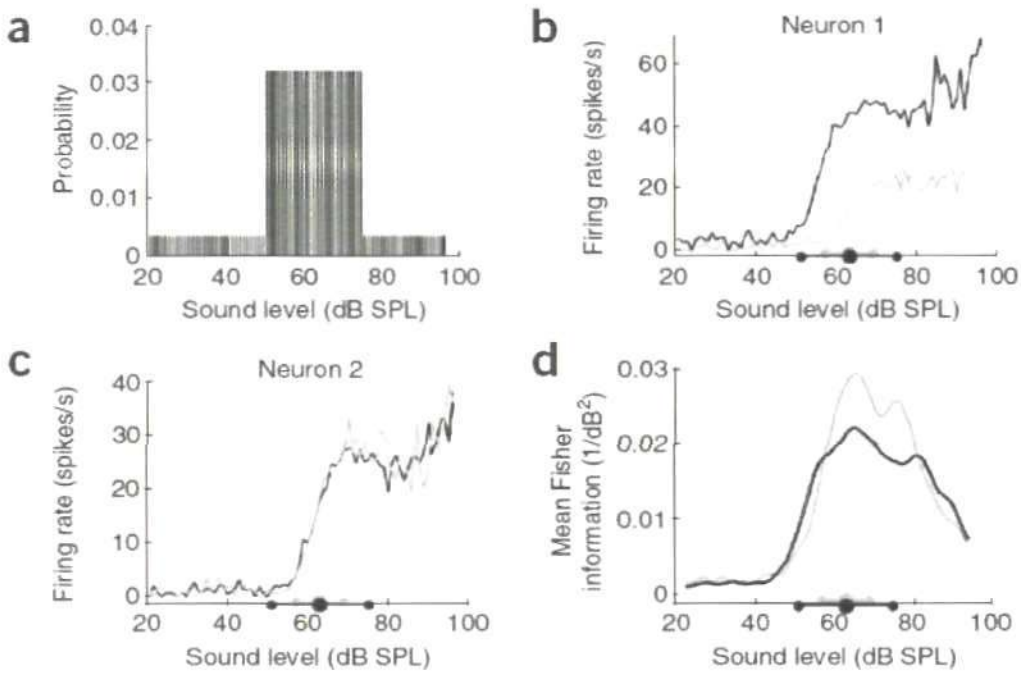

Figure 3.2: Adaptation to increased stimulus variance in inferior colliculus
Plot (**a**): stimulus distribution, high-probability region is centred on 63 decibels but wider than stimulus distribution in Figure 3.1c, thereby increasing stimulus variance independently of the stimulus mean. (**b**): example of a neuron showing adaptation. (**c**): example of a neuron showing no adaptation. (**d**): Population Fisher function. For plots b-d, the filled circles and lines on abscissa plot the extent of high-probability region; light and dark grey lines represent low- and high-variance, respectively. Taken from Figure 3 in (Dean et al. 2005)

In brief, each adapting stimulus sequence lasted 30s, whilst the amplitude of whisker deflection was selected randomly from a predetermined set of 25 amplitude values (equal, integer intervals of between 25-500μm for the first animal and 30-750μm for the rest) and refreshed every 40ms.

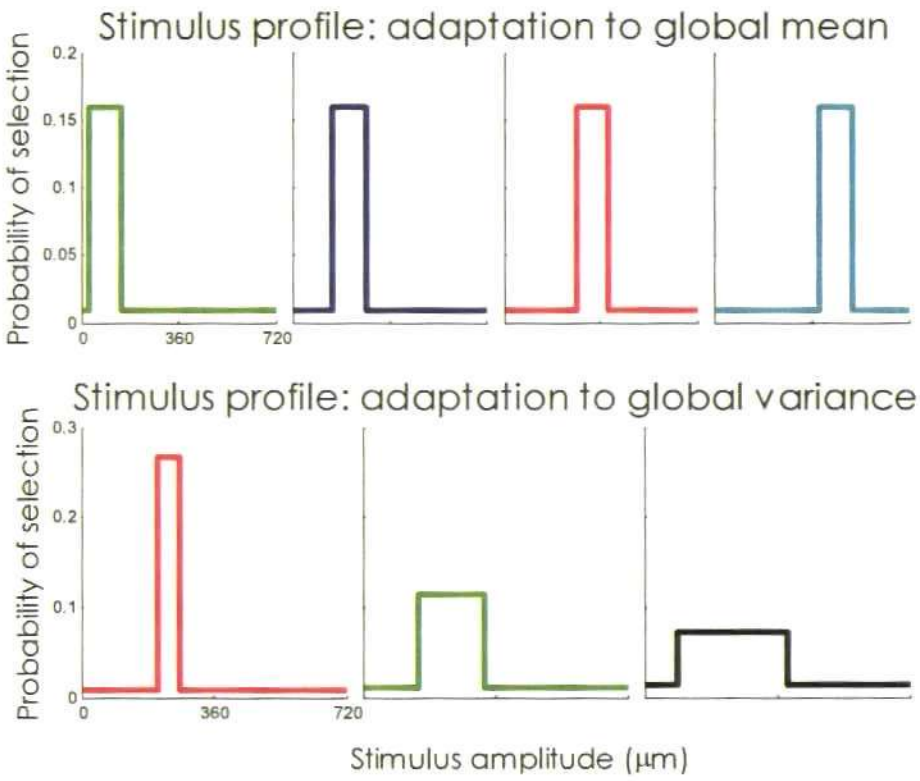

Figure 3.3: Stimulus distributions, under adaptation to mean and variance.
Plot (A): Adaptation to mean stimulus amplitude; mean stimulus amplitude is dependent on the high-probability region (high-probability region) which is shifted along the stimulus axis (green – 90μm, blue – 210μm, red – 330μm cyan – 450μm). (B): Adaptation to stimulus variance; the high-probability region is wider under each condition, increasing the number of amplitudes most likely to be selected and this increasing noise within the stimulus (red - 210-270μm/Low variance, green - 150-330μm/Mid variance, black - 90-390μm/High variance

As in the study outlined above (Dean et al. 2005), the probability distribution of amplitude selection was highly non-uniform (see Figure 3.3) and consisted of a high-

probability region (high-probability region) where the probability of selection equalled 0.8 and a low-probability region (low-probability region) where selection probability was 0.2. Thus, the global mean and variance of the stimulus could be manipulated independently by varying the location (mean, Figure 3.3A) or the width of the high-probability region (variance, Figure 3.3B). The data used to calculate the rate-functions was collected during, rather than post, adaptation. For every given stimulus amplitude, the adaptive neural-response was taken to be the average number of spikes elicited by the input, within the 40ms time bin following stimulus onset (given a response latency of 20ms).

Figure 3.4A shows how the responsiveness of a representative barrel cortex neuron was affected by adaptation to the different high-probability region levels. Following adaptation, the rate-function was shifted laterally (relative to the un-adapted response; black line and circles in Figure 3.4A), with the degree of displacement dependent on the global, averaged amplitude of whisker deflection. However, the adaptive neural-response functions and thresholds (stimulus amplitude at the half-maximal response, S_{50}) were not centred on the prevailing stimulus levels but onto regions in the stimulus space that were of higher amplitude than the global mean of the adapting stimulus. The Fisher Information of the response was calculated for each neuron and across the population as whole (for the method of calculation see Section on data analysis, below).The peak of the Fisher function (Figure 3.4B) was also displaced, relative to the global mean of the adapting stimulus, such that the cell was most sensitive to stimuli just outside the adapting high-probability region (assuming that Fisher Information is a measure of neural sensitivity). This type of adaptive response was typical, as can be observed from the population Fisher Information (Figure 3.4C) and the lateral displacement of S_{50} values across the population (Figure 3.4D-E). In

general the relationship between slope and firing rate was roughly linear and remained fairly constant for all adapting high-probability region regions (Figure 3.4F).

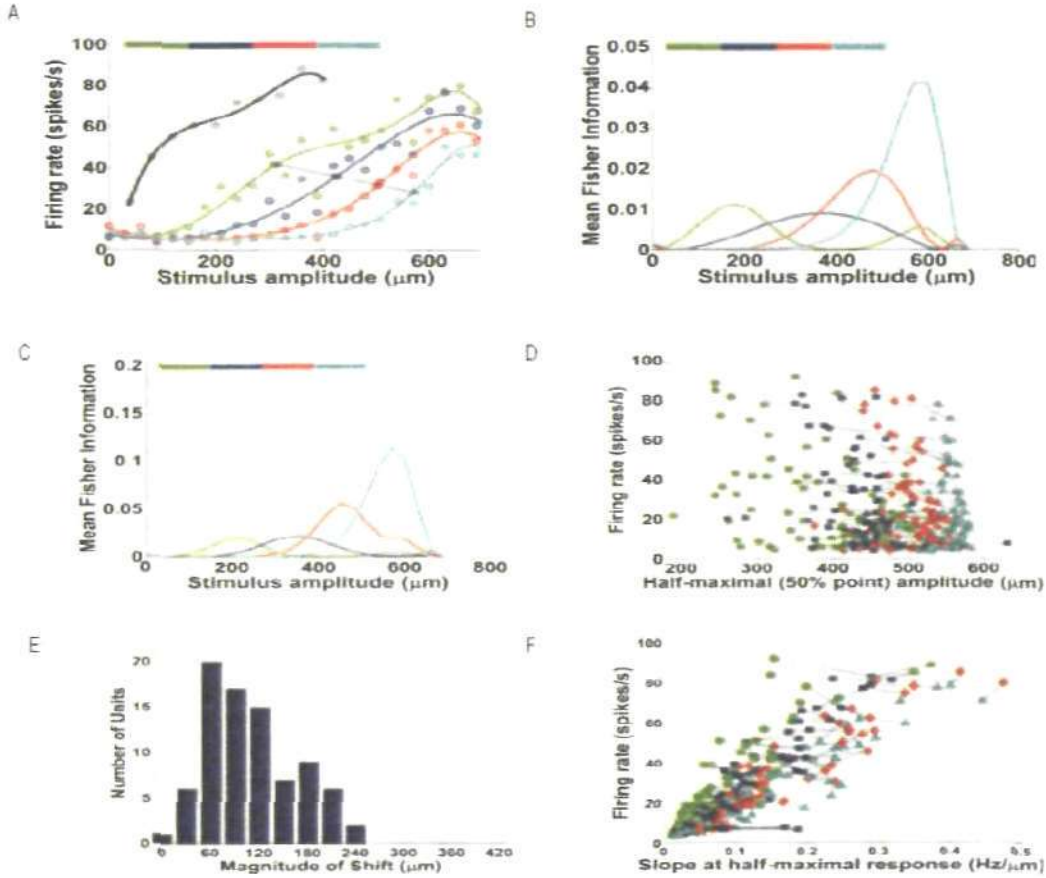

Figure 3.4: Adaptation to the global mean level of whisker deflection.
A: The rate-functions of a single neuron to 40ms bursts of whisker vibration presented in isolation (black line/circles, 200Hz frequency, presented at 1s intervals) and following adaptation to four high-probability region levels (green, blue, red and cyan lines/circles, colour coding as for stimuli in Figure 3.3A); empty circles plot actual spike counts for each amplitude, whilst smooth lines plot best-fit to the data, filled circles are the S_{50} constant. **B**: Corresponding Fisher functions for the neuron in A. **C**: Population Fisher function, pooled across all neurons. **D**: Scatter plot of relationship between stimulus amplitude and firing-rate at S_{50} for all units in the population. **E**: Magnitude of shift in amplitude at S_{50} between the 1st (red) and 3rd (green) high-probability region condition, for all units. **F**: Scatter plot of relationship between slope and firing rate at S_{50}, across population. Taken from Garcia-Lazaro et al. (2007).

In terms of adaptation to stimulus variance, the results from a single neuron are

given in Figure 3.5A-B. For this cell, as stimulus variance increased (Figure 3.5A; red lines/circles give the adaptive response to the low-variance stimulus, green = mid-variance and black = high-variance), the slope of the rate-function became steeper and the maximum response increased; the stimulus amplitude at the S_{50} also decreased.

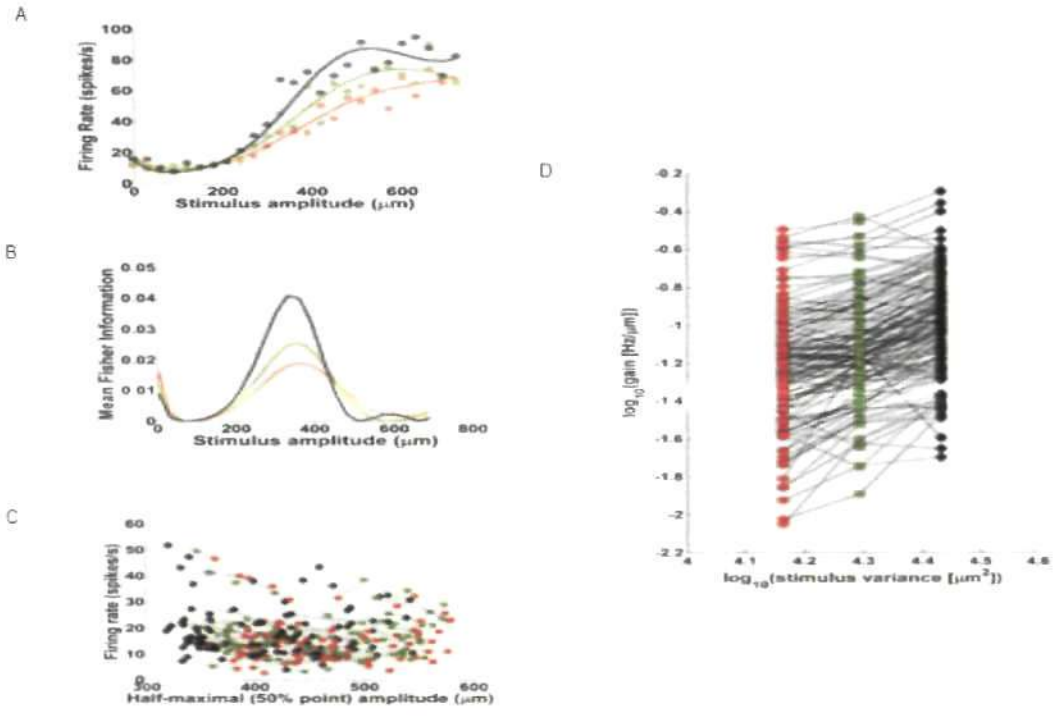

Figure 3.5: Adaptation to stimulus variance in rat barrel cortex (taken from Garcia-Lazaro et al. 2007)
A: Neural rate-function for a single neuron following adaptation to stimuli of increased variance, but constant mean; colour-coding: red, green and black are low-, mid- and high-variance, respectively, circles are the actual spike-counts and smooth lines are the best-fit to the data. **B**: Corresponding Fisher Information for the neuron in A. **C**: Scatter plot showing relationship (circles, colour coding as for A) between S_{50} and firing rate for all cells tested (n = 131), black lines connect up S_{50} values for each cell, under the three variance conditions. **D**: Scatter plot of slope at S_{50} (or neural-response gain), for all cells tested, under each level of stimulus variance (colour coding as in A, black lines are connecting the slope values for each neuron).

This cell clearly underwent an adaptive increase in gain, in response to increased variance within the adapting stimulus. The bandwidth of the corresponding Fisher function

(see Figure 3.5B) became progressively narrower as stimulus variance increased, implying that (for this cell at least) the adaptive response to increased levels of stimulus variance tends to improve stimulus resolution for whisker deflection amplitudes around the global mean, relative to lower variance conditions.

Using the change in stimulus amplitude and gradient of the stimulus response function at the S_{50} point as a measure of gain change, it would appear that an increase in gain (decrease in S_{50} amplitude and increase in slope) was the typical adaptive response to increases in stimulus variance (see Figure 3.5C-D). The general trend was for S_{50} amplitudes to be displaced leftwards (i.e. to lower S_{50} amplitudes, see Figure 3.5C) and stimulus response function gradients to become steeper (Figure 3.5D).

Thus, these results are broadly in line with adaptation to stimulus statistics in the IC (Dean et al. 2005), with two important exceptions:

1. Increased gain, following adaptation to stimulus variance, was apparently more common in the barrel cortex than the IC and

2. Adaptation to the stimulus mean displaced the stimulus response function further than observed in IC.

These two studies reveal that adaptation to stimulus statistics is an amodal phenomenon, as it is elicited by both auditory and somatic stimuli. However, this type of adaptation is more pronounced in the barrel cortex. Whether this is due to differences in cortical and sub-cortical adaptation or the modality of stimulus presentation cannot be answered here; what can be stated is that these adaptive responses conflict with the principles of Barlow's Efficient Coding Hypothesis.

3.2: Re-analysis of experimental data showing adaptation to stimulus statistics in rodent barrel-cortex

In order to explore this more thoroughly, Jan Schnupp and colleagues kindly provided published experimental data (from Garcia-Lazaro et al. 2007) to allow for further analysis of the population rate codes under the two types of adaptation regimes.

3.2.1: Experimental details, methodology & stimulus parameters

The experimental details of the data collection are taken from Garcia-Lazaro et al. (2007) and given below:

Experimental details

The results from Garcia-Lazaro et al. (2007), outlined above, were derived from the pooled results taken from two mature Long Evans rats; however, for the purpose of this analysis, only data from the single animal (LE270705) that was used in both experimental set-ups (i.e. adaptation to stimulus mean and variance) was analysed.

The animals were anaesthetised using halothane, during initial surgery, and ketamine, during recording (for full details of the experimental set-up, such as surgical procedures etc, please see Garcia-Lazaro et al. 2007). A 4x4 silicon 'Michigan probe' 2.5MΩ electrode array was used to record neural responses from the right barrel cortex. Signals from the array were band-pass filtered (between 0.5-3 kHz), amplified and digitized (25 kHz) using specialist software (Brainware, Tucker Davis Technologies, Alachua, FL) and analysed using MatLab (The MathWorks, Inc., Natick, MA, USA).

Layer 4 neurons of the rat somatosensory cortex are organised into distinct anatomical and functional areas, which resemble barrels (in terms of their shape) when

stained with cytochrome oxidase. Neurons within these areas respond preferentially to a given vibrissa, or principle whisker, of the mystacial pad. Potential principle whiskers where identified initially by the location of the electrode, within the cortex, and then confirmed by the response of neurons to stimulation of the suspected principle whisker and its neighbours. Stimuli of varying amplitudes (see below for a description of stimulus parameters) were presented in isolated, 40ms, bursts. The vibrissa that elicited the strongest response was deemed the principle whisker. This method of stimulation (i.e. presenting stimuli in isolation) was also used to determine the latency of the onset response (which was found to be, on average across the population, 20ms) and the un-adapted stimulus-response function for each individual neuron. Units with less than 1Hz firing rate were excluded from the analysis.

Whiskers were stimulated by attaching them, via a pair of glass capillaries, to a loud-speaker diaphragm which vibrated sinusoidally at 200Hz. The amplitude of oscillation was refreshed every 40ms and transitions were smoothed by a 3ms wide, running average, *filter. Amplitudes were drawn from a selection of 25 possible values ranging from 0-*720μm, in 30μm steps.

As outlined earlier, the probability distribution of amplitude selection was highly non-uniform and consisted of a high-probability region (high-probability region) where the probability of selection equalled 0.8 and a low-probability region (low-probability region) where selection probability was 0.2. Thus, the global mean and variance of the stimulus could be manipulated independently by varying the location (mean) or the width of the high-probability region (variance).

Stimulus parameters

The stimulus parameters were as follows:

In terms of adaptation to the mean, each high-probability region spanned 5 stimulus amplitudes and was centred on 90, 210, 330 or 450μm (Figure 3.3A), actual mean amplitudes were 161 μm (standard deviation or SD: 170μm^2), 252μm (SD: 136μm^2), 343μm (SD: 116μm^2) & 433μm (SD: 117μm^2) respectively. Differences in standard deviations were due to the skewed nature of the stimulus distribution when the high-probability region was centred on the lower stimulus amplitudes (see below for a further discussion). For adaptation to stimulus variance the high-probability region was fixed at 240μm and spanned amplitude ranges (with 30μm increments): 210-270μm (Low-variance condition), 150-330μm (Mid-variance condition) and 90-390μm (High-variance condition); with mean amplitudes of 272,279 & 290μm, respectively.

The experimental data, provided by Jan Schnupp and colleagues, consisted of spike times, from individual cells, elicited during 30s long stimulus presentations. Each trial consisted of a continuous series of 750, 40ms, amplitude presentations. For every high-probability region condition (under a given adaptation regime), there were three stimulus sequences and these were repeated at least five times. The order of presentation for a given sequence was randomly predetermined and distinct from the stimulus order of the remaining two sequences. Trials within and between high-probability region conditions were randomly interleaved with a 2s latency between the 30s presentations.

Data analysis

For each amplitude presentation, all spikes elicited within a 40ms time-window (located 20ms after stimulus-onset to account for response latency) were summed into

40ms time-bins and used to calculate the response to the given whisker stimulus. Stimulus-response functions (stimulus response function) for each high-probability region condition were determined by taking the average firing rate for a given whisker deflection amplitude, calculated across stimulus sequences and their repeats.

The half-maximal (S_{50}) amplitudes, gradient and firing rates were taken from the best-fits to the stimulus response function. The best-fit to the data was calculated by fitting a fifth-order polynomial to the data, using Poisson regression with an exponential link function (this was achieved by using the MatLab function *glmfit* and MatLab code generously provided by Jan Schnupp and colleagues).

Population Fisher functions were calculated using the methodology of Garcia-Lazaro et al. (2007). The estimated Fisher information, f, for a given cell, a, and stimulus, s (see Equation 3.3 below), was taken to be the differential, y', of the best-fit to the rate function divided by the best-fit to the spike-count variance, σ (see Figure 3.6, below). As mentioned above and discussed fully in Chapter 2 (see Section: 2.6), Fisher Information is dependent upon both the slope of the neural stimulus response function and variance in the response. For Gaussian noise, variance is assumed to be either fixed or proportional/equal to mean activity; thus Equation 3.3 gives the approximation of the Fisher Information when assuming Gaussian noise.

$$f_a = y_a'(s)^2/\sigma_a^2 \qquad \textbf{3.3}$$

An example of the best-fit to spike count variance, from a single neuron, can be seen in Figure 3.6, below. Response noise increased in proportion to mean activity, which was the typical response across the population. The proportional increase was seen for all high-probability region levels, thus was irrespective of the variance in the stimulus

distribution (variance tended to decrease as global mean amplitude increased, due to the stimulus configuration; see Section on Stimulus Parameters, above). The best-fit to the spike-count variance was calculated by fitting a fifth-order polynomial to the data (this was achieved by using the MatLab functions *polyfit* and *polyval*).

The population Fisher information (F) was taken to be the mean of the Fisher Information functions across all cells (N) tested:

$$F(s) = \frac{1}{N}\sum_{a} f_a(s) \qquad 3.4$$

One can only calculate the population Fisher information by Equation 3.4 under the assumption of independence, i.e. one must assume that the activity of all neurons in the population is uncorrelated. In order to test this assumption the spike-count data was analysed by plotting cross-correlation histograms for all possible neuron pairs in the population. Cross-correlograms were constructed for all spike-trains elicited in response to the same 30s stimulus presentation, across all neurons in the population, in a pair-wise manner. For each neuron pair (one was arbitrarily assigned as the target neuron and the other as the reference neuron), cross-correlograms were constructed from corresponding target and reference spike-trains elicited to a given 30s stimulus sequence and trial. The cross- correlograms were normalised, such that a perfect, positive, correlation at the zero time lag would be equal to 1; a time delay of up to ±100ms was used. This was achieved by using the MatLab function *xcov* (with the normalisation parameter set to *'coeff'*, which returns correlation coefficients for each time delay by normalising the data such that a perfect, positive correlation at the zero time-lag is equal to unity); the function *xcov* first subtracts the mean from the inputs vectors (i.e. the reference and test spike-trains of the

neuron pairs) and then calls the MatLab function *xcorr*, which calculates the cross-correlation function of the spike-trains.

The cross-correlation function counts the number of times in which spikes from the target and reference trains occur within a given time-lag of each other (in this case time delays where assessed from 0 to ±100ms) with a strong peak indicating that firing rates are correlated. On average, the correlation-coefficient at the zero time-lag, across all spike trains, high-probability region conditions and neuron pairs was 0.0137, with a standard deviation (SD) of ±0.0226; the peak correlation-coefficient was, on average, 0.0200 ±0.0242 SD. Even with the relatively large deviations in coefficients observed across the population, the correlation coefficients were low enough to assume independence.

Adaptation to global mean stimulus amplitude

As mentioned, only experimental data from a single animal was used in this analysis. Thus, it seemed prudent to first confirm that the adaptive response to increases in mean amplitude was broadly in line with the pooled, population response outlined in Section 3.1.2 above (see Figure 3.4).

Following adaptation, there was a rightward shift in the population rate-function as the mean of the adapting stimulus was increased (Figure 3.7A). This was in line with the general observed trend as reported in Garcia-Lazaro et al. (2007). Rate-functions were displaced laterally to centre on stimulus amplitudes that were stronger than those within the high probability region of the stimulus space.

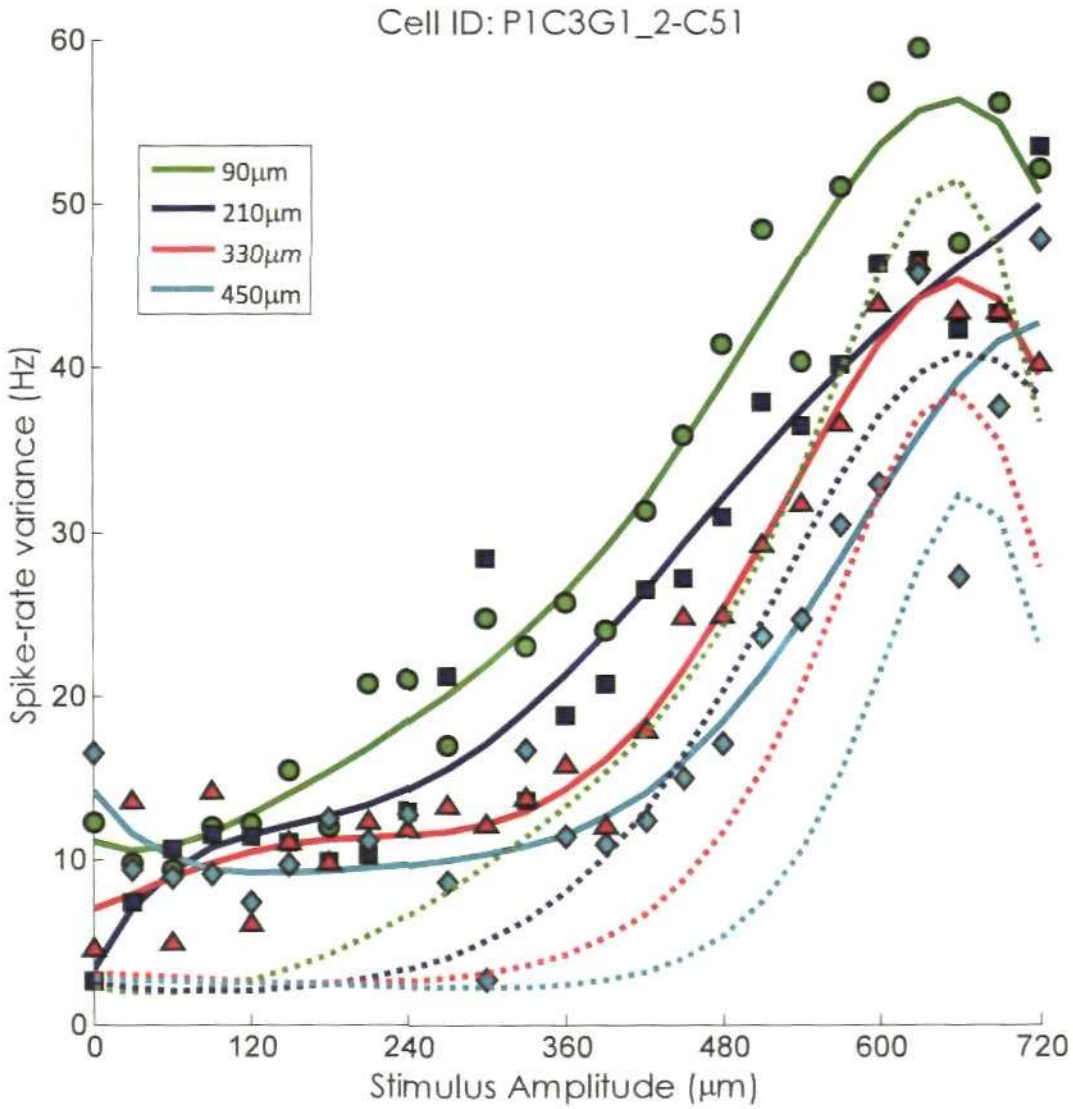

Figure 3.6: Best-fit to the spike-rate variance of a single cell
This figure plots the variance in the spike-count response, as a function of stimulus amplitude, of a single cell to all 4 adapting high-probability region conditions. Solid lines give the best fit to the actual spike-count variance data (symbols); with green line/circles plotting spike-count variance as a function of the stimulus for the 90µm condition, blue line/squares: 210µm, red line/triangles: 330µm & cyan line/diamonds: 450µm adapting stimulus condition. For comparison, the corresponding colour-coded dotted lines plot the corresponding best-fits to the average spike-count data (see Figure 3.7), as a function of stimulus (i.e. the stimulus response function).

The population Fisher function (Figure 3.7B) was bimodal for the lowest mean amplitude condition (90μm: green line), in that there were two peaks: a broad peak centred on the mean adapting stimulus and a main peak, with a narrower bandwidth, centred on 570μm. This amplitude was stronger than the stimulus that elicited peak Fisher information under the 210μm condition (blue line). Distribution was unimodal for all the other stimulus levels and bandwidth narrowed with increased mean stimulus amplitude.

Under the 90μm mean amplitude condition, the low-probability 'tail' covering amplitudes greater in strength than the high-probability region was longer than the other tail (see Figure 3.3), resulting in a strongly asymmetric stimulus distribution. Thus, under this condition, the adapting stimulus, relative to the higher mean amplitude conditions, had more inherent variance and this could account, in part, for the bimodal nature of the Fisher curve. The high slopes at the extremes of the Fisher Information functions are an artefact of the decrease in slope observed in the corresponding stimulus-response functions once maximum firing rate is reached.

3.2.2: Adaptive population stimulus-response function and Fisher Information function in response to increased levels of stimulus variance

The experimental evidence, as presented in Garcia-Lazaro et al. (2007), implied that increased gain was the typical response to increased stimulus variance; however this was not explicitly stated, i.e. no averaged, population data was shown.

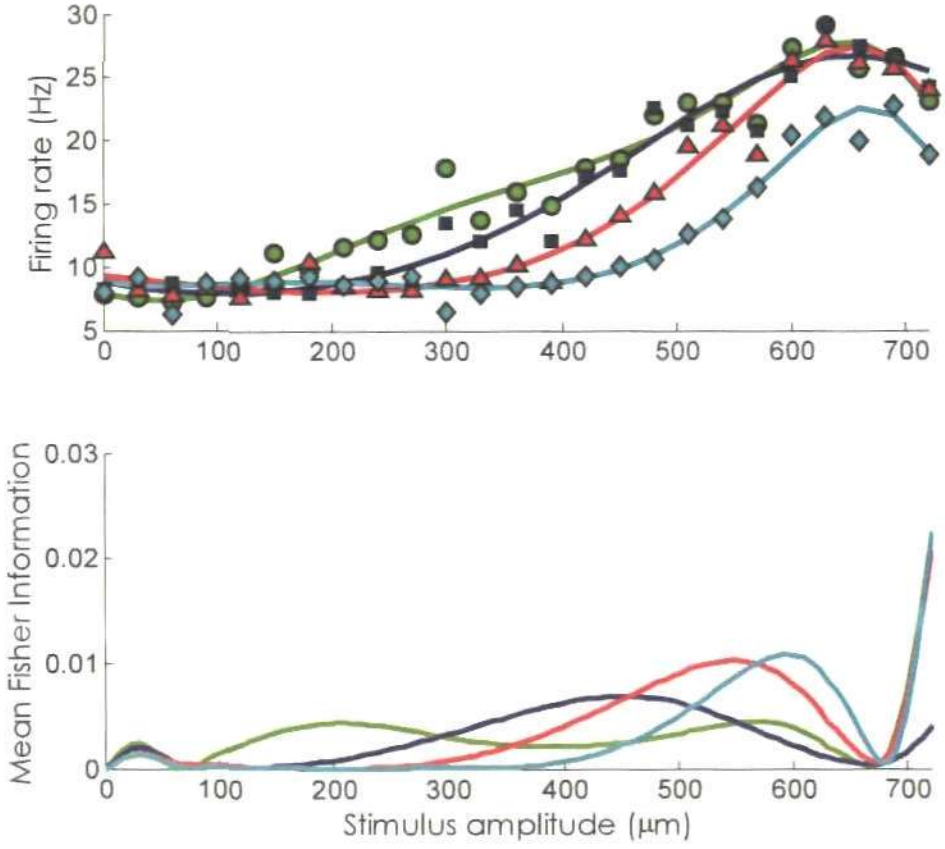

Figure 3.7: Adaptive population stimulus response function and Fisher function from animal LE270705 (number of cells in population = 81)
Top: Adaptive stimulus response function under four different high-probability region levels (colour-coding and amplitude at high-probability region centre is given in inset); symbols are actual, mean spike-counts whilst smooth lines give the best fit to the data.
Bottom: Pooled Fisher Information calculated from the best-fit to the data in the top plot

The population rate-function was taken as the average of the individual stimulus response function, pooled across the whole population of cells tested [n = 132].

Following adaptation to stimuli of increasing variance, the population stimulus response function decreased in gain (see Figure 3.8A). Under the high-variance condition (Figure 3.8A, black line), both the maximum response and the gradient of the stimulus

response function decreased, whilst the amplitude at the S_{50} increased, relative to the low-variance stimulus response function (Figure 3.8A, red line).

In terms of the population Fisher information, there was little difference in bandwidth between the mid- and high-variance conditions, and only slight increases in amplitude for the mid-variance condition (Figure 3.8B). Both stimulus conditions yielded Fisher functions that could be described as binomial, with the strongest peak located at a stimulus amplitude of 570μm and 600μm (for the mid- and high- variance conditions, respectively), and a second smaller, broader peak between 270-330μm.

Under the low-variance condition, the Fisher function had a single peak located at 330μm and a narrower bandwidth, relative to the two lower variance conditions.

The relative broadening of the Fisher function was not unexpected given the difference in slope between the population rate functions under the three conditions (Fisher information is strongest when the stimulus response function slope is steep and spike-count variability is low).

All conditions elicited Fisher functions with peaks near the stimulus mean (272, 279 and 290μm, low-, mid- and high- variance, respectively) and the binomial nature and broad bandwidth of the Fisher functions under the higher variance conditions implies that the population response was able to effectively encode the wider stimulus distributions.

It was found that the majority of cells in the population displayed an adaptive decrease in stimulus response function gain. Overall, following adaptation to stimuli of high-variance, a decrease in stimulus response function S_{50} gradient was observed for 93.18% of adapted neurons, relative to the low- variance condition (control). Under the high-variance condition, an increase in stimulus amplitude at the S_{50} was observed in

87.12% of cells, relative to control; 80.30% displayed a concomitant decrease in slope and increase in S_{50} amplitude.

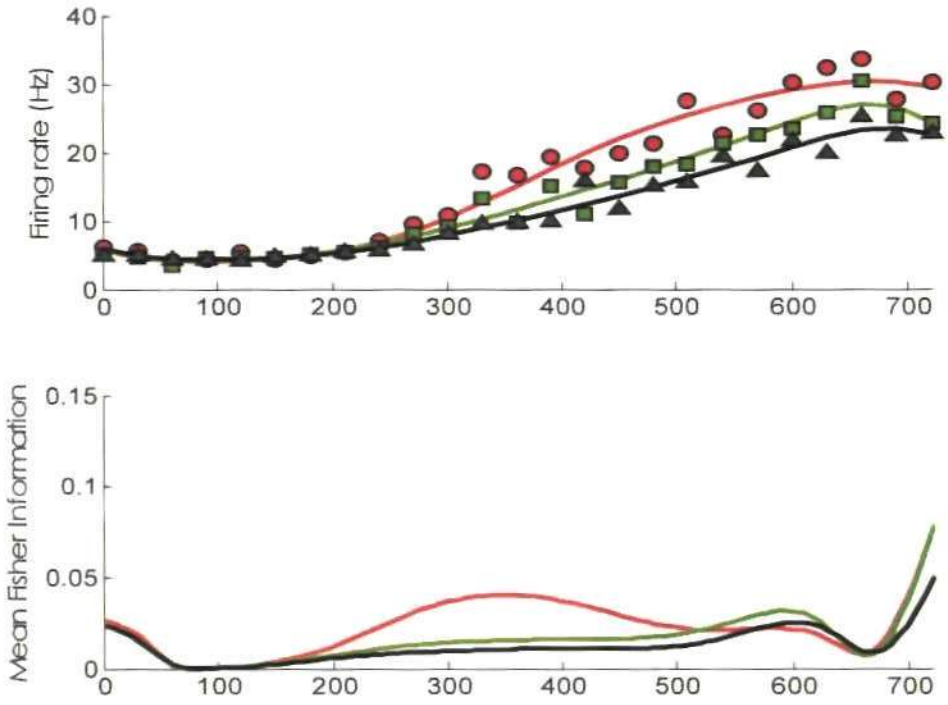

Figure 3.8: The population response to increases in variance within the stimulus distribution.
Top: the mean, pooled rate functions for all cells tested ($n = 132$) under the three variance regimes (low -- red line/circles, medium -- green line/squares, high -- black line/triangles), the smooth curves are the best-fits to the data. **Bottom**: The population Fisher Information, calculated from the population stimulus response function in the uppermost plot

The mean decrease in slope between the low and high variance conditions was 0.0515 ± 0.004 standard error (SE) and highly significant, at the 5% level, with $p < 0.001$ using a one-tailed paired t-test (MatLab function *ttest2*). The average increase in amplitudes eliciting the S_{50} was 51.73μm $\pm$ 2.25SE and was also highly significant, at the 5% level (one-tailed t-test, $p < 0.001$).

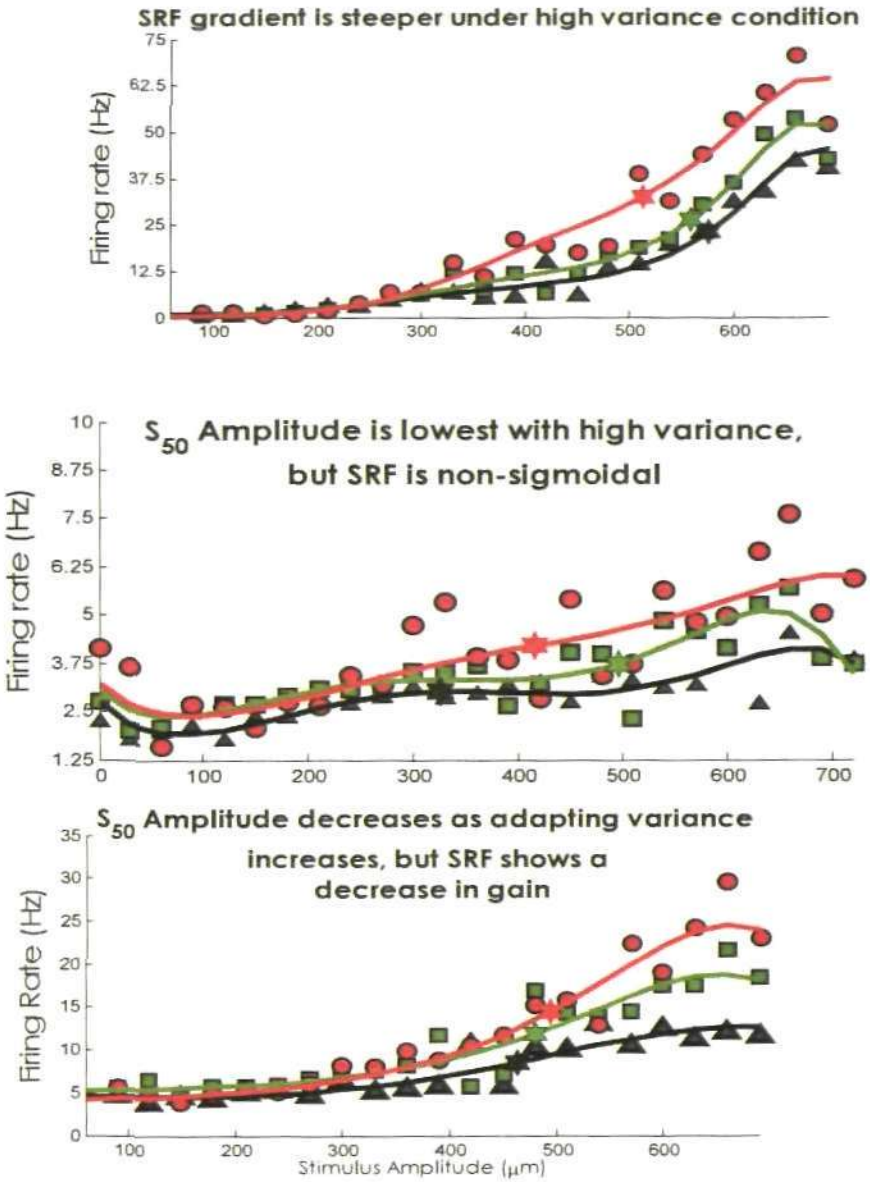

Figure 3.9: Example stimulus-response functions (stimulus response function) from three cells that displayed either an adaptive increase in gain (stimulus response function gradient) or decrease in amplitude at the S_{50}
Adaptation regime: Low-variance: red line/circles; Mid-variance: green line/squares; High-variance: black line/triangles. Smooth lines plot best-fit to the actual spike counts; stars indicate S_{50} point, all other symbols plot actual spike counts.

Overall, the results, from the re-analysis of the experimental data, indicate that the

population response was in line with Information Theory, in that increased stimulus variance resulted in an adaptive **decrease** in gain across the population. However, what was surprising was the degree of agreement with current theory given the results published in Garcia-Lazaro et al. (2007). If one examines the relationship between S_{50} slope and stimulus variance in the published data (Figure 3.5D), the overall trend was for neural gain (as measured by the steepness of stimulus response function slope at the S_{50}) to **increase** with levels of adapting variance.

In fact, during this analysis, no cells were found that displayed an adaptive increase in stimulus response function gain (as in Figure 3.5A).

Due to the differences in the results presented here and the published findings, it was decided to re-examine the relationship between the individual stimulus response function S_{50} amplitudes and slopes for all cells in the population and compare the results from the re-analysis with the published data of Garcia-Lazaro et al. (2007). Scatter plots are presented in Figure 3.10.

In comparing the scatter plots in the left column (re-analysis) of Figure 3.10 with the right column (published data) it is immediately apparent that the results from this re-analysis, in terms of adaptation to variance only (Figure 3.10A-D), are exactly opposite to those of the published data; whilst the results relating to adaptation to the mean are in agreement with the published findings (Figure 3.10E-F).

Following an in-depth analysis of the methods used, by the authors of the published data, to read the experimental spike-count data it was concluded that the published data was in fact presented back-to-front, i.e. the high-variance data was mistakenly presented as the low-variance data and *vice-versa*.

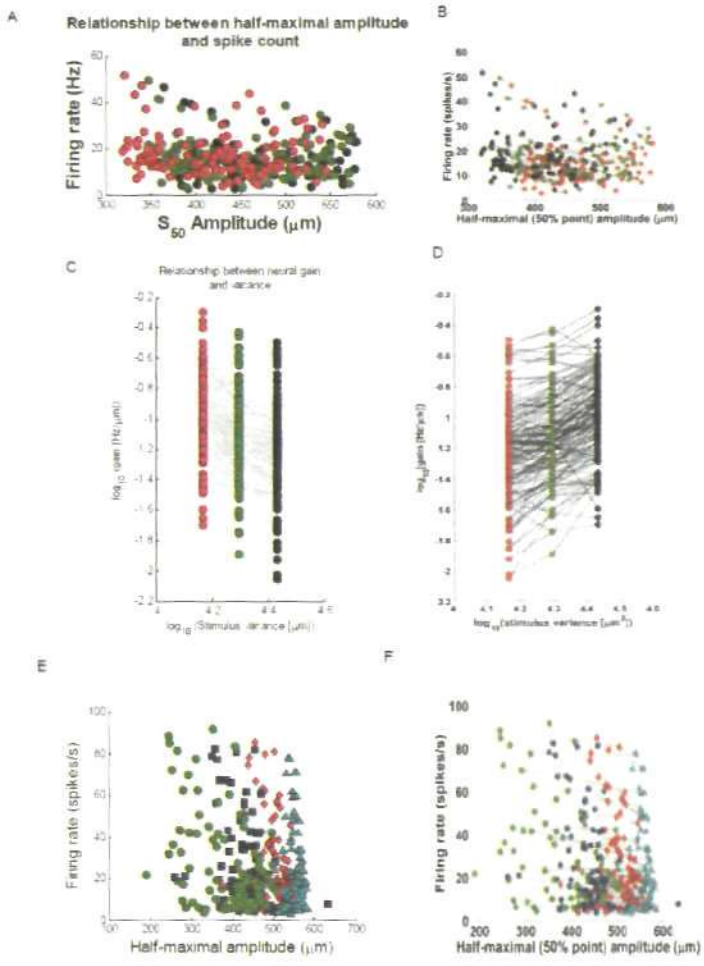

Figure 3.10: Comparison of results from this re-analysis of the variance and mean adaptation data and the published results in Garcia-Lazaro et al. (2007)
All plots in the left hand-side column (A, C & E) result from the re-analysis of the experimental data performed here; all plots in the right hand-side column (B, D & F) are taken from the published data (Garcia-Lazaro et al. 2007). **Top row:** Comparison of relationship between S_{50} amplitudes and firing rates under conditions of adaptation to variance; red circles = Low-, green = Mid-, black = High-variance; black (B)/ grey (A) lines connect the S_{50} points for an individual cell. **Middle row**: Comparison of neural gain (stimulus response function gradient at S_{50}) as a function of stimulus variance; colour coding as for the top row. **Bottom row**: Comparison of relationship between S_{50} amplitudes and firing rates under conditions of adaptation to stimulus mean; green circles = high-probability region centred on 90μm, blue = 210μm, red = 330μm,cyan = 450μm; black (F)/ grey (E) lines connect the S_{50} points for an individual cell.

This has been confirmed by the authors of the published data (personal communication, Jan Schnupp, 2008). As such the results presented here reflect the actual effects of adaptation to increasing levels of stimulus variance in the rodent barrel cortex

3.3: Research looking at time course of adaptation

Adaptation can occur on many time scales, e.g. there can be a rapid reorganisation of neural output (100s of ms scale) following stimulus onset, suppression of firing rate during adaptation (time scale of stimulation - 100s of ms to seconds) and adaptive responses can be observed some time (order of seconds to minutes) post adapting stimulus offset (e.g. Dean et al. 2008; Ulanovsky, Las, Farkas & Nelken, 2004).

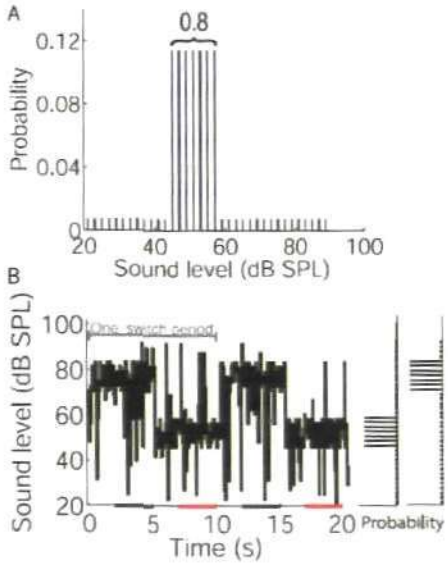

Figure 3.11: An example of the switching stimulus used in to assess the time course of adaptation to global mean in Dean et al. (2008)
A: The distribution of the adapting stimuli with high-probability region (high-probability region) centred at 51 dB SPL. **B:** Example of a stimulus sequence over two 10 s switch periods, drawn from distributions shown to the right (high-probability region centred on 51 and 75dB). The high-probability region switched every 5 s between the 75 and 51 dB stimuli. The bars on abscissa indicate sampling periods (repeated throughout stimulus) that were used for plotting rate level functions from the final 3 s of each high-probability region condition

In a more recent study (Dean et al. 2008) the time scales inherent in adaptation to stimulus statistics were explored. Adaptation to noise burst stimuli, in the inferior colliculus, was elicited through a switching stimulus in which the overall mean sound level was switched from 51 ± 6 db to 75 ± 6 db every 5s (over a period of 10mins, see Figure 3.11 for an example of one transition between global mean levels); the amplitude of the noise burst was refreshed every 50ms and drawn from a highly skewed distribution comprising a high probability region (high-probability region) and two low probability tails.

The study found that the median time constant of adaptation (across the population) was ≈ 160ms, following an increase in the mean sound level, and ≈ 330ms, following a reduction in mean sound level. Reorganisation of stimulus-response functions (stimulus response function) was apparent in the 1st 300ms following the switch to 75db stimulus and for some neurons the stimulus response function reached its steady state (averaged from last 3s of each 5s section) within this time window, see Figure 3.12.

Some neurons also showed a slow, long form of adaptation that had a time constant of ≈ 17s and suppressed firing rates during the 10mins of recording time.

3.3.1: Time course of adaptation in the barrel cortex

Within the study, highlighted above, there was no attempt to measure the time course of adaptation to variance, given that very few cells displayed variance adaptation in the IC. Thus it was decided to attempt to analyse the barrel cortex data to compare and contrast the adaptive time constants under the different conditions, of adaptation to mean and variance, as well as to determine if there are differences across modality and processing stage. The firing rate as function of stimulus amplitude was averaged over the 1st and 2nd 400ms of the stimulus presentation and compared with the steady-state response, averaged

over the final 10s. In general, across the population, no differences in responsiveness could be found, between the first 400ms epoch and any successive time bins, under either adaptation to global-mean or adaptation to global-variance.

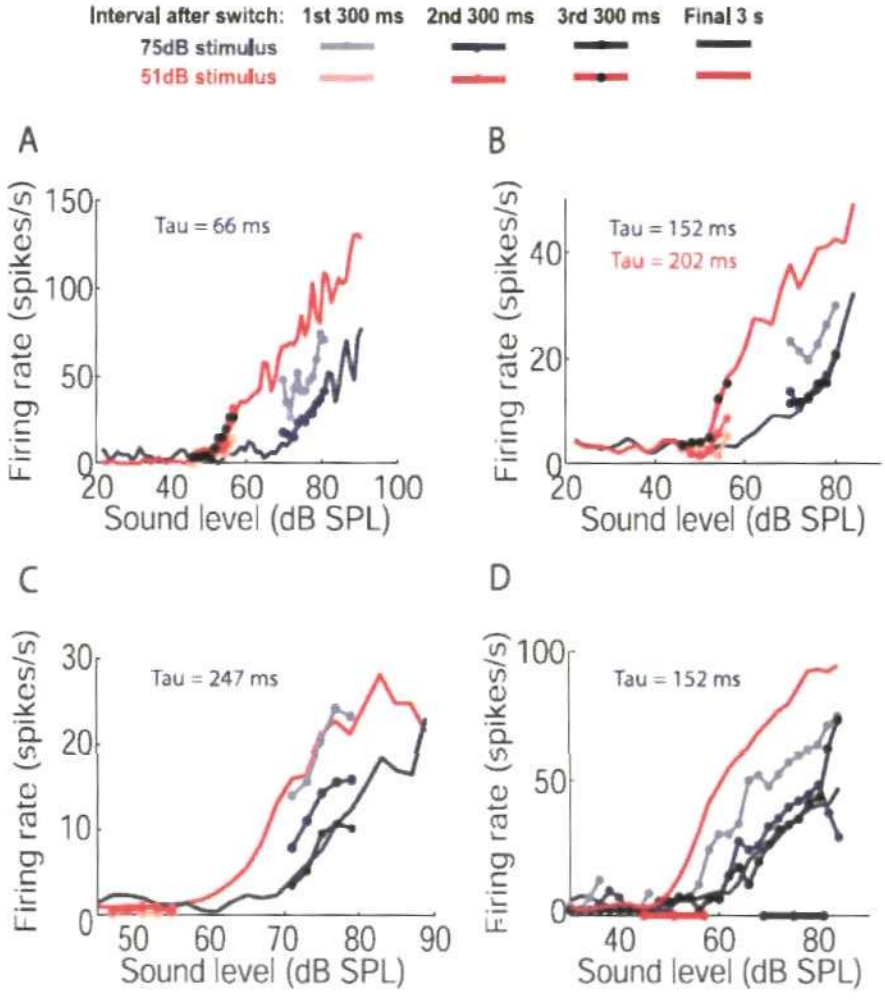

Figure 3.12: Time course of adaptation, during transition between high-probability region levels

Time course of adaptation to mean decibel level of noise bursts, for four representative neurons in IC. The stimulus response functions for amplitudes within the high-probability region are plotted as a function of the first 3, 300ms time bins, following the switch between mean adapting levels, and the final 3s. Different shades of blue represent the response to the 75dB stimulus and red plots the response to the 51dB level. Taken from Dean et al (2008)

The results from a sample single cell in the population (adaptation to mean) are plotted in Figure 3.13; only the responses to high-probability stimuli were analysed. The black dashed line represents the average response (across and within trials) during the last 10s of the each 30s adaptation trial, for a given high-probability region level. The red line gives the averaged response during the 1st 400ms of adaptation and the blue line plots the response for the next 400ms epoch. As can be seen, there is very little difference in the responses during all three time bins.

Thus one could argue that adaptation is happening rapidly (with a time constant of <400ms) as per the inferior colliculus. However, due to the constraints of the data this cannot be stated unequivocally. In terms of measuring the adaptive time course, the switching stimulus (as applied in the inferior colliculus, Dean et al. 2008) has two advantages over the method used to collect the barrel cortex data. Firstly there are only two global mean adapting levels and many repeats, therefore it's possible to get a statistically valid average of adaptation for each 5s section. Secondly, there were no gaps between switches, thus making it possible to compare rapid changes in responsiveness following a mean level transition, i.e. 300ms before and after the switch.

In terms of the barrel cortex data, it was collected in the following manner: responses to a given high-probability region (high-probability region) was recorded for 30s with a 2s break in stimulation between trials. For each high-probability region, 3 random sequences were generated and each sequence was run at least 5 times. Therefore for each high-probability region, only a minimum of 15 presentation times, per neuron, could be used to determine the time constant of adaptation onset (compared to 60 times in Dean et al. 2008); also there could be no comparison between previous firing rates and those

following stimulus onset. Unfortunately then, due to the constraints of the data, it was impossible to determine the time course of the development of the adaptive response with any certainty.

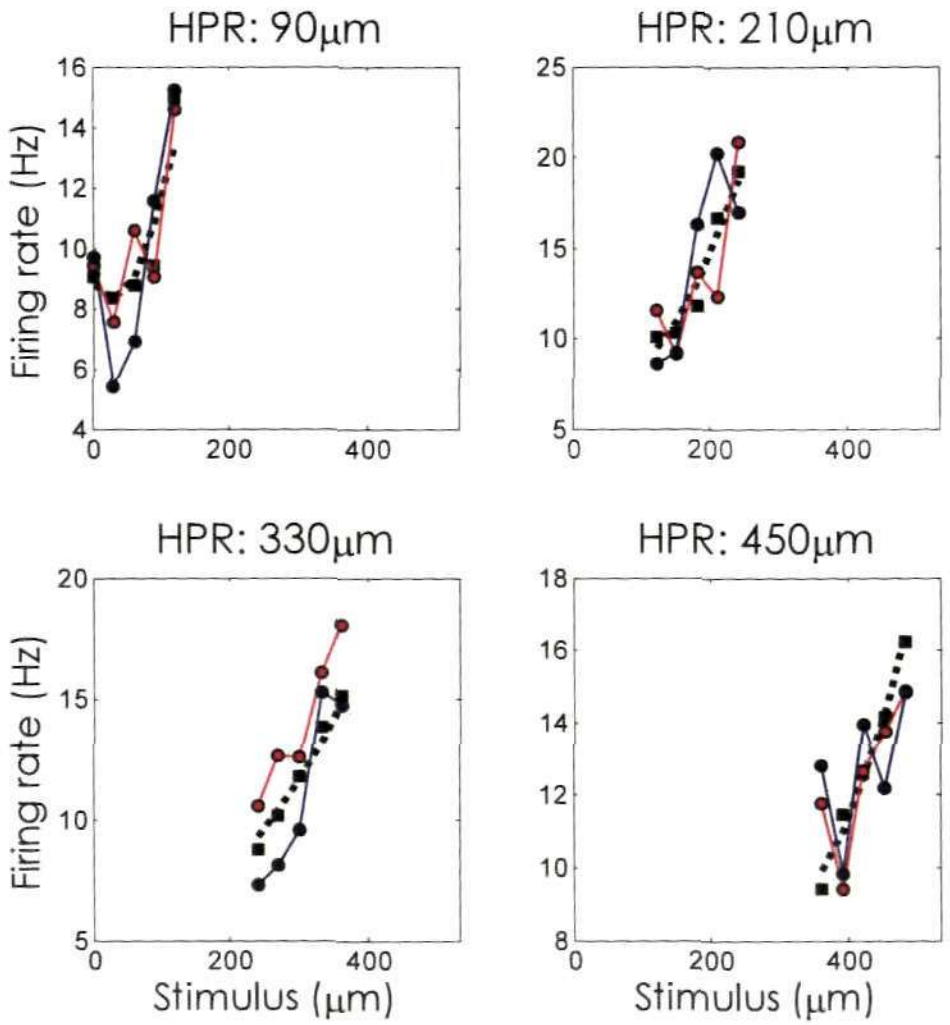

Figure 3.13: The time course of adaptation.
Example of time-course analysis from a single, example neuron for all four adapting conditions. Black dashed line/squares: the average steady-state response of the neuron to high-probability stimuli. Red line/circles: The response averaged over 1st 400ms of stimulus presentation. Blue lines/circles: The response averaged over the 2nd 400ms of stimulus presentation.

As trials testing adaptation to each high-probability region were not presented successively, it was not possible to determine the time course of adaptation over any longer time scales. Therefore to elucidate the actual time course of adaptation to stimulus statistics in the barrel cortex will require further experimentation.

3.4: Discussion

The main findings of the re-analysis of the experimental data from Garcia-Lazaro et al. (2007) were that:

1. The mean neural-response function, averaged over the whole population of cells tested, displayed a decrease in gain, i.e. was scaled downwards, following adaptation to increasing levels of stimulus variance.

2. These findings are in line with the expectations of Barlow's Efficient Coding Hypothesis.

3. In terms of Fisher information, the bandwidth of the population Fisher-function, under the low-variance condition, was unimodal, narrowly tuned and maximal at or close-to the mean adapting stimulus amplitude.

4. Under the mid/high-variance adapting conditions, the shape of the Fisher Information functions were very similar; coding efficacy was broadly tuned and bimodal.

5. In terms of adaptation to stimulus variance, the findings presented within this Chapter directly contradicted the published data. This arose due to issues in the published analysis of the data that assigned the output from the low-variance experiments to the high variance condition

6. Following personal communication with the authors of the published data the consensus was reached that the results reported in this chapter reflect the actual

experimental findings.

These results present further confirmation of Barlow's theory of sensory coding, in that adaptation (in rat barrel cortex and at the population level, at least) serves to adjust the sensitivity of a cell, or even a population of cells, in a manner that is dependent upon the prevailing stimulus levels. This re-analysis of the experimental data revealed that independently increasing stimulus variance led to a decrease in gain of the population response and that this was contrary to the published reporting of the experimental data.

It was previously established that, under conditions where there is adaptation to mean amplitude of whisker deflection, sensitivity is greatest for stimuli outside the high-probability region Garcia-Lazaro et al. (2007). Dean et al. (2005) suggested that the slight displacement of neural rate-functions, from their expected location, meant that adapted cells were primed for novelty detection. Within somatosensory cortex, stimulus response function displacements were more extreme than observed in the IC (Garcia-Lazaro et al. 2007), implying that cells here were primed to respond to stimuli far outside the expected range of stimulus amplitudes. As such, the next phase of the analysis was to explore other possible coding strategies and to see which best describes the adaptive response to global mean.

In conclusion, adaptation to variance results in a decrease in gain and not an increase, as previously published (Garcia-Lazaro et al. 2007). Overall, this suggests that adaptation to stimulus statistics, within the barrel cortex, is in line with the principles of the Efficient Coding Hypothesis. However, adaptation to variance within the Inferior colliculus (IC) was of course rarer and in the opposite direction (Dean et al. 2005 and see Figure 3.2). This raises the question of whether the opposite effects of adaptation on neural output

between the IC and the barrel cortex is a reflection of the differences in adaptation to variance between sub-cortical and cortical regions generally or is specific to sensory modality, or even species. Another question to raise also is whether the experimental data from the IC (Dean et al. 2005) was subject to the same issues of analysis as the data from the barrel cortex (Garcia-Lazaro et al. 2007); in which case adaptation to variance in the IC would result in decrease in gain, as per the expectations of the Efficient Coding Hypothesis. Further experimentation is required to answer these questions.

The displacement of the neural-rate function, within the barrel cortex, places the sensitivity of the cell in a region of the stimulus space that is higher than the global mean of the input levels. It could be argued that this adaptive shift reduces firing rates overall and is an energy-saving strategy (as the most common stimuli elicit low firing rates), however, it also has the added affect of reducing the cell's ability to discriminate (with a rate-code) between those stimuli which are presented most frequently. Thus it was decided to analyse the data further in order to asses the optimal encoding strategy underpinning adaptation to stimulus statistics; the results are presented in the next chapter and suggest that the lateral displacement of the stimulus response function following adaptation to the global mean input optimizes the encoding of relative changes in stimulus amplitude on the 40ms time scale, rather than the stimulus distribution.

Chapter 4: Adaptive strategies

One of the aims of the last chapter was to ascertain whether adaptation to stimulus statistics in the barrel cortex was in line with the predictions of information theory. Whilst some of the effects of adaptation on neural output were within expectations (specifically the mean-amplitude-dependent lateral displacement of the stimulus-response functions, see Figure 4.1A below, and the decrease in neural-gain as a result of increased variance in the adapting stimulus) there was one feature of the adaptive response that did not fit the predicted neural-output: the lateral displacement of the stimulus-response function tended to place its dynamic region outside the high-probability stimulus space, thus resulting in Fisher information functions that peaked outside the high-probability region (see Figure 4.1B, below). Both the Infomax principle and the efficient coding hypothesis would predict that information transfer should be maximised around the high-probability region (see Figure 4.1A, below). The aim of this chapter, then, is to explore different possible coding strategies that may underpin the experimental data.

4.1: Is information maximisation the strategy underpinning adaptation in the barrel cortex?

As one can recall from Chapter 2, the adaptive responses to velocity and acceleration, in the fly H1 neuron, were scaled when normalised by the standard deviation of the adapting stimulus (Brenner et al. 2000).

However, the normalised responses, plotted in Figure 4.1, are clearly not scaled versions of each other.

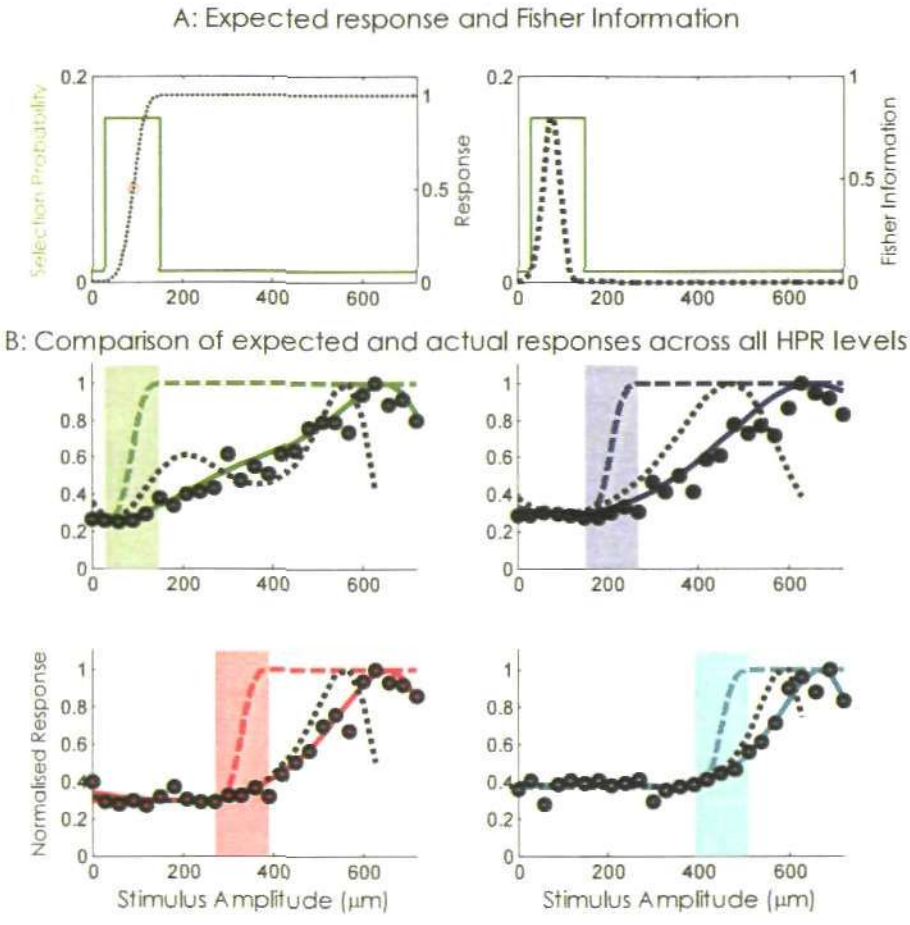

Figure 4.1: Expected versus actual adaptive response.
A: Predicted output according to the Information maximisation principle. Solid green lines plot the stimulus probability distribution and dotted lines describe the expected stimulus-response function (SRF, left plot) and Fisher information (right plot, calculated as: $f'(s)^2 / \sigma^2$, where $f(s)$ is the predicted response, see left-hand plot, as a function of the stimulus, s, and under the assumption that the variance, σ^2, is equal to $f(s)$; see Chapter 2, Equation: 2.28). **B**: Comparison of actual versus predicted response; shaded areas cover the high-probability regions (HPR) in the stimulus space: Green: HPR centred on 90μm, blue: 210μm, red: 310μm, cyan: 450μm. Correspondingly coloured dashed lines plot the predicted SRF; the corresponding solid lines and circles plot the best-fit and **actual** average, population stimulus-response function, respectively. Black dashed-lines give the **actual**, population, mean Fisher function; both the population SRF and Fisher Information curves were normalised to give a maximum response of unity.

In fact, normalising the data by any method (dividing by either the mean or standard

deviation of the stimulus or normalising the data to have zero mean and a standard deviation of unity) did not result in a scaling-up of the normalised stimulus-response functions such that the responses were invariant with respect to the global mean of the adapting stimulus (data not shown, but please refer to Figure 4.1). There is therefore an argument to say that the degree of lateral displacement was not only influenced by the global mean of the adapting stimulus but by other aspects of the stimulus, such as rarity or deviation.

Within Chapter 2, the predicted effects on neural output of an information encoding strategy based on the principles of background suppression, , was discussed for a multi-modal stimulus distribution (Wimmer et al. 2008). Under a mechanism of background suppression, whereby responses to the weakest intensity inputs are suppressed, an adapted stimulus-response function would only cover the highest-intensity mode of a multi-modal stimulus distribution (see Figure 4.2, right plot, below), rather than covering the whole distribution (as per adaptation along information theoretical principles, see Figure 4.2, left plot, below). This is similar to the actual adaptive response, within the barrel cortex, observed in response to increases in global mean whisker deflection amplitude; the stimulus-response function is displaced to centre on stimuli that are greater in intensity than the global mean stimulus amplitude (see Figure 4.1B, above).

However, in the example of background suppression just discussed (see Figure 4.2), the stimulus-response function was displaced to centre on the stimuli of greatest intensity but also of high presentation-probability. In the case of adaptation in the barrel cortex, the stimulus-response functions were displaced to centre on the strongest but low-probability inputs; this suggests that barrel cortex neurons were adapting to respond selectively to rare,

but high amplitude stimuli (see also Figure 4.3, bottom row).

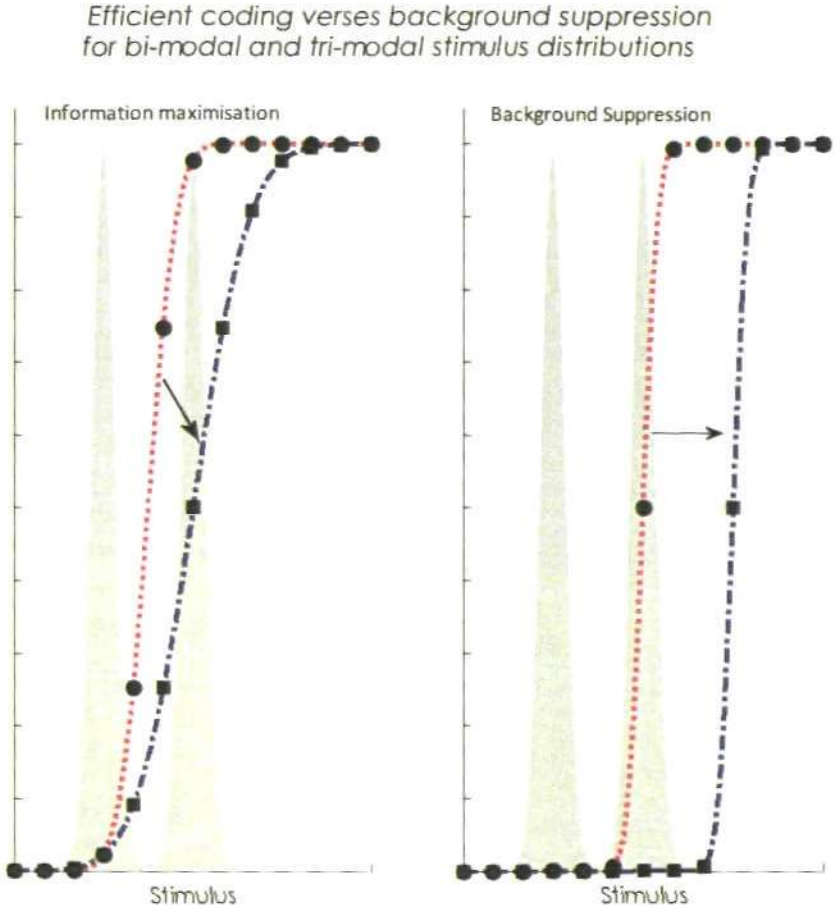

Figure 4.2: Information Maximisation versus Background Suppression as encoding strategies
Dark- and light grey regions represent adapting stimuli with bi- and tri-modal distributions, respectively. **Left:** Expected influence of adaptation, mediated by an Information Maximisation strategy, on stimulus-response curves of a hypothetical neuron. Red line and circles gives expected response following adaptation to an arbitrary stimulus with a bi-modal distribution; blue line and squares give expected adaptive response to a tri-modal distribution. **Right:** As for A, but for a strategy of Background Suppression. Left and right plots both adapted from Wimmer et al. (2008).

Novelty detection, whereby adaptation serves to ready a cell for deviant stimuli, is another possible strategy for the encoding of sensory input. This type of adaptive process would result in higher firing rates for novel stimuli; as, for example, the study described in

Chapter 2 which found that cells in the primary auditory cortex fired most strongly to rare (or oddball) stimuli (Ulanovsky et al. 2003), irrespective of the relative relationship in value between the common and rare inputs and their presentation probability (see Chapter 2 and below).

An optimal adaptive response for a sensory cell primed not only to detect but also to discriminate between novel stimuli, would be for the regions of greatest slope, in the neural-response function, to be centred on the tails, rather than the peak, of the stimulus distribution (see Figure 4.3, middle row, below). This would ensure that the discrimination threshold, as determined by the inverse of the Fisher Information, would be lowest for less common stimuli, irrespective of intensity. There is an argument to say that the tails of the stimulus distribution is where the Fisher Information of the input signal itself would be strongest; thus, there would be a direct correspondence between the strength of Fisher Information in both the input distribution and the response to it.

With reference to the adaptive response to global-mean whisker-deflection amplitude, one could argue that the slope of the stimulus-response function tends to begin accelerating in a linear manner (see Figure 4.1B), on the edge of the high-probability region of the stimulus. Thus, the barrel cortex cells could be adapting to reflect the region of highest Fisher Information within the stimulus distribution. However, it should be noted that the slope of the rate-function was centred only on stimuli greater than global-mean and both neural threshold and peak Fisher Information, as a function of neural response, were displaced to regions of low stimulus-information. Therefore a different strategy, e.g. background suppression or deviation detection, may be underpinning adaptation to intensity, in the barrel-cortex.

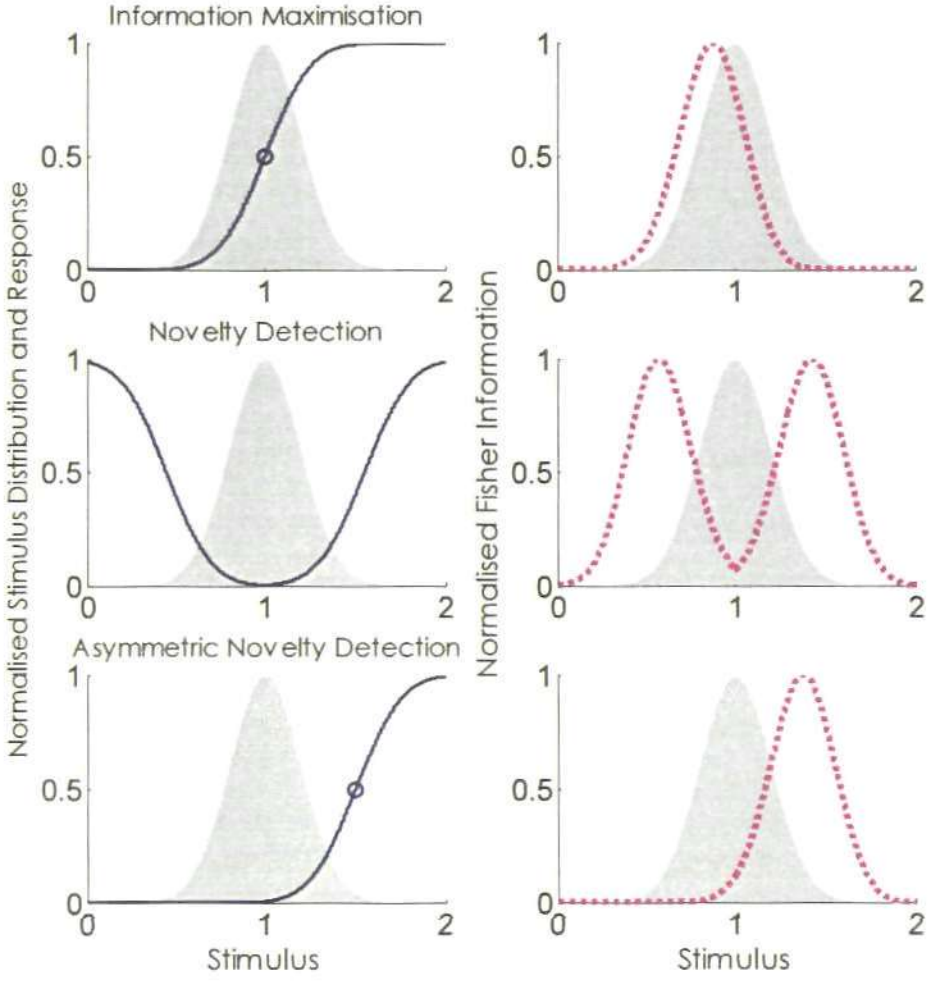

Figure 4.3: Novelty detection as an encoding strategy
Left Column: Idealised stimulus-response functions under three encoding strategies: Information Maximisation (**top row**), Novelty detection (**Middle row**) and Asymmetric Novelty detection (or a form of background suppression, **bottom row**); the stimulus-distribution is shaded in grey, stimulus-response functions are plotted in blue solid lines and the empty circle is the half-maximal point. **Right Column**: Idealised levels of Fisher Information from the stimulus-response functions plotted in the corresponding row of the left column; stimulus-distribution is shaded in grey, Fisher Information functions plotted with magenta dotted lines.

As one can recall from Chapter 2, the study comparing the possible underlying strategies for the encoding of information in the cricket auditory system (Wimmer et al.

2008) could not conclusively state whether information maximisation or background suppression was the underlying strategy. It would appear that the same can be said about the adaptation data from the barrel cortex (as described in Chapter 3 and highlighted in Figure 4.1): Information maximisation cannot fully account for adaptation to stimulus statistics in the barrel cortex. Thus a pertinent question to ask at this point is can background suppression or novelty detection (or a combination thereof) be the strategy or mechanism underpinning adaptation to stimulus statistics in the barrel cortex?

The experimental set-up in the barrel cortex data certainly contained, what could be considered, novel stimuli. One could view all amplitudes outside the high-probability region to be oddball stimuli as they were only presented on average 7 times, within each trial, whereas amplitudes within the high-probability region were presented approximately 100 times.

If improving novelty detection was the driving force behind adaptation one might have expected to see responses increase in both low-probability tails of the stimulus space. Overall, across the population, a monotonic increase in the stimulus-response function was only observed for amplitudes within and greater than the high-probability region. Thus it would seem at first glance that novelty detection is not the primary function of adaptation to stimulus statistics. However, individual cells may have displayed novelty detecting traits, following adaptation, that were lost in the population analysis.

In order to explore this further, the data was re-analysed (see Section 4.2: Novelty detection, below) using a methodology inspired by the experimental paradigm of the study that observed novelty detection in the auditory cortex (Ulanovsky et al. 2003; described fully in Chapter 2).

4.2: Novelty detection

In order to explore whether adaptation to stimulus statistics primes the barrel cortex for novelty detection, one must first ask whether the experimental data can be analysed in terms of novelty detection.

An earlier chapter (see Chapter 2) described a form of stimulus-specific adaptation (SSA) in the auditory cortex that resulted in auditory cells preferring novel, over common, stimuli (Ulanovsky et al. 2003). In brief, adaptation was achieved by presenting two auditory stimuli of equal amplitude but separate frequency. One frequency was assigned as the standard (and was thus the most commonly presented stimulus) and the other the deviant, or oddball, stimulus. The strength of the adaptive effect was dependent on both the difference in frequency value between the stimulus pairs and on the ratio of presentation. Under all conditions, the oddball stimulus elicited the greater response (please see Figure 4.4 and the relevant Figure in Chapter 2).

In comparing SSA *(Ulanovsky et al. 2003) with adaptation to stimulus statistics* (Dean et al. 2005; Garcia-Lazaro et al. 2007) one must first acknowledge the key differences in experimental paradigm, as the barrel cortex experiments were not designed with novelty detection in mind. For example, in the case of the barrel cortex data, amplitude was varied and frequency fixed and the adapting stimulus did not strictly have an oddball design. In general, an oddball experiment will have one `common' and one `rare' stimulus; in the case of the experiments looking at adaptation to stimulus statistics all 20 amplitudes within the low-probability region could be considered as `rare' (each amplitude is presented on average 7 times) and all 5 amplitudes within the high-probability region could be the `common' stimulus (presented on average 100 times, per trial). Thus, the adapting regime

can be considered as a more general form of the SSA paradigm.

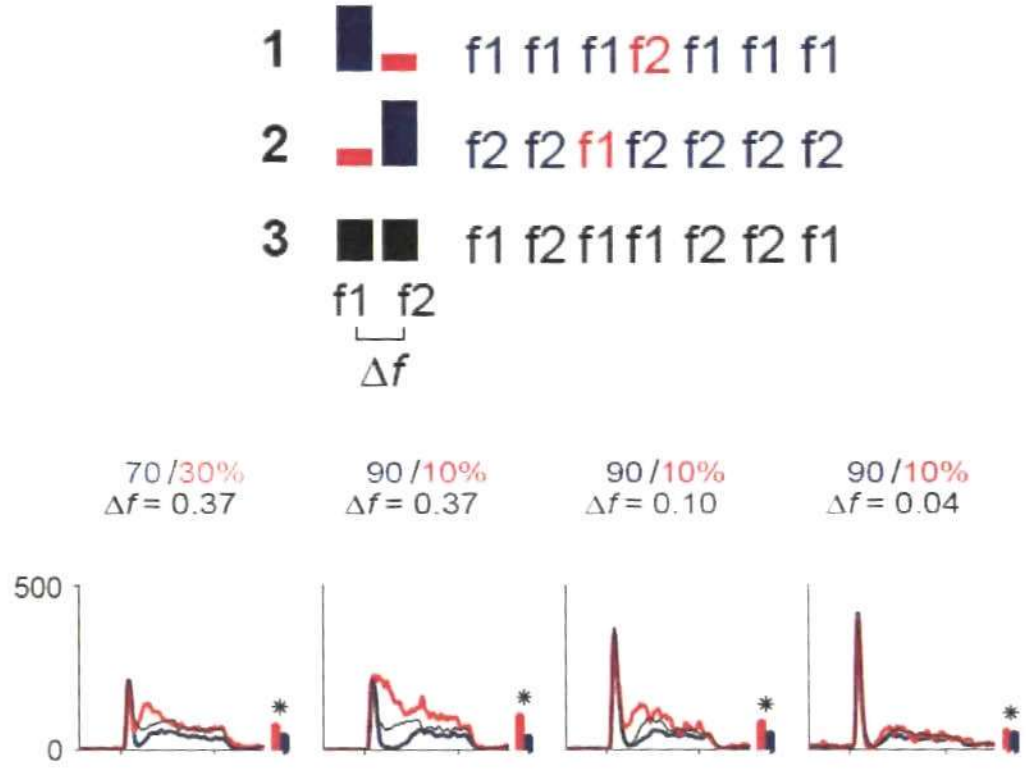

Figure 4.4: Adaptation to novel and standard auditory stimuli.
Top: Schematic illustration of odd-ball experimental paradigm; red indicates the oddball frequency-stimulus and blue the standard input. **Bottom:** The mean response, averaged across *f1* and *f2*, to the deviant (red), standard (blue) and control stimuli (black, 50/50% presentation) for four stimulus configurations (ratio of presentation and difference in frequency between the stimulus pairs (Δf) given above each plot); bars indicate spike counts and asterisk indicates a significant difference between the firing rate-response to deviant versus standard frequency-stimulus (all plots taken from Ulanovsky et al. (2003)).

In the experiments outlined above, testing adaptation to the stimulus mean in the barrel cortex, there were four high-probability region regions, centred on the amplitudes 90μm, 210μm, 330μm & 450μm. Thus each of the amplitudes within the four high-

probability region could be classed as either common or rare depending on the mean adapting level and the responses could be compared to ascertain whether cells respond most strongly when a stimulus is rare, irrespective of its amplitude value.

Of course the response to amplitude, in the barrel cortex, is a monotonically increasing function and distinct from the Gaussian type response that is typical of frequency preference in the auditory system (e.g. Ulanovsky et al. 2003). Thus, in order to facilitate comparison between the rare and common response it was decided to restrict the analysis to amplitudes that were within one of the four high-probability region only and, further, to select the amplitudes at the centre of each high-probability region. If one recalls from the previous chapter, responses to high-probability region amplitudes were generally located at the low end of the firing- response spectrum, thus any preference for deviant stimuli may override the general tendency for firing-rates to increase with stimulus intensity.

The responses of individual cells to the four sets of amplitude pairs were compared, across trials where one stimulus in the pair was rare (i.e. located in the low-probability region), the other common (i.e. within the high-probability region) and *vice versa*. The stimulus pairs were as follows: 90/210μm, 90/330μm, 90/450μm, 210/330μm, 210/450μm and 330/450μm. Responses to the oddball/common stimuli were averaged over both amplitudes in the stimulus pair to give the overall response as a function of presentation frequency. If a cell's adaptive strategy was based on novelty detection then one would expect deviant amplitudes to elicit significantly stronger firing rates than common stimuli.

Figure 4.5, below, presents a graphical representation of the analysis of 90/210μm pairing from a sample cell. The mean response to both amplitudes, across all instances and

all trails where 90μm was the common stimulus and 210μm the rare, was calculated (top row, Figure 4.5). For this cell, and under this configuration, the strongest response was elicited to the rare amplitude; as the deviant was the higher amplitude in the pair, one would expect the response to be greater. However, in the opposite configuration the strongest response was also elicited by the oddball stimulus, which was the lower value amplitude of the pair (see middle row of Figure 4.5). The responses were then averaged as a function of commonality and rarity (i.e. the responses to the amplitudes 90μm and 210μm, across all trials where they were the common stimulus, were averaged together and *vice versa* for the rare stimuli). Any significant difference between the common and rare responses was tested for by either the non-parametric Wilcoxon ranked sum test (MatLab function: *ranksum*; test for significant difference in median response at 5% significance level), if the distribution of responses was not normal and a paired t-test otherwise (MatLab function: *ttest2*; test for significant difference between mean responses to rare or common stimulus at 5% level). For this cell, the response to the rare stimulus was not significantly greater (*p-value* = 0.861, non-parametric test), for the 90/210μm stimulus pair (see Figure 4.5: bottom row, below)

The common and rare responses for this cell, across all six stimulus pairs are plotted in Figure 4.6. The overall rare response was on average greater than the common response for all stimulus pairs (Figure 4.6, bottom row); however in terms of the individual presentations (see Figure 4.6, top/middle rows), the response to the deviant was always greater for three out of the six stimulus conditions only.

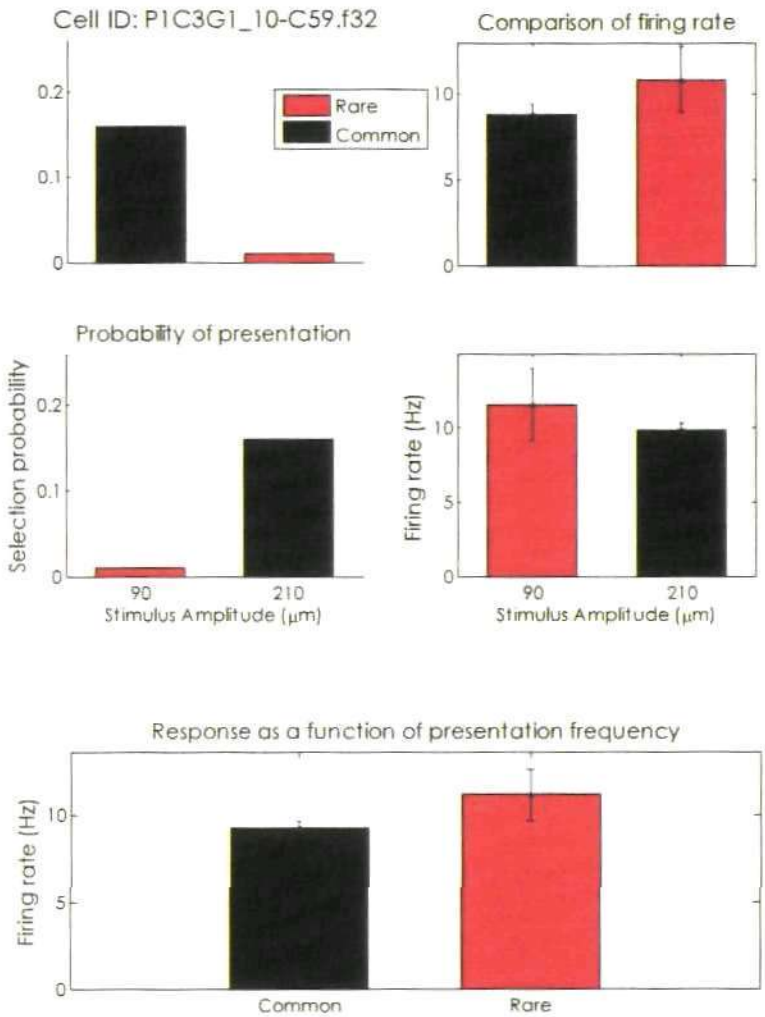

Figure 4.5: Novelty detection for a single stimulus pair (single cell example). **Red**: rare stimulus, **black**: common. **Top and middle row, Left**: Probability of selection for the 90/210μm stimulus pair, when the 90μm stimulus is common (top) and rare (bottom) and *vice versa* for the 210μm stimulus. **Top and middle row, Right**: Firing response of the cell to the two stimulus amplitudes as a function of their relative rarity (as given in the corresponding left-hand plots). **Bottom row**: Firing rate response of the cell, averaged across both stimulus amplitudes within the pair, as a function stimulus presentation levels; the response to the rare (red) stimulus was on average stronger than the response to the common (black) one.

The relative increase in the rare response was significant for two stimulus pairs

only (90/330μm and 210/330μm [indicated by the asterisk above the plots in Figure 4.6, bottom row], non-parametric test, $p = 0.015$ & < 0.001, respectively).

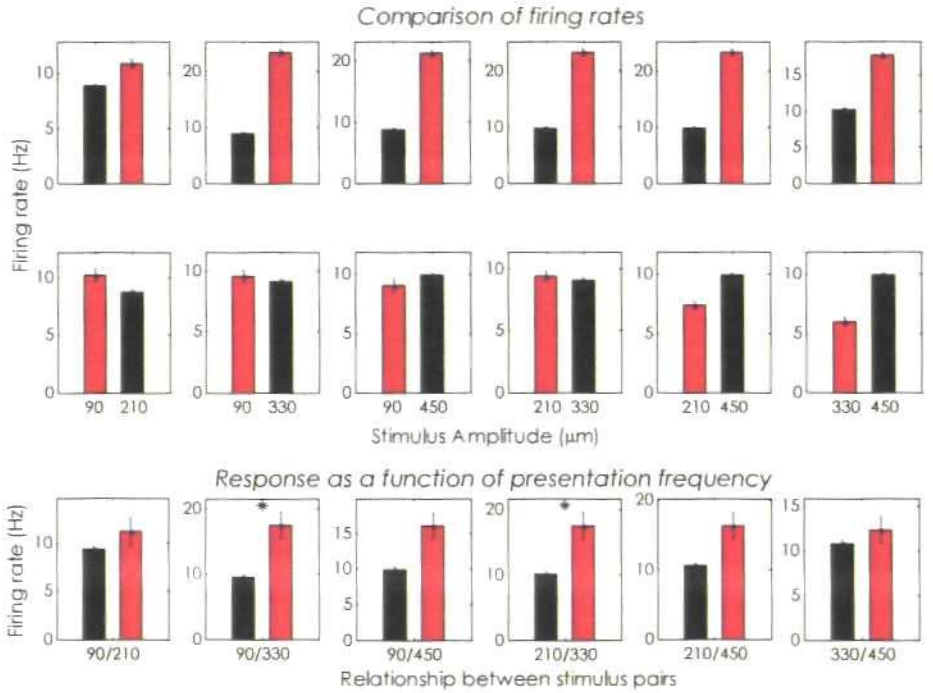

Figure 4.6: Novelty detection for all stimulus pairs (single cell example).
Red = rare stimulus, **black** = common, error bars give the standard error in the average firing rate for a given amplitude. **Top and middle row**: Firing rate response of the cell in Figure 4.5 to all pairs of stimuli, as a function of their relative rarity. For the top row, the weakest amplitude in the pair was always the common stimulus in the top row and *vice versa* for the middle row). **Bottom row**: The average firing rate response of the cell in Figure 4.5, for all stimulus pairs, as a function of stimulus presentation levels, rather than intensity; where the response to the rare (red) stimulus was significantly stronger than the response to the common (black) one is highlighted with an asterisk.

Across the population (Figure 4.7), the response was similar to that of the single cell case, highlighted above (c.f. Figure 4.6). Whilst the overall response to rare stimuli was greater, for all pairs bar the 330/450μm combination, the response was only significant for the 90/210μm and 210/330μm stimulus pairs (non-parametric test, $p < 0.001$, for both cases).

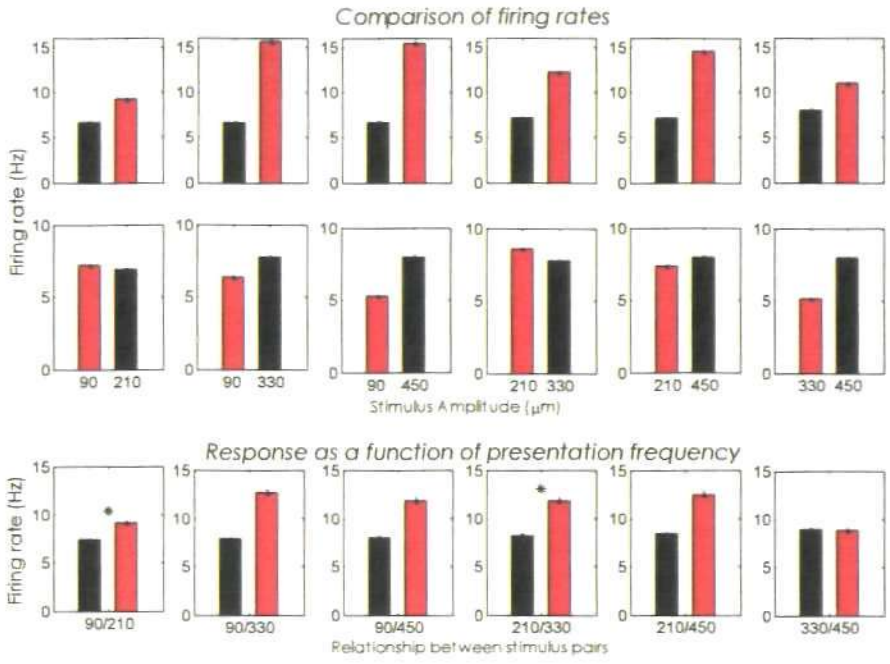

Figure 4.7: Novelty detection for all stimulus pairs, across the population (n=81). **Red** = rare stimulus, **black** = common, error bars give the standard error in the average firing rate for a given amplitude. **Top and middle row**: Firing rate response of the population to all pairs of stimuli, as a function of their relative rarity. For the top row, the weakest amplitude in the pair was always the common stimulus in the top row and *vice versa* for the middle row). **Bottom row**: The average firing rate response of the population, for all stimulus pairs, as a function of stimulus presentation levels, rather than intensity; where the response to the rare (red) stimulus was significantly stronger than the response to the common (black) one is highlighted with an asterisk.

Both the significant pairs were composed of the 210μm stimulus and the amplitude level from either the preceding or successive high-probability region. Thus it could be argued that for amplitude relationships of ±120μm, adaptation acts to elicit a preference for deviant stimuli, irrespective of the relationship between amplitudes. However for this to hold true one would have expected there to be a significant preference for the oddball stimulus in the 330/450μm pairing; instead, the response to the high amplitude (450μm) was always greater. This then could reflect the emergence of the dominant tendency for firing-rates to increase with stimulus amplitude.

It would appear, from the above analysis, there is a degree of novelty detection underpinning adaptation to stimulus statistics in the barrel cortex, but no definite pattern as to its behaviour emerges. A tentative suggestion is that the distance between paired amplitudes can influence the development of novelty detection, but it is easily overridden by the tendency of cells to increase firing rates with increased amplitude, irrespective of presentation frequency.

4.3: Stimulus Context

The results above, relating novelty detection and adaptation, were equivocal, thus it was decided to apply an alternative approach to assessing the underlying encoding strategy; specifically, that of exploring the effect of context on adaptive responses, in terms of the relationship between output and recent stimulus history. Adaptation is not a static process and can act on several time scales ranging from the 10ms to 10s scale; therefore it is useful to consider what's happening on a local, as well as at the global, scale. The dependency of firing rate on stimulus history can be quantified by plotting the spike-triggered average or STA (de Boer & Kuyper 1968); the STA is essentially the average stimulus over time, which elicits a spike.

4.3.1: The spike-triggered average stimulus

The STA can be quantified as:

$$STA = \frac{1}{n}\sum_{i=1}^{T} s_i x_i$$

4.1

Here n is the total number of spikes in the spike train, s is the stimulus vector preceding the

spike in the *ith* time-bin and x is the number of spikes within the *ith* time bin.

In calculating the STA, spikes were sorted into 40ms time-bins in a manner identical to that employed to calculate the firing rate response (see Chapter 3). For each cell in the population ($n = 81$), stimulus vectors (from 0- 320ms [equivalent to 8 time-bins] prior to each time-bin), for every spike in each time-bin, were summed and then averaged across all trials of a given mean adapting level. The population response was the average STA across all cells.

The population STA following adaptation to all four mean amplitude levels is plotted in Figure 4.8 . Overall, across all mean adapting conditions, spikes were elicited when there was a relatively large difference between the stimulating and preceding amplitude (with respect to step-wise differences between the average stimuli at times of < -40ms). With the exception of the 90μm adapting condition, all the amplitudes preceding the stimulating amplitude were within the high-probability region.

The overall trend, for all conditions, was for stimulus amplitudes to decrease from the global stimulus mean to reach a minimum level at ≈120ms prior to the spike, and then increase in small step-sizes until the amplitude preceding the spike (-40ms time-bin) was reached. The average stimulus amplitude that elicited the spike (in Figure 4.8 the average stimulating amplitude was plotted at time > 0) was always greater than the amplitudes located within the current high-probability region level.

A similar trend was also found for the variance data (Figure 4.9), though the steady-state amplitudes (latency of < -240ms) were generally higher than the global mean amplitude of the stimulus and the difference between the stimulating and preceding amplitude was greater than observed under adaptation to stimulus mean.

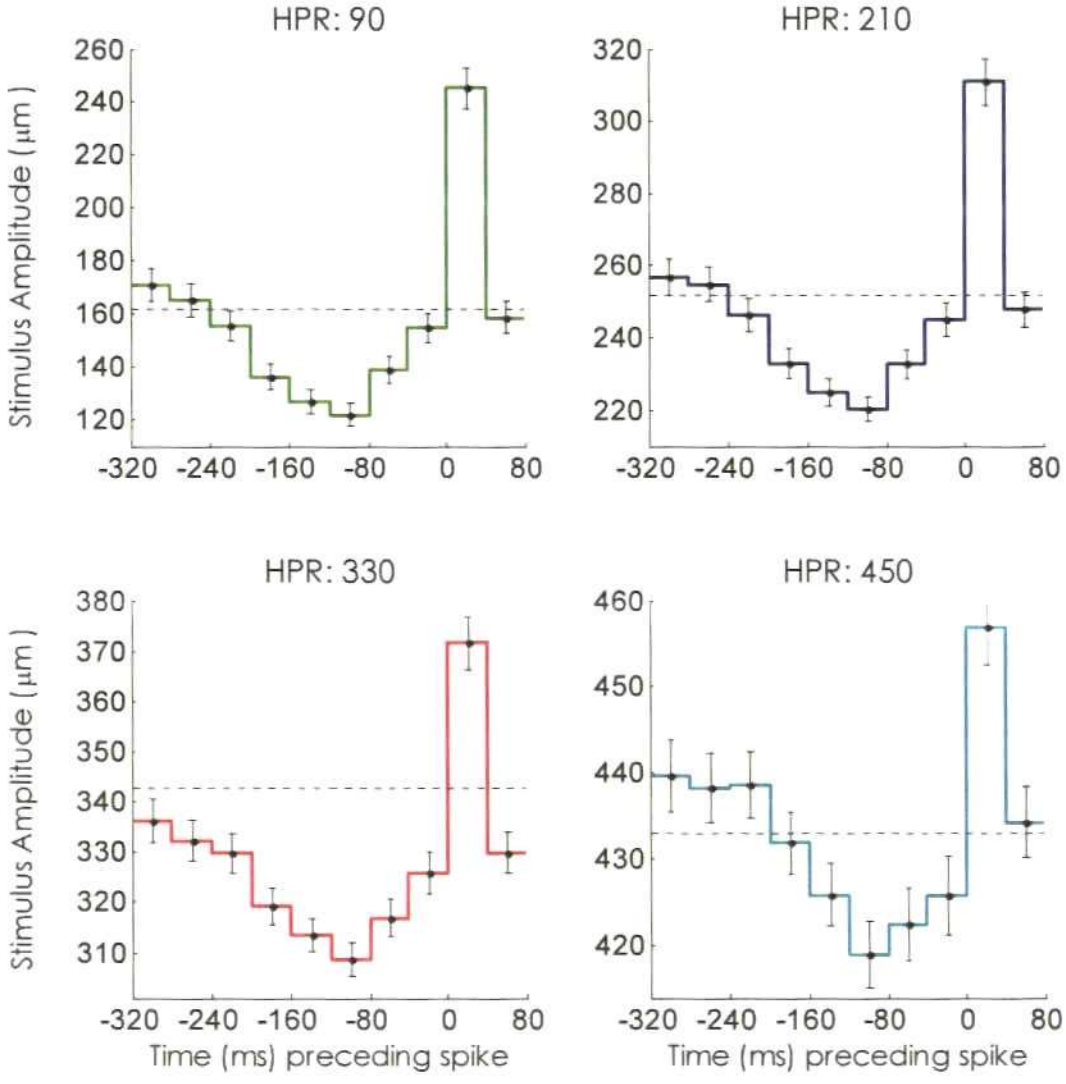

Figure 4.8: Spike-triggered average (STA) during adaptation to global mean, across the population (*n=81*).

Error-bars give the standard error, solid lines plot the STA, dashed line indicates global mean of the stimulus (the central amplitude (in µm) of the high probability region, is given above each plot). The amplitude at time from 0-40ms indicates the average stimulus amplitude that elicited each spike; the input that followed the spike-eliciting stimulus is plotted from 40-80ms.

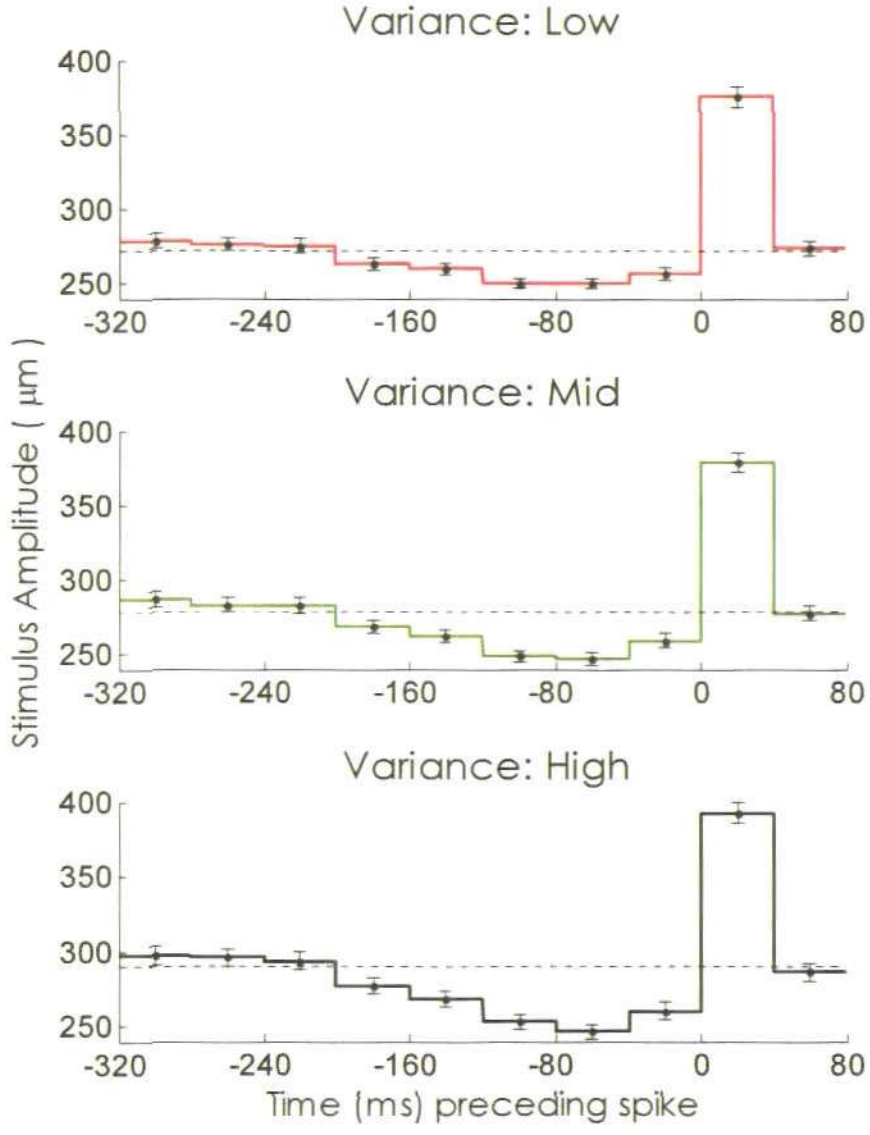

Figure 4.9: Spike-triggered average (STA) during adaptation to global variance, across the population (n=81).
Error-bars give the standard error, solid lines plot the STA, dashed line indicates global mean of the stimulus (the level of adapting variance is given above each plot). The stimulus amplitude at time from 0-40ms indicates the average stimulus amplitude that elicited each spike; the average input that followed the spike-eliciting stimulus is plotted from 40-80ms.

In contrast with other studies looking at the relationship between a stimulus feature

and neural response, within the barrel cortex (e.g. Maravall, Petersen, Fairhall, Arabzadeh & Diamond 2007 [see below and Section 4.5:, this Chapter]; de Kock & Sakmann 2009) the spike-triggered average revealed a clear dependence, of spiking behaviour, on stimulus history. The most interesting trend in the STA was the apparent dependence of spiking output on a general decrease in stimulus amplitudes prior to firing and the large, positive difference in step-size between the preceding and stimulating amplitude.

It was thus decided to explore whether the adaptive response was actually a function of step-size between the current amplitude and the one preceding it, rather than the global stimulus statistics. This question has already been addressed in the Inferior colliculus (Dean et al. 2005), where adaptation to stimulus statistics was not found to be linked to step-size (see Figure 4.10).

In Figure 4.10, the step-size functions of three sample cells, from the inferior colliculus, were plotted. As a brief reminder, these cells were adapted using noise burst amplitudes drawn from a highly skewed distribution in which the overall adapting mean level was determined by the location of the high-probability region in the stimulus space (Dean et al. 2005). Neuron 1 (Figure 4.10a) showed a clear preference for positive stimulus steps (i.e. the cell only fired when the stimulus in the current epoch was greater than that of the preceding stimulus epoch), across all conditions whilst the step-size functions of Neuron 3 (Figure 4.10c) were almost flat with a slight bias for positive step-sizes; Neuron 2 (Figure 4.10b) responded mainly to step-sizes above a threshold level that was dependent on the global mean of the stimulating condition.

The constant feature of all the cells in Figure 4.10 is that adaptive step-size functions were not invariant, with respect to mean adapting amplitude, thus adaptation to

stimulus-mean, in the inferior colliculus, was not a function of the stimulus step-size. One might expect then, that as there are parallels between adaptation to stimulus mean in both the barrel cortex (Garcia-Lazaro et al. 2007) and inferior colliculus (Dean et al. 2005), that a similar non-dependence on step-size would be observed in the barrel cortex.

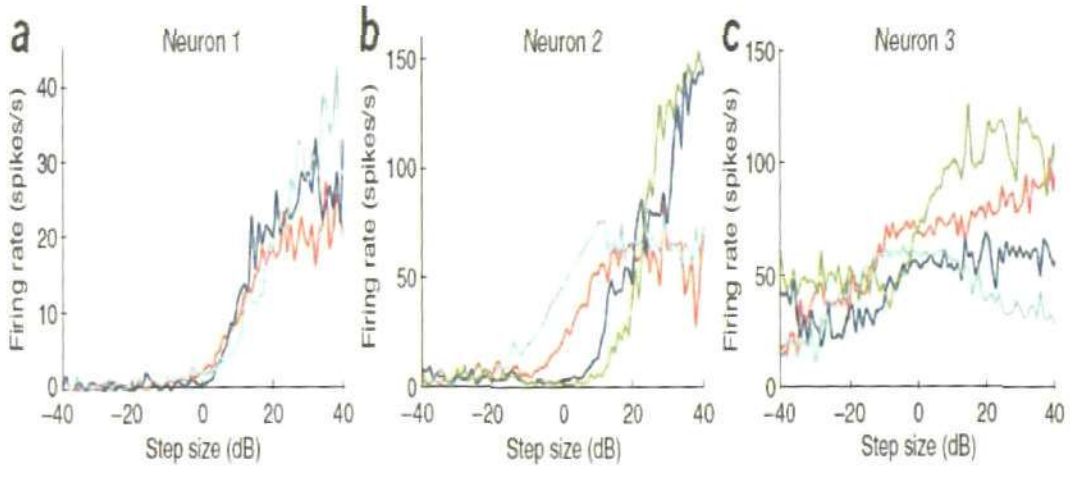

Figure 4.10: Firing-rate as a function of the step-size difference between the current and immediately preceding stimulus epoch for three example neurons from the inferior colliculus.
Firing rate as a function of the step-size difference in noise burst amplitudes (between the stimulating decibel level and the preceding stimulus) for three example Inferior Colliculus cells. Cells were adapted to shifts in the high-probability region of the stimulus space and each coloured line represents the results following adaptation to a given high-probability region level: green – high-probability region centred on 39dB, blue - 51dB, red – 63dB & 75dB. See text in this Chapter and Chapter 3 for full details; Figure 4.10 taken from Dean et al. (2005).

However, it could be argued that the step-size difference in amplitude, from one time-bin to the next, is a measure of stimulus velocity, in that one is describing the relationship between neural output and the direction (either positive or negative) and magnitude of change in intensity levels. A recent study (Maravall et al. 2007) has

demonstrated adaptation to whisker deflection velocity (as a function of stimulus variance) in rat barrel-cortex and found that adaptive neuronal stimulus-response functions were scaled in proportion to changes in the stimulus distribution. Whiskers were deflected along a single dimension, with the position randomly assigned using a Gaussian distribution and refreshed every 5ms; the global variance, rather than mean, of the distribution was varied every 5s in a switching stimulus paradigm. Thus barrel cortex cells were adapted to two levels of variance only (high and low). The experimental paradigm differed from the adaptation to stimulus statistics study (Garcia-Lazaro et al. 2007) in several key aspects:

1. The stimulus distribution was truly Gaussian
2. The adapting stimulus feature was velocity and not amplitude
3. The nature of the stimulus meant that the variance and mean were dependent on each other, thus adaptation to one, i.e. variance, implied adaptation to the stimulus mean.

The spike-triggered average did not reveal any stimulus/response correlations (due to the rapidly fluctuating and noisy stimulus) thus spike-triggered covariance analysis was preformed on the data. Overall, the study (Maravall et al. 2007) found that the majority of barrel cortex cells displayed adaptation to stimulus velocity (with a time course of 280ms ±180ms). An example from two cells is given below in Figure 4.11. The normalized firing rate as a function of the stimulus projection onto the significant feature (velocity), extracted from the covariance matrix, is plotted in the top row, with blue representing low variance and red the high variance condition. In the bottom row (Figure 4.11), the adapted rate functions (plotted in the top row of Figure 4.11) were plotted as a function of normalised input units (stimuli were normalised by the corresponding standard deviation of the stimulus distribution) and almost perfectly overlapped each other.

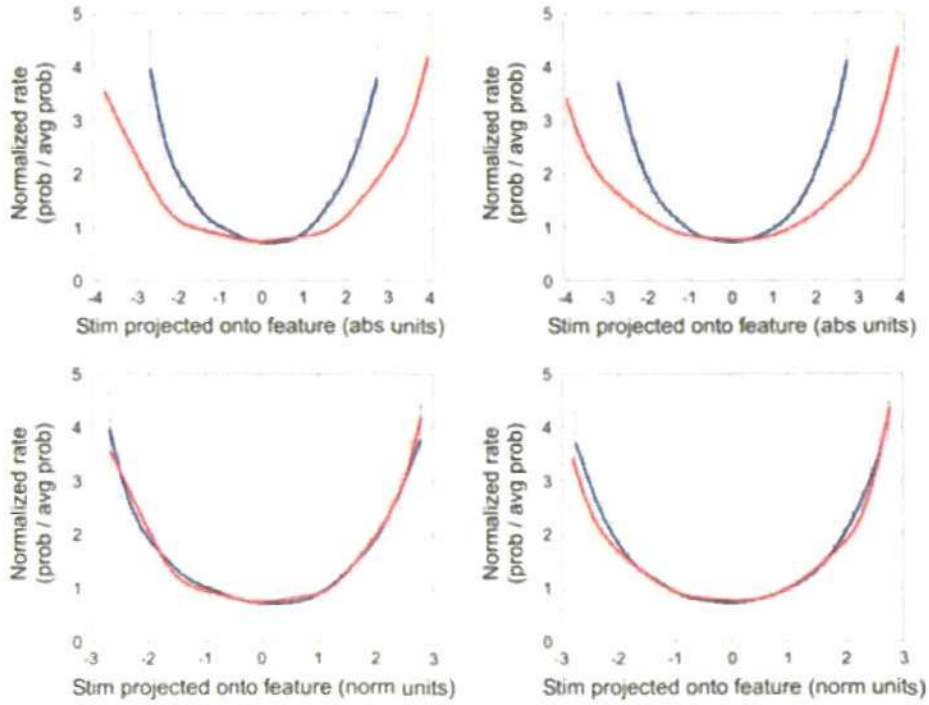

Figure 4.11: Adaptation to velocity in the barrel cortex.
Top row: *The nonlinear, stimulus-response function for low-variance (blue) and high-*variance (red) stimuli. **Bottom row:** Plots the same functions as above, with inputs normalized by their corresponding standard deviation. Error bars are the standard deviation from 30 repetitions of the estimation procedure.

Accordingly, if one considers the step-size changes, every 40ms, in whisker-deflection amplitude (Garcia-Lazaro et al. 2007), to be a velocity measure, then one might expect to see the step-size functions, in response to adaptation to global variance, to be scaled in proportion to levels of stimulus variance.

4.3.2: Response as a function of step-size

The average firing rate, as a function of the difference in stimulus amplitudes

between the stimulating and preceding 40ms stimulus epoch (for each high-probability region stimulus, across trials and every presentation within a trial) was calculated for each individual cell and across the population as whole.

For every 40ms time-bin ($T_{i=2 \ldots N}$, where N is the total number of stimuli presented during a 30s stimulus sequence), the spikes elicited (by a given neuron) during T_i were re-sorted into bins according to the difference in amplitude between the stimulus being presented during T_i (given a 20ms latency, see Chapter 3) and the stimulus presented in the preceding time-bin (T_{i-1}). As there were only ever 25 possible stimulating amplitudes (0-720μm, in 30μm steps) that could be presented during T_i, the step-size difference in amplitude could only be one of 49 possible values (-720 to 720μm, in 30μm steps). The response to the step-size was taken to be the average spike count across all instances in which the difference between the stimuli presented at T_i and T_{i-1} was equal to the given step-size difference in stimulus amplitudes.

As in the inferior colliculus, responses across the barrel cortex population were varied. In Figure 4.12 the step-size functions of four sample cells are plotted. For the two neurons in the top row of Figure 4.12, firing rates increased monotonically for step-sizes greater than $\approx$ 0μm; in general, responses were roughly invariant with respect to the adapting mean amplitude (except for the cell in right-hand plot which was roughly invariant with respect to the three highest mean adapting amplitudes only). Both neurons in the bottom row of Figure 4.12 displayed a preference for step-sizes in either direction away from 0μm, with the cell in the right-hand plot showing a bias towards positive step-sizes and the left-hand cell having an almost symmetrical function.

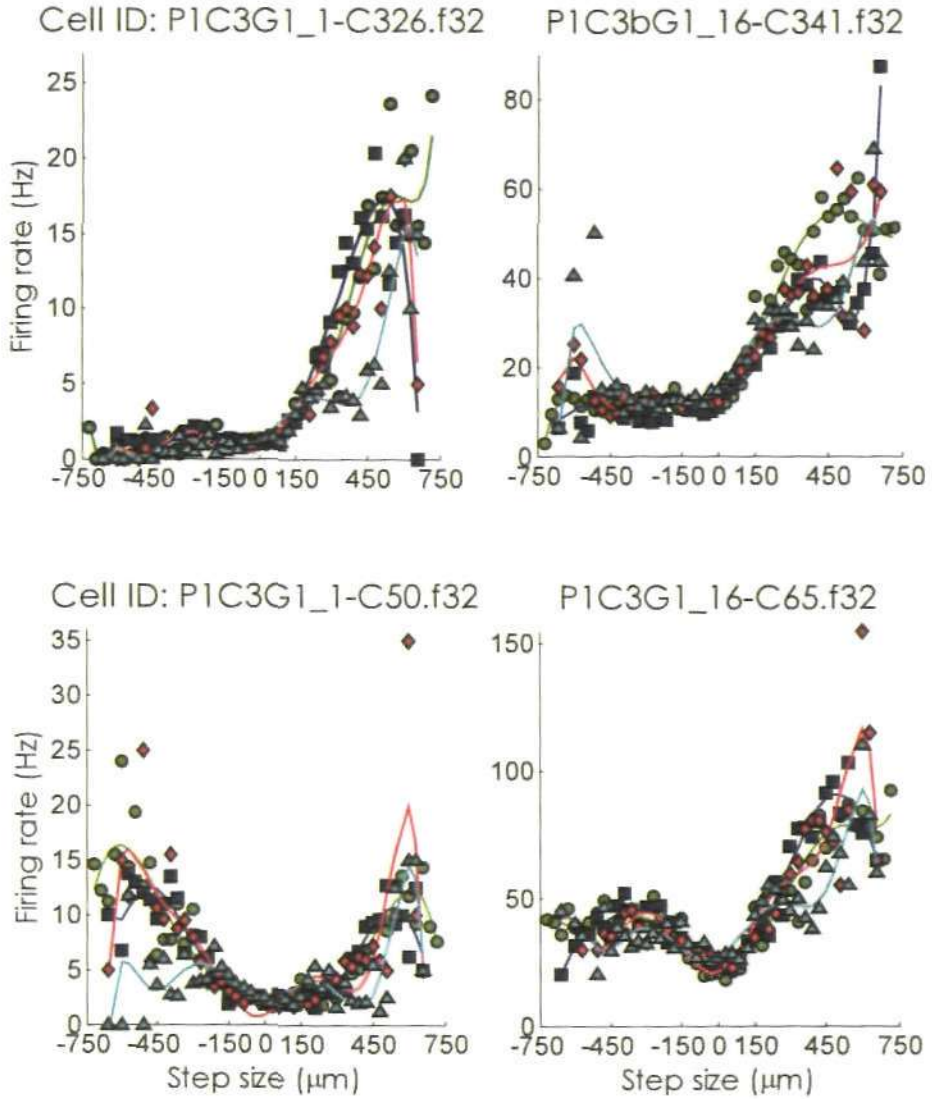

Figure 4.12: Step-size functions under conditions of adaptation to global mean, for four sample cells.
Green lines/circles plot the data for the 90μm high-probability region level, blue lines/squares: 210μm, red lines/diamonds: 330μm and cyan lines/triangles: 450μm; the solid lines are the best fit to the data using a 6th order polynomial and symbols plot the actual data.

The response across the population as a whole is plotted in Figure 4.13. The best-fit to the population response, as a function of stimulus step-size, was roughly invariant across

all adapting mean-amplitude conditions; the closest overlap between step-size functions was observed between the two middle high-probability region conditions (red and blue solid lines, Figure 4.13).

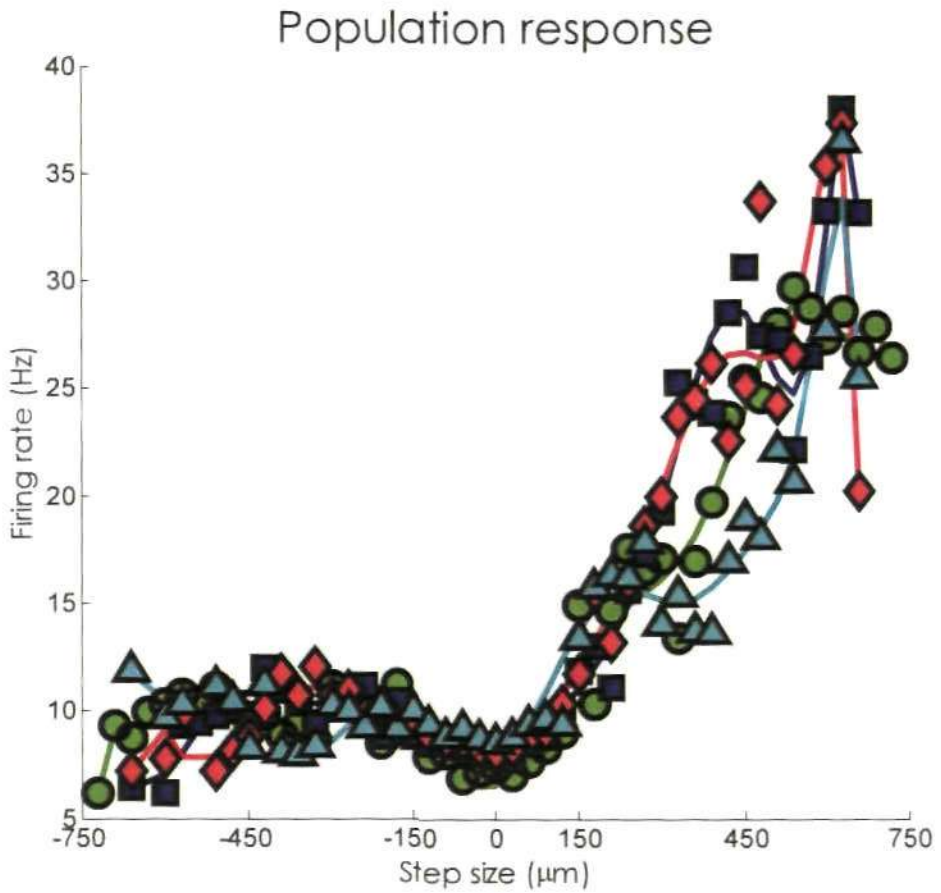

Figure 4.13: Step-size functions under conditions of adaptation to global mean, averaged across the population (*n=81*).
Green lines/circles plot the data for the 90μm high-probability region level, blue lines/squares: 210μm, red lines/diamonds: 330μm and cyan lines/triangles: 450μm; the solid lines are the best fit to the data using a 6th order polynomial and symbols plot the actual data.

The only significant difference ($p=0.004$, using a Kruskal-Wallis, one-way ANOVA, MatLab function *kruskalwallis*) was found between the population step-size

function for the highest mean amplitude condition (cyan lines, Figure 4.13) and the middle two high-probability region conditions (red and blue solid lines, Figure 4.13). Firing rates tended to increase as a function of distance from the 0μm step size (up to approximately ±600μm); with a strong response-bias for positive step-sizes (i.e. where stimulus-amplitude at T_t was greater than that presented during T_{t-1}).

The selection probability of individual amplitude transitions (Figure 4.14, top) and their associated firing rates (Figure 4.14, bottom), across the whole population of cells ($n = 81$) are plotted in Figure 4.14, below, for two of the four mean adapting levels. Rows represent the amplitude level of the current epoch (the stimulating amplitude), whilst columns represent amplitude levels in the preceding stimulus epoch. The two uppermost plots show the average number of times, in a trial, stimuli were presented concurrently. The likelihood of presentation clearly depends on the location of the stimuli, relative to the high-probability region; the most numerous amplitude interactions were restricted to amplitudes within the high-probability region.

The bottom row plots the average firing rate across the population for each interaction. For the lowest high-probability region level (90μm condition, Figure 4.14, left, top & bottom), the strongest firing-rates were found only when the stimulating amplitude (i.e. the amplitude of the current stimulus epoch) was greater than the high-probability region. This is unsurprising given the fact that, under this condition, the lowest amplitude levels were within the high-probability region. However, this trend is also apparent for the high-probability region level with the strongest intensity (450μm condition, Figure 4.14, right, top & bottom). Thus, over the population as a whole, the response bias is for positive step-sizes where the stimulating amplitude is greater than the amplitudes located within the

high-probability region of the stimulus space.

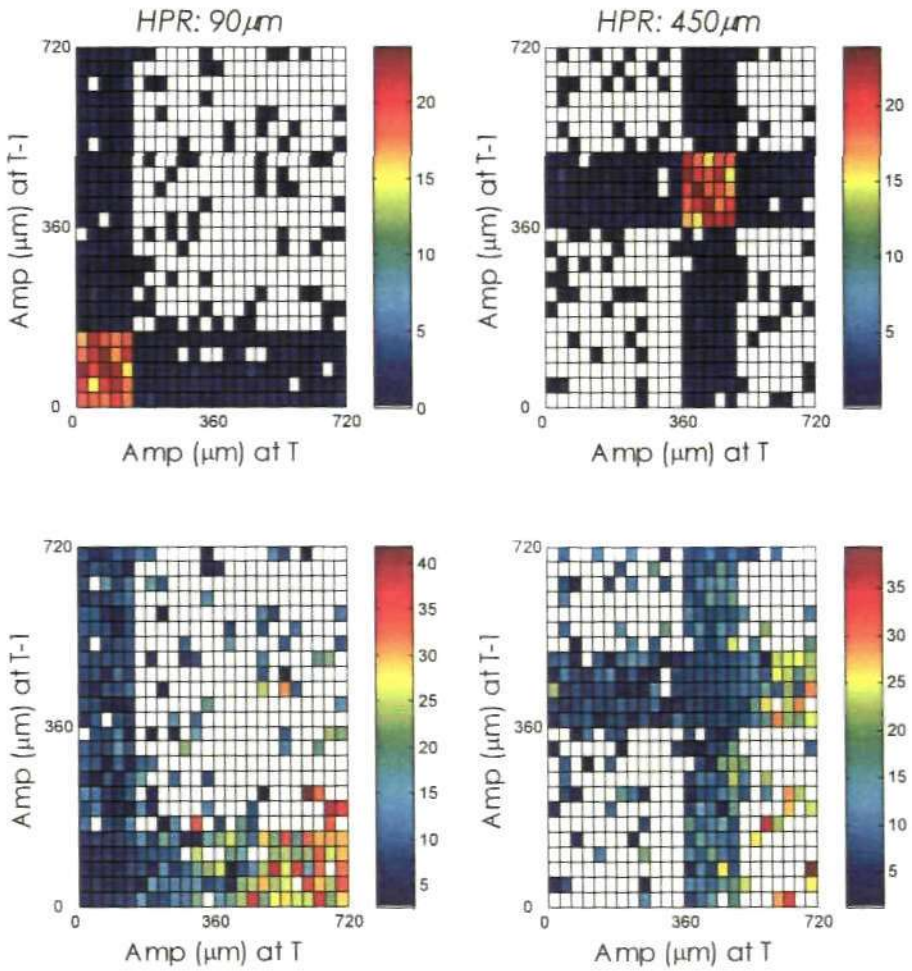

Figure 4.14: Relationship between presentation frequency of amplitude steps and neural output, across the population (n=81).
Average number of presentations, per trial (top), and firing-rate (bottom) as a function of the current (time T, rows) and preceding (time T-1, columns) stimulus amplitude. Colour bars to the side of each plot indicate the number of presentations (top) and firing-rate in Hz (bottom). White indicates that there was no interaction between amplitudes

The first major point of interest is the rough invariance of the population step-size function across all global mean levels. From this result, it would seem that adaptation acts

primarily to maintain the relationship between firing-rate and step-size, rather than to transmit stimulus distribution information.

Before exploring this issue further, there was one other interesting facet of adaptive behaviour that emerged from the analysis of the step-size functions, namely that some cells (and enough cells overall to influence the population response, see below) appeared to be behaving as novelty detectors, or at the very least - novelty-preferring cells (c.f. step-size functions plotted in Figure 4.12, bottom row). For these neurons, firing rate was dependent on the absolute difference between successive amplitudes and not the direction of amplitude transition. Due to the experimental paradigm, step-sizes near to 0μm were most likely to be derived from amplitudes that were drawn from the high-probability region of the stimulus space, with a maximum difference in step-size of only ±120μm. Overall, within each trial, consecutive pairs of high-probability region amplitudes were presented 64.75% of the time; concurrent high-probability region and low-probability region amplitudes made up 31.2% of the stimulus presentation, with the remaining 4.05% composed of successive pairs of low-probability amplitudes. Thus cells displaying a preference (in terms of firing rates) for step-sizes away from ±120μm are not only responding to large amplitude transitions but also the presence (either as the stimulating or preceding amplitude) of a stimulus from the low-probability region of the stimulus space. Across the population of cells tested (n = 81), roughly a quarter of all cells (22/81 neurons) displayed some preference for step-sizes away from 0μm. Preference was assessed by taking the average gradient of the best-fit to the data, from the 0μm point to the local maximum, in either direction. A cell was assigned as having novelty preference if the average gradient (for all adapting conditions) of the best-fit to the data had a value less than

-0.1, in the negative, amplitude-step, direction, and a value > 0.1, in the positive direction. Of those cells displaying novelty preference, only 4/22 neurons displayed roughly symmetrical step-size functions (these are plotted in Figure 4.15 (top row), along with their corresponding rate-functions, Figure 4.15, bottom row).

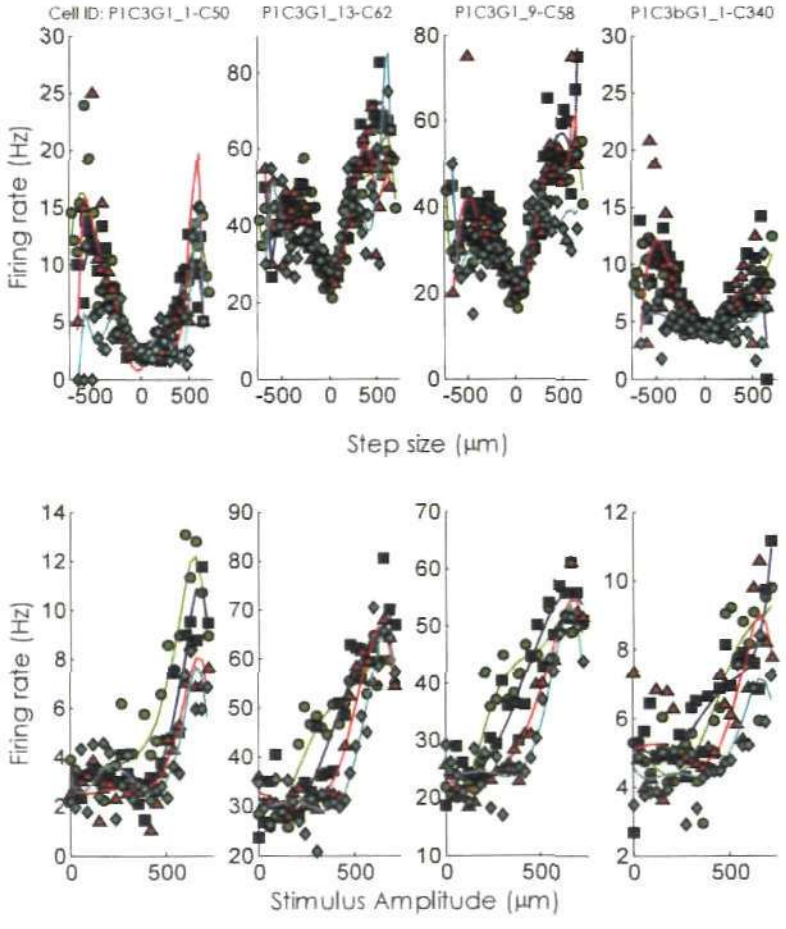

Figure 4.15: The step-size functions for four symmetrical, novelty detecting cells and corresponding rate functions.
Top row: Step-size functions for four symmetrical novelty type cells. **Bottom row:** Corresponding rate-functions for the cells in the top row. Solid lines: best fit (8^{th} and 5^{th} order polynomial were used to fit step-size and rate function data, respectively), symbols: actual data (green: 90μm, blue 210μm; red: 330μm; cyan: 450μm).

One can observe a tendency towards novelty detection in all the 'symmetrical' cells (Figure 4.15, bottom row); rate-functions tend to increase away from a minimal point located close to the global-mean of the stimulus. This trend is also evident in the rate-functions of non-symmetrical cells that show a preference for novel stimuli (see Figure 4.16).

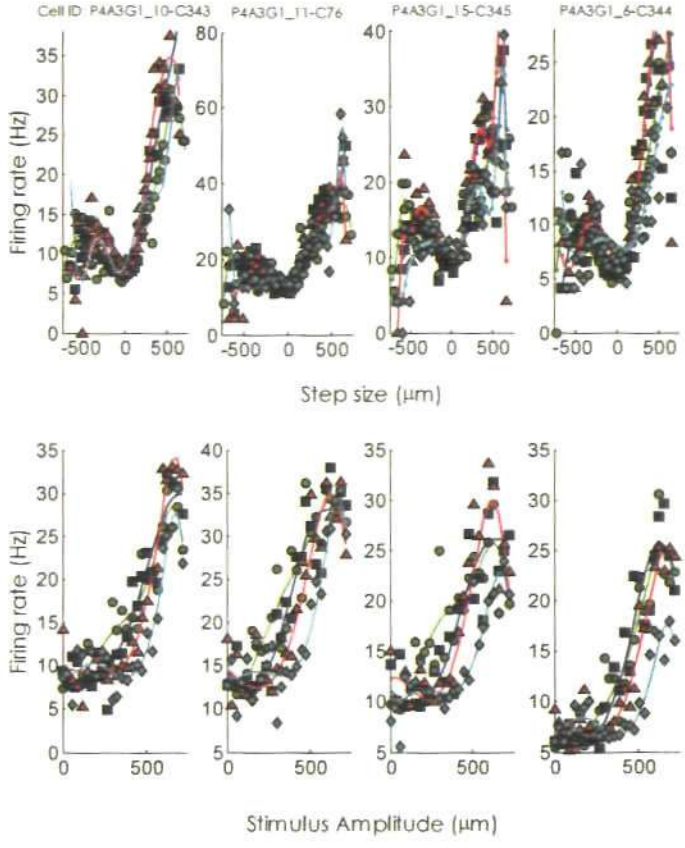

Figure 4.16: The step-size functions for four `asymmetrical novelty detecting cells' and their corresponding rate functions.
Top row: The step-size functions for four asymmetrical novelty type cells. **Bottom row:** Corresponding rate-functions for the cells in the top row. Solid lines: best fit (8^{th} and 5^{th} order polynomial were used to fit step-size and rate function data, respectively), symbols: actual data (green: 90µm, blue 210µm; red: 330µm; cyan: 450µm).

Conversely, the rate-functions of cells with little or no novelty detecting traits (see

Figure 4.17) displayed the sigmoidal shape that was typical of the population response, i.e. the function is flat until the amplitude space reaches the high-probability region; firing rates then increase in a linear, accelerating manner, until the maximum response is reached.

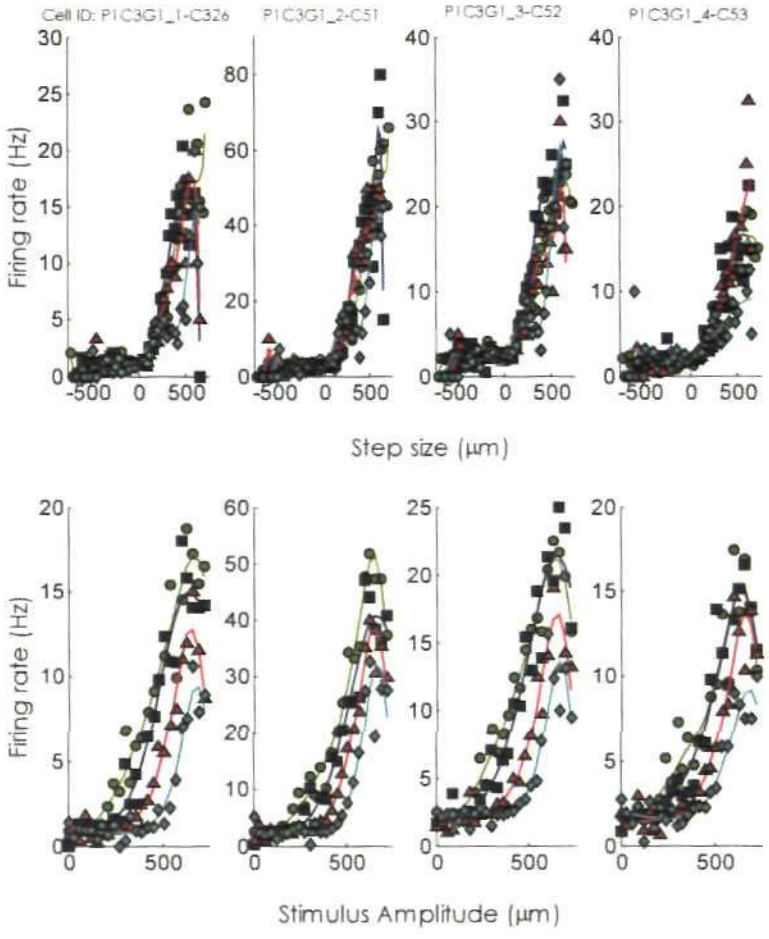

Figure 4.17: The step-size functions for four example non-novelty detecting cells and their corresponding rate functions.
Top row: Step-size functions for four sample non-novelty type cells. **Bottom row:** Corresponding rate-functions for the cells in the top row. Solid lines: best fit (8^{th} and 5^{th} order polynomial were used to fit step-size and rate function data, respectively), symbols: actual data (green: 90µm, blue 210µm; red: 330µm; cyan: 450µm)

Step-size functions under conditions of adaptation to global-variance

The same analysis was performed on neurons that were adapted only to stimulus

variance. The population response is plotted in Figure 4.18, below. Interestingly, the population response was not invariant with respect to step-size, thus adaptation to variance was not dependent on local stimulus history.

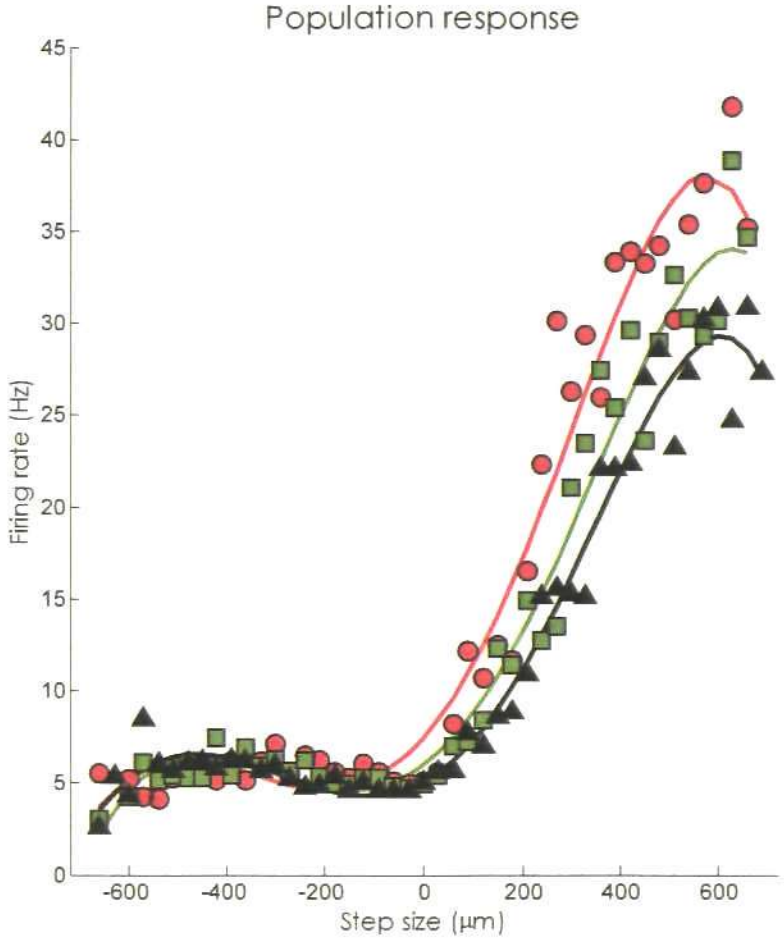

Figure 4.18: Step-size functions under conditions of adaptation to global variance, averaged across the population (*n=132*).
Red lines/circles plot the data for the low variance level, green lines/squares: mid-variance ,black lines/triangles: high-variance: the solid lines are the best fit to the data using a 6th order polynomial and symbols plot the actual data.

A variety of responses were observed across the population. Neurons typically had step-size functions that either displayed a monotonic increase in firing rate from a threshold

amplitude-step (in general, located at ≈0μm, see Figure 4.19, leftmost plot for an example) or firing rates increased either direction from a local minimum (≈0μm) up to amplitudes steps of between ±500-600μm (see Figure 4.19, middle and right).

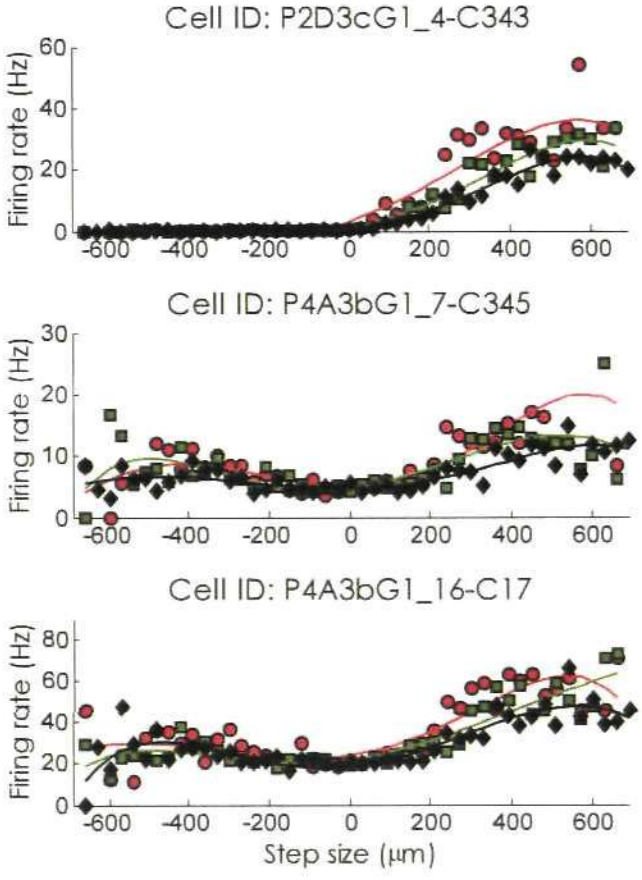

Figure 4.19: The step-size functions for the three example cells, under conditions of adaptation to variance.
Top, middle and bottom row: Step-size functions of three sample cells. Solid lines (red lines: low-variance, green lines: mid, black lines: high-variance) plot best-fit to the data (8th order polynomial); colour coordinated symbols: actual data (circles: low-; squares: mid-; triangles: high-variance adaptation).

Out of the population, roughly 11% of cells (14/132) displayed a preference for large step-sizes in either direction. Of these thirteen only three neurons displayed a degree

of symmetry in the step-size function, i.e. held no particular bias for positive or negative stimulus steps. The 'symmetrical' cells are plotted in Figure 4.20, along with their corresponding rate functions.

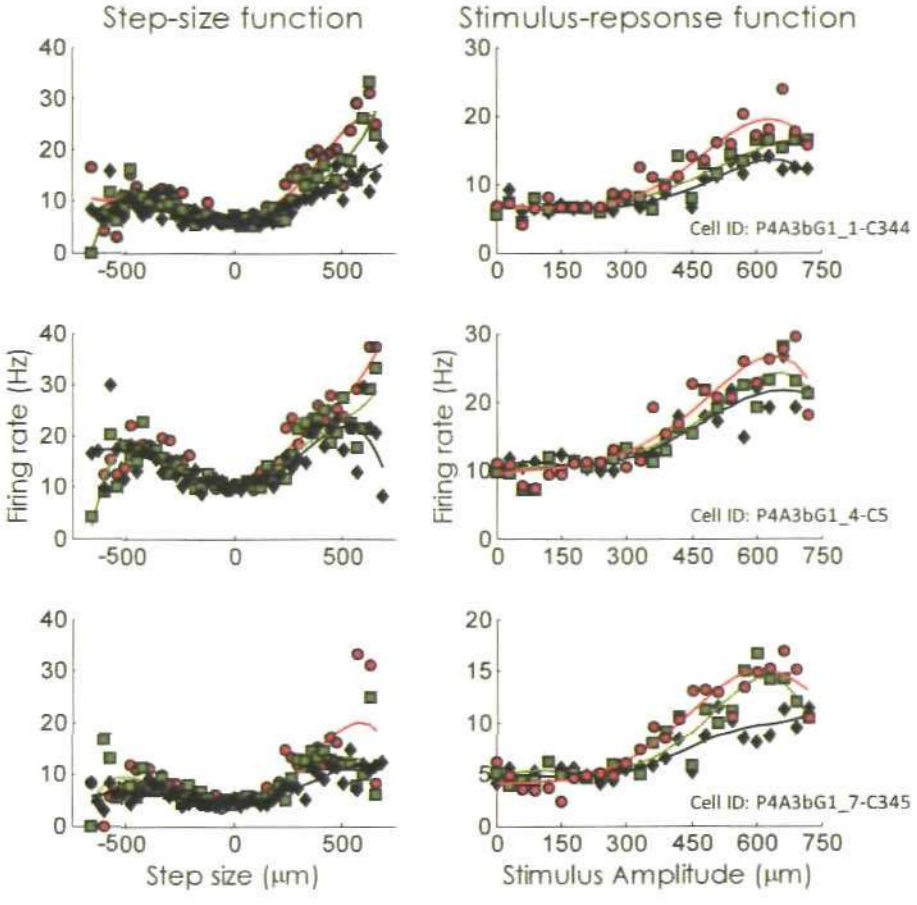

Figure 4.20: The step-size functions for the three, identified, `symmetrical novelty detecting cells' and their corresponding rate functions, under conditions of adaptation to variance.
Left column: The step-size functions for the three identified, symmetrical novelty type cells. **Right column:** Corresponding rate-functions for the cells in left column. Solid lines (red: low-, green: mid, black: high-variance) plot the best-fit to the data (8^{th} and 5^{th} order polynomial were used to fit step-size and rate function data, respectively)

The first observation to note is that the rate-functions, for these 'symmetrical' cells, display some preference for novel stimuli (Figure 4.20, right column), especially for the

lower variance conditions, in that responses tend to increase, in either direction, away from a local minimum (located between 90-150μm).

In comparison, the rate-functions of neurons, with no sensitivity for negative step-sizes, displayed little tendency for novelty preference (Figure 4.21, right column).

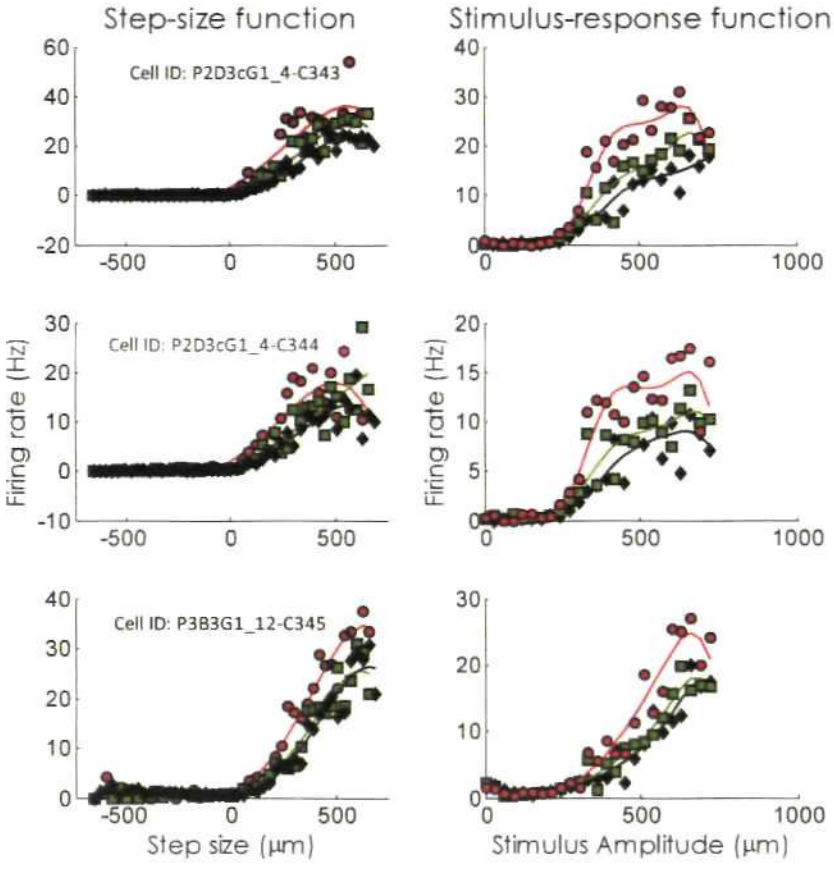

Figure 4.21: The step-size functions for three example non-novelty detecting cells and their corresponding rate functions, under conditions of increasing stimulus variance. Left column: The step-size functions for three non-novelty type cells. **Right column:** Corresponding rate-functions for the cells in left column. Solid lines (red: low-, green: mid, black: high-variance) plot the best-fit to the data (8^{th} and 5^{th} order polynomial were used to fit step-size and rate function data, respectively).

In general, responses to amplitudes below the high-probability region were fairly flat and constant across all conditions.

The step-size functions plotted in Figure 4.21 (left column) were typical of those observed across the population of cells that only responded to positive step-sizes, in that firing rate increased for step-sizes greater than 0μm until step-sizes of roughly 600μm. Across the population, the step-size function showed some dependence on the width of the high-probability region. For example, the curve around the 0μm point, for novelty-preferring cells (Figure 4.20, left column), was narrowest under the low-variance condition and broadened with stimulus variance (especially with reference to the left-most and middle plots in Figure 4.20, left column). For the majority of cells with no preference for novel stimuli, the maximum firing rate tended to decrease as adapting variance increased and there was also a decrease in step-size function slope.

If one considers the stimulus configuration, the range of likely amplitude interactions increased with each successive increase in the width of the high-probability region (see Figure 4.22, top row). Step-sizes from 0μm to ± 60μm, under low-variance, to ± 300μm, under high-variance conditions were most likely to be derived from concurrent presentations of amplitudes within the high-probability region. This type of interaction tended to elicit relatively low firing rates (Figure 4.22, bottom row) thus possibly resulting in the lateral displacement of the step-size function, observed in the population response (see Figure 4.18), with increasing stimulus variance.

As well as increasing the range of high probability to high-probability amplitude interactions, the number of likely low-probability and high-probability amplitude transitions also increased (these type of interactions occurred ≈30% of the time), with stimulus variance. Likely step-size values ranged from ±90-510μm under the low-variance condition, to ±330-630μm, under the high-variance environment.

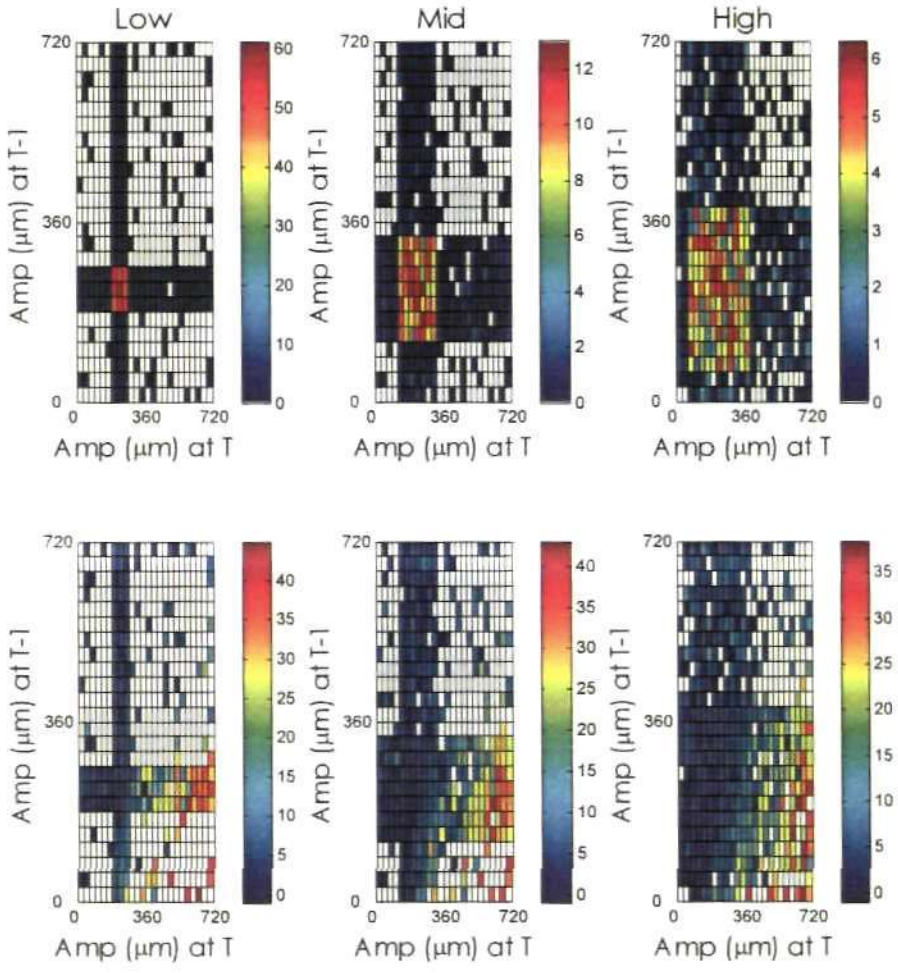

Figure 4.22: Relationship between presentation frequency of amplitude steps and neural output, across the population (n=81), under conditions of increasing stimulus variance.

Plots of the average number of presentations, per trial (top), and firing-rate (bottom) as a function of the stimulating (T_i, columns) and preceding (T_{i-1}, rows) amplitude levels, for each variance condition. Colour bars to the side of each plot indicate either the number of presentations (top) or firing-rate in Hz (bottom). White indicates that there was no interaction between amplitudes

The broadening of the step-size functions (and the adaptive decrease in the maximum response) could be a reflection of the increased diversity in amplitude

interactions.

Spiking behaviour, on the ms scale, during amplitude transition

The evidence above implies that adaptation is not dependent on the global mean level of whisker deflection but on local changes (on the 40ms scale) in stimulus amplitude. Neural responses can be classified into two groups, relative to the step-size function: neurons that only respond to positive changes in amplitude transitions (the majority of cells) and neurons that respond to the step-size irrespective of the direction of transition. Cells that fall within the latter classification, which were termed novelty-preferring, can be subdivided into two further groups according to the degree of symmetry in the step-size functions. Novelty-preferring cells could also be classified by their stimulus-response functions which tended to show an increase in firing away from a local minimum response, in both directions; the local minimum was located at or near the mean adapting amplitude.

It was decided to analyse the relationship between step-size and firing-rate further, by plotting the frequency of amplitude interactions and the corresponding response (see Figure 4.23). Overall, the strongest firing-rates were found for positive step-sizes where the stimulating amplitude (Figure 4.23, bottom, all rows) was greater than the maximum high-probability amplitude and for negative step-sizes where the stimulating and preceding amplitudes were within and higher than the high-probability region, respectively.

It was decided, therefore, to perform an analysis of the firing rate, as a function of time, during the actual transition between amplitudes to see how the direction and distance of the step-size affected responsiveness in both novelty-preferring cells and the majority of the neural population.

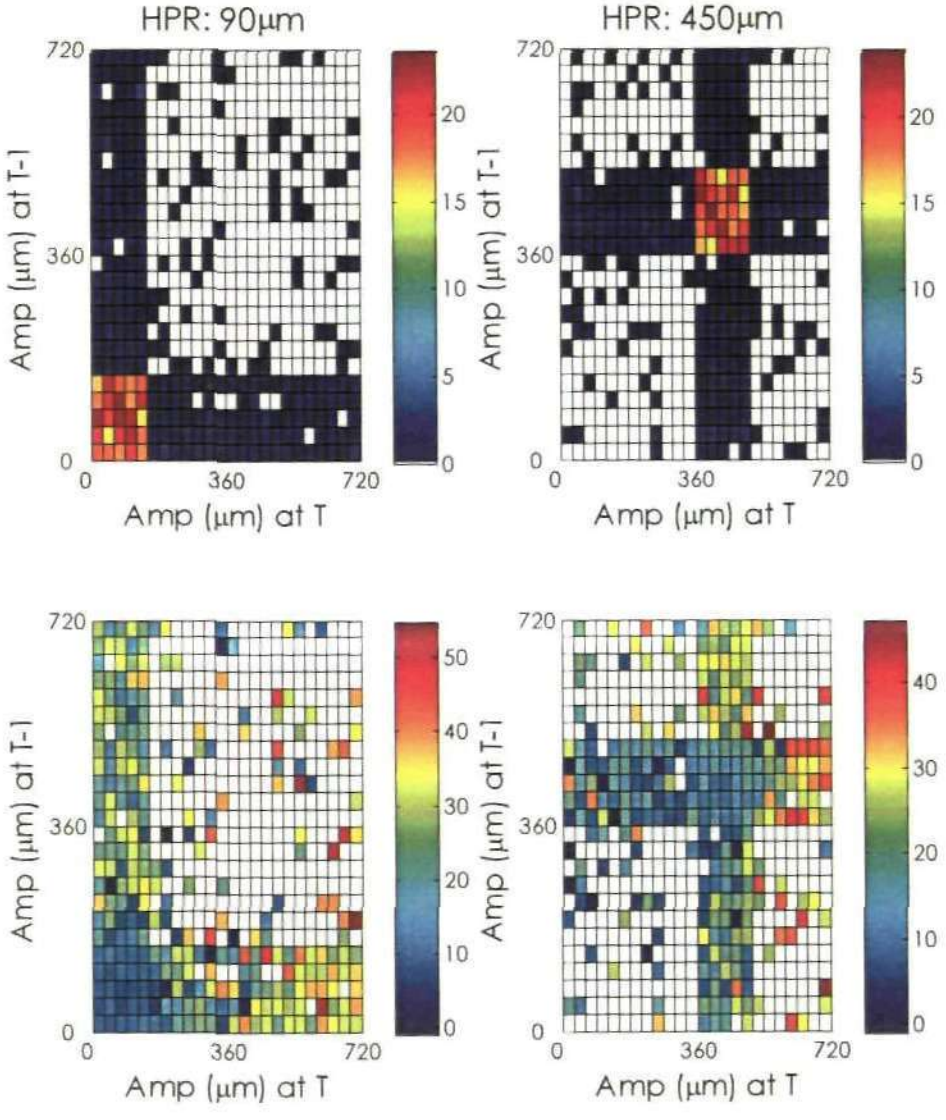

Figure 4.23: Relationship between presentation frequency of amplitude steps and neural output, averaged over the four symmetrical, novelty cells, under adaptation to global mean values of 90µm and 450µm.
Average number of presentations, per trial (top), and firing-rate (bottom) as a function of the stimulating (T_t, columns) and preceding (T_{t-1}, rows) amplitude levels. Colour bars to the side of each plot indicate the number of presentations (top) and firing-rate in Hz (bottom). White indicates that there was no interaction between amplitudes

In deciding which amplitudes to compare, it was felt that the best method of

maximising spike count information would be to compare the transition between all amplitudes within the high-probability region (and treat them ostensibly as a single amplitude) and all amplitudes within one of the low-probability region (LPR) 'tails', i.e. either the low-probability tail in the stimulus space where stimuli were stronger in intensity than the high-probability region (now termed LPR-high) or the tail where stimulus amplitudes were lower in intensity than the high-probability region (now termed LPR-low). As the response-onset latency was assumed to be 20ms and response-decay latency assumed to be 40ms (in Garcia-Lazaro et al. (2007), and all previous analyses, presented within this thesis, spikes were counted during the 20-60ms time-window from stimulus onset), the spike trains across 120ms of stimulus presentation (starting from the presentation of the pre-transition stimulus amplitude at 0ms) were collated so that the response to both the pre- (0-40ms into spike train) and post-transition stimuli (40-80ms into spike train) would have time to both develop and decay. The spike counts for each presentation, and across all trials of the relevant adapting condition, were sorted into 1ms time bins (with the onset time of the pre-transition amplitude being set to 0ms), and then averaged across trials and all cells used in the analysis. Only the first two amplitudes (pre- and post-transition stimulus amplitudes) of the 120ms epoch were pre-selected, i.e. were set to be either from the high-probability region or LPR-high/LPR-low. Responses were then normalised, and the maximum response, across all conditions, was set to unity.

Figure 4.24 plots the average firing rate plotted over 120ms, for a transition from LPR-high to the high-probability region (negative step-size), for all mean adaptation levels and across all the cells within the population that did not display any preference for novel stimuli.

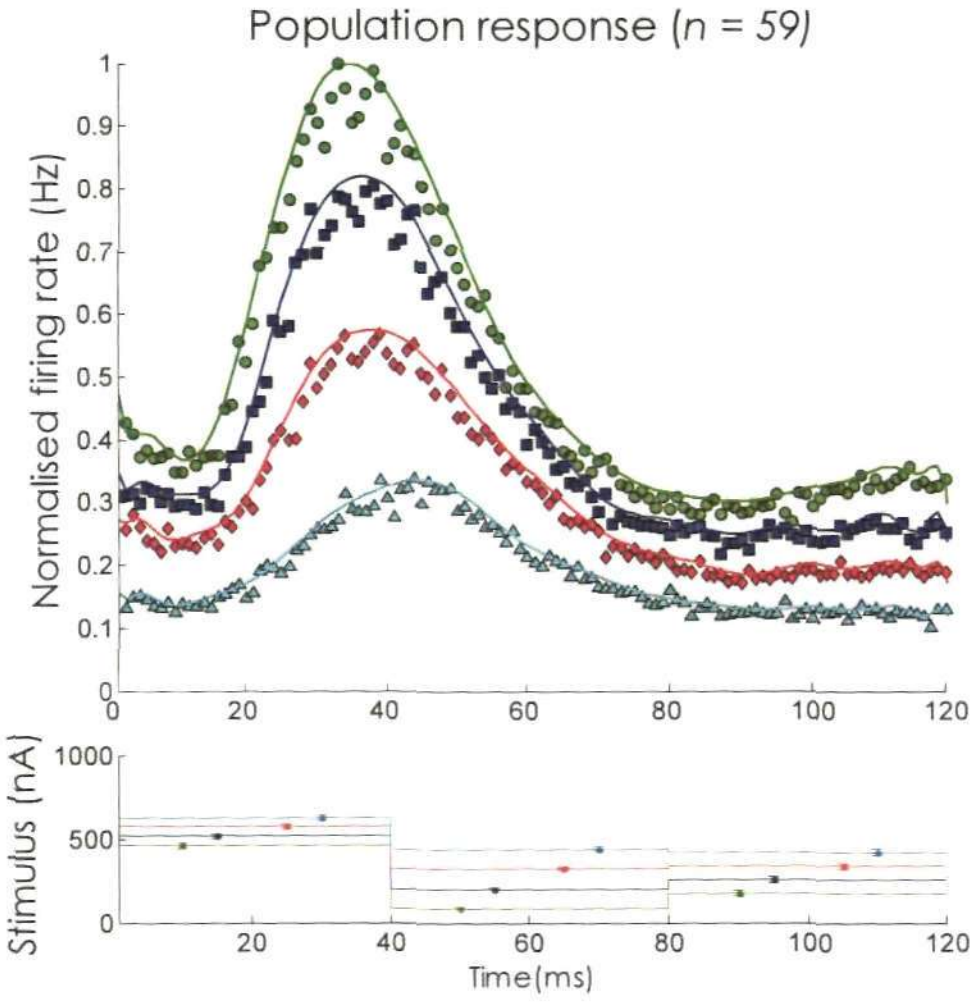

Figure 4.24: Transition from a low-probability stimulus, of strong intensity, to a high-probability stimulus, adaptation to mean.
Firing rate over time, averaged over all transitions between amplitudes drawn from the high end of the low-probability region and amplitudes drawn from the high-probability region. Green circles: 90μm mean adapting level, blue squares: 210μm, red diamonds: 330μm, cyan triangles: 450μm. Colour coordinated solid lines plot the best-fit to the data

The firing rate, over time, followed a similar pattern across all adapting conditions. The response to the LPR-high, pre-transition stimulus started approximately 10s after stimulus onset, then peaked at 35-40ms (≈42ms for highest mean adapting amplitude [cyan], see Table 4.1, below); the response then began to decay and was fully extinguished

roughly 80ms after stimulus onset. The average latency values, across the population, for time-to peak and decay of the response to the preceding amplitude are given in Table 4.1, below.

Table 4.1: Average latency times for the high amplitude/low-probability stimulus

Latency (Δ) from stimulus onset	Mean Amplitude Level			
	90μm	210μm	330μm	450μm
Δ Peak (ms)	35.81 ± 0.91	37.15 ± 0.85	38.03 ± 0.97	42.53 ± 0.83
Δ Decay (ms)	73.76 ± 1.86	71.64 ± 2.23	70.17 ± 2.16	68.42 ± 2.02

The latency of the peak-response to the pre-transmission stimulus was determined by calculating the time taken from stimulus onset (at $T = 0ms$) to the first 'localised' peak response. The peak was defined as the point at which the curve of best fit to the data stopped increasing from its value at $T = 20$ms (assuming an onset-response latency of 20ms, see above). The latency of the decay response was taken to be the time from stimulus onset next local minimum response following the peak response described above, i.e. where the curve of the best fit to the data stopped decreasing from the value of the peak-response and either began increasing or attained the value of the steady state response for the post-transition stimuli (taken to be the average response to the last 80-120ms of stimulus presentation).

The response to the high-probability, post-transition stimulus (from ≈60ms into the spike train) is difficult to extract from the decay response to the pre-transition stimulus. However, the firing rate reached a local minimum value ≈70ms into the spike train and

either reached a plateau (330μm & 450μm mean adapting level) or increased slightly (90μm & 210μm condition). The average pre- and post- transition stimuli levels are displayed in the bottommost plot of Figure 4.24. The increase in firing rates from the local minimum was only observed when the following stimulus transition was, on average, higher in amplitude.

The overall, relative, response levels decreased with each successive increase in the mean of the adapting stimulus. As global mean increased, the step-size difference between the high-probability and low-probability amplitudes (Figure 4.24, bottom) became smaller as the number of amplitudes within the low-probability 'tail' decreased (due to the highly-skewed nature of the stimulus distribution; *c.f.* Chapter 3). Thus a decrease in response, as a function of step-size, is to be expected. However, it should be noted that the decrease in available amplitudes would also affect the number of presentation times, thereby reducing the number of data points available for analysis.

A further example of a negative step-size was examined (Figure 4.25), this time the amplitude transition was from a high-probability to a low-probability stimulus, of lower amplitude.

Overall, response levels were fairly flat until ≈90ms, across all mean-adapting levels except for the 450μm condition; it was not possible to get any data for the 90μm condition as no post-transition stimuli were of lower amplitude than the pre-transition, high-probability stimulus.

The quantitative analysis, below, is based on the adaptive output for the 450μm condition only (Figure 4.25, cyan triangles and solid line). The average peak response latency was 41.75ms (±1.98 Standard Error or SE), and the average decay response latency

was 96.41ms (±4.34 SE) from stimulus onset at 0ms (bearing in mind that the decay response was a relaxation of the firing-rate response to the amplitude transition).

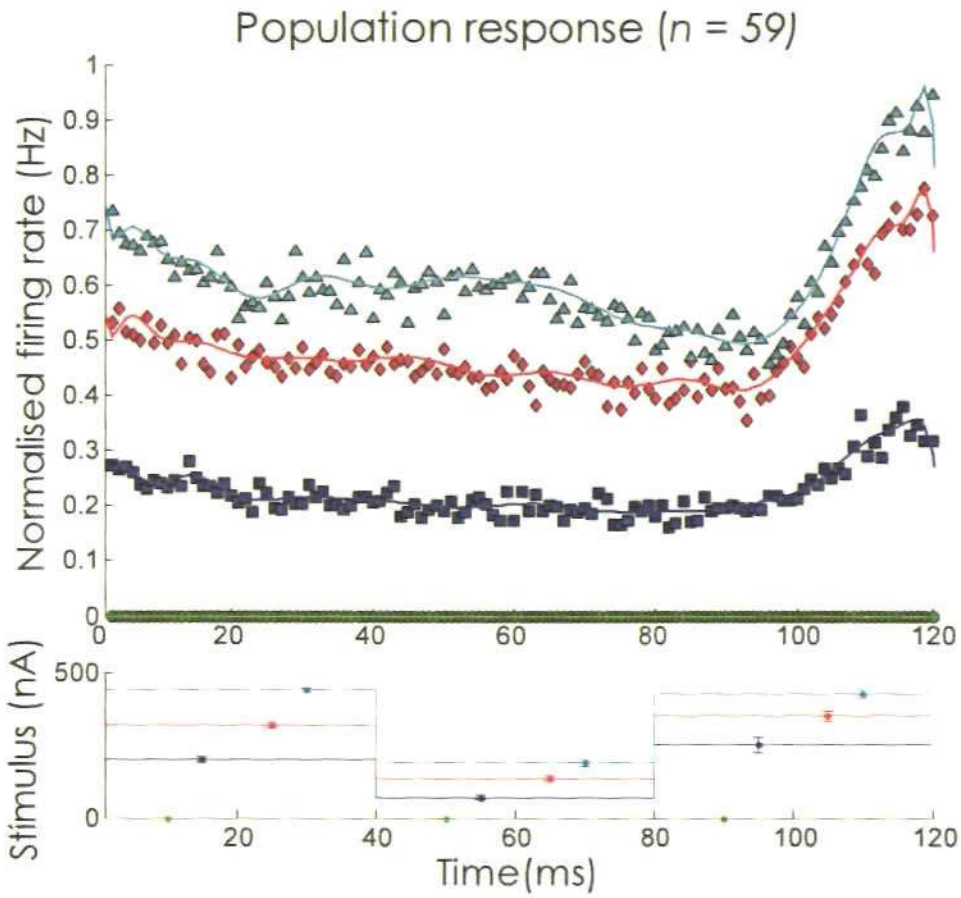

Figure 4.25: Transition from a high-probability stimulus to a low-probability stimulus of lower intensity (adaptation to mean).
Firing rate over time, averaged over all transitions between amplitudes drawn from the high-probability region and amplitudes drawn from the low end of the low-probability region. Green circles: 90µm mean adapting level, blue squares: 210µm, red diamonds: 330µm, cyan triangles: 450µm. Colour coordinated solid lines plot the best-fit to the data

The response to the post-transition stimulus either didn't reach or remained at a steady-state level, however the response to the following amplitude transition (from 80ms into the stimulus presentation) was characterised by a sharp increase in firing-rate that reached a peak at or close to the 120ms time-bin.

If one examines the plot of the average stimulus levels for each transition, it is clear that the sharp increase in firing rate was in response to the positive step-size change in stimulus amplitude at the 80ms point (in general, due to the stimulus configuration, it is most likely that the stimulus to be selected, following the transition to a low amplitude and low probability stimulus, would be of a higher intensity).

The response to this class of transition, i.e. from a low-probability, low amplitude stimulus to high-probability stimulus, is plotted in Figure 4.26; the average time-to-peak response ranged from 34-42ms from the onset of the amplitude transition (average responses, across conditions are given in Table 4.2, below, the ± values gives the standard error). Thus the time-to-peak latency, for positive amplitude transitions, increased with step-size.

Table 4.2: Average latency times, from amplitude transition at the 40ms time-bin to peak response, for the stimulating, high-probability amplitude, across all neurons with no novelty-preference

Latency (Δ) from stimulus onset	**Mean Amplitude Level**			
	90μm	210μm	330μm	450μm
Time-to-peak (ms)	N/A	34.86 ± 1.71	41.17 ± 1.38	41.66 ± 1.09

A similar analysis was performed on only those cells in the population that displayed a novelty preference and had either a symmetrical (Figure 4.27, population = 4) or a non-symmetrical function step size function (Figure 4.28, population = 18)

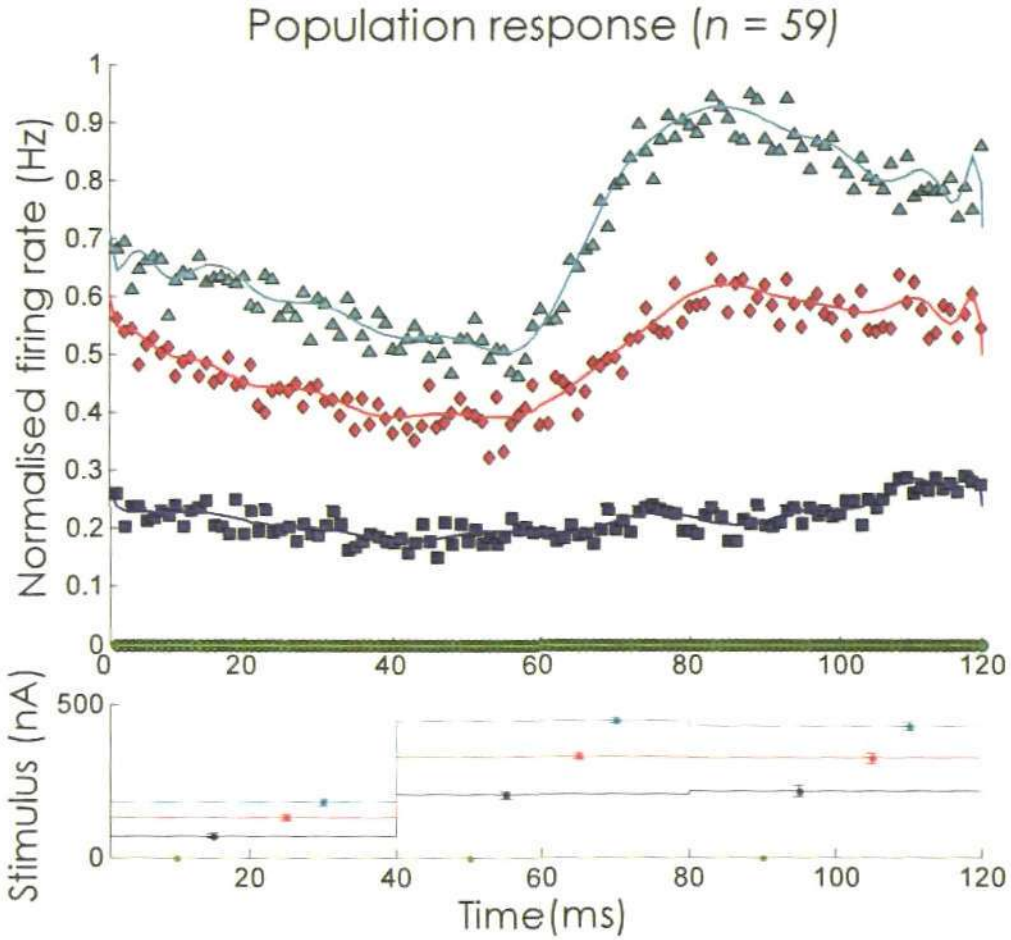

Figure 4.26: Transition from a low-probability stimulus, of lower intensity, to a high-probability stimulus; adaptation to mean.
Firing rate over time, averaged over all transitions between amplitudes drawn from the low end of the low-probability region and amplitudes drawn from the high-probability region. Green circles: 90μm mean adapting level, blue squares: 210μm, red diamonds: 330μm, cyan triangles: 450μm. Colour coordinated solid lines plot the best-fit to the data

The latency of time-to peak and decay for the pre-transition amplitude is given in Table 4.3, below.

Table 4.3: Latency times for the low-probability/high amplitude stimulus for all four novelty-preferring cells with symmetrical step-size functions

Latency (Δ) from stimulus onset	Mean Amplitude Level			
	90μm	210μm	330μm	450μm
Δ Peak (ms)	56.00 ± 4.65	60.50 ± 5.50	57.25 ± 4.05	56.00 ± 4.50
Δ Decay (ms)	94.00 ± 13.0	84.25 ± 9.37	101.0 ± 2.48	81.75 ± 13.0

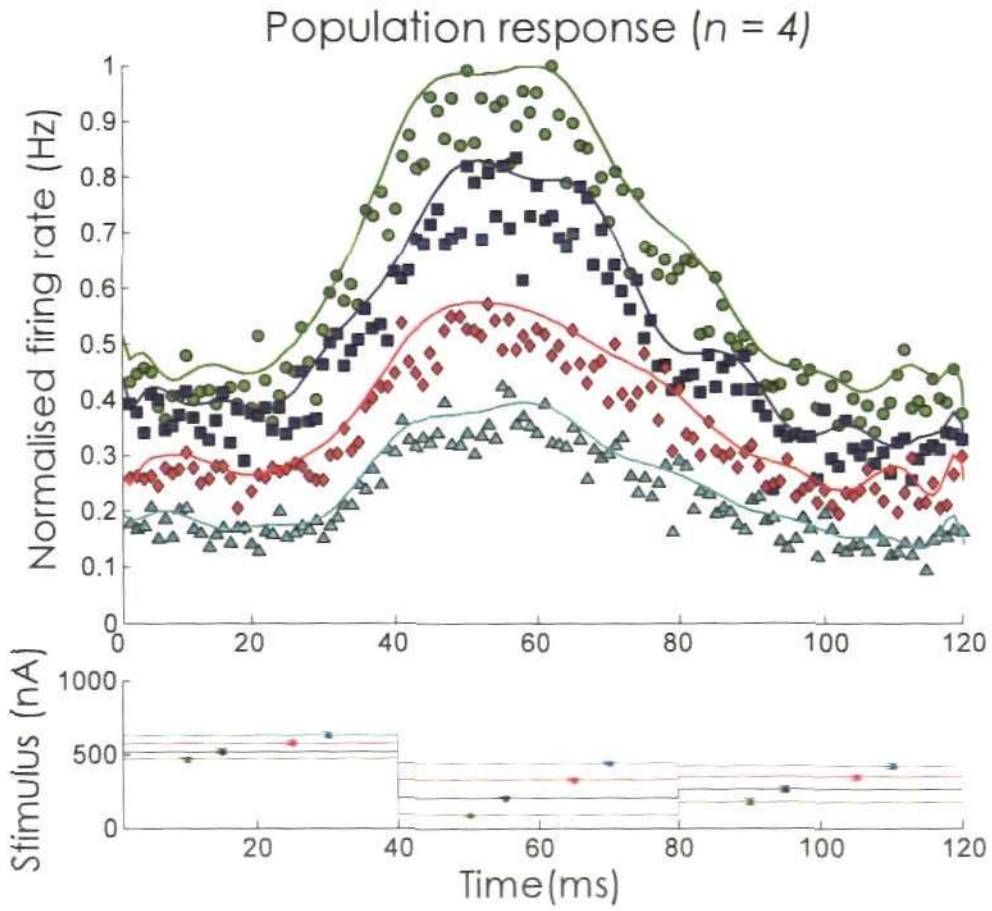

Figure 4.27: Transition from a low-probability stimulus, of strong intensity, to a high-probability stimulus, adaptation to mean, for the four symmetrical novelty-type cells. Firing rate over time, averaged over all transitions between amplitudes drawn from the high end of the low-probability region and amplitudes drawn from the high-probability region. Green circles: 90μm mean adapting level, blue squares: 210μm, red diamonds: 330μm, cyan triangles: 450μm. Colour coordinated solid lines plot the best-fit to the data

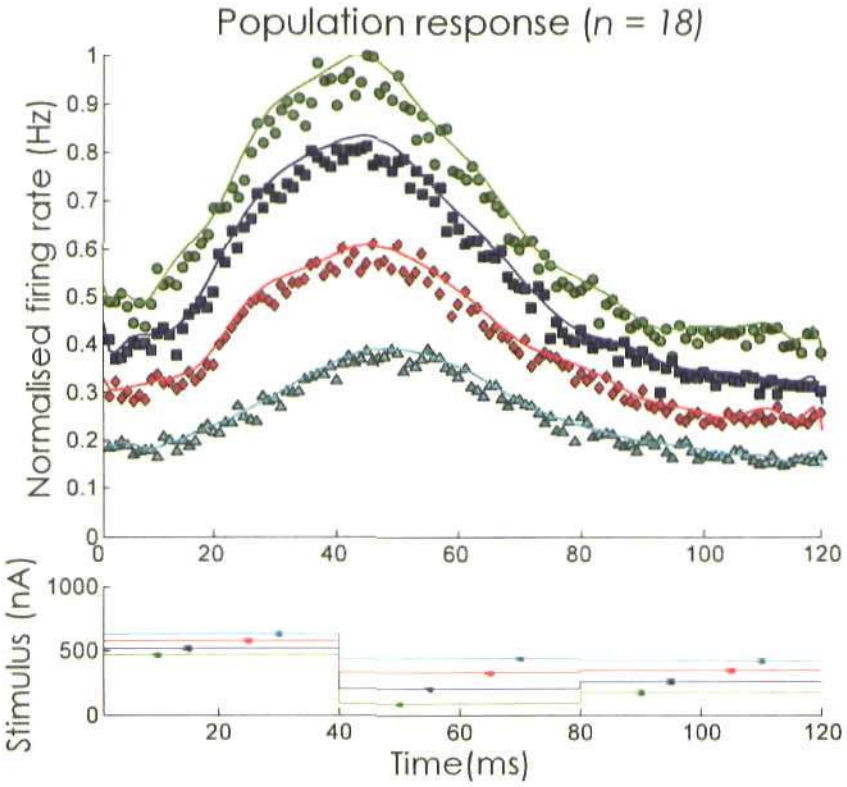

Figure 4.28: Transition from a low-probability stimulus, of strong intensity, to a high-probability stimulus, for all asymmetrical novelty-type cells.
Firing rate over time, averaged over all transitions between amplitudes drawn from the high end of the low-probability region and amplitudes drawn from the high-probability region. Symbols/lines as for previous Figure.

The average latencies, across the population of non-symmetrical novelty-preferring cells, are given in Table 4.4, below.

Table 4.4: Latency times for the pre-transition stimulus (low-probability but high amplitude) for all asymmetrical novelty-preferring cells

Latency (Δ) from stimulus onset	Mean Amplitude Level			
	90μm	210μm	330μm	450μm
Δ Peak (ms)	38.50 ± 3.00	41.33 ± 3.14	41.61 ± 2.79	44.16 ± 2.64
Δ Decay (ms)	79.17 ± 4.21	80.89 ± 4.10	76.22 ± 5.27	71.50 ± 5.40

All the populations analysed shared some common features in the pattern of firing rate over time. Namely that the overall level of response was relatively lower with each increase in global mean and that the response to the pre-transition stimulus appeared to dominate the response within the first 80ms or so after stimulus onset (at 0ms). However, it should be noted that the decay of the pre-transition stimulus included the relaxation of the firing rate back to levels typical of the response to post-transition stimulus (as, if one recalls, these were always lower in amplitude than the pre-transition stimulus). Nevertheless, it was virtually impossible to determine where the response to high-probability, post-transition stimulus began from the firing rate patterns above; instead one could say that the response to the post-transition stimulus reached its steady state level at the decay latency for the high-amplitude, pre-transition stimulus.

Within the population of novelty-type cells, those cells with symmetrical step-size functions showed the clearest divergence from the general trend in firing-rate patterns (see Figure 4.27), in that their response to the pre-transition stimulus was relatively broad. The latency differences were also most pronounced for this group of cells with latencies for time-to-peak (see Table 4.3) being roughly 15-20ms longer than the remaining population. Novelty-type cells with non-symmetrical step-size functions did display longer latencies than the general population, but these differences were not significant (as can be seen from standard errors, if one compares Table 4.1 & Table 4.4). The quantitative analysis of latency differences between the three putative groups of cells, of course, should be viewed with some caution due to the small sample size of both types of novelty-preferring cells.

The time to decay was longer for both groups of novelty-preferring cells, relative to the remainder of the population. Thus, this could account for the strong response to step-

sizes that were comprised (in either direction) of a high-probability amplitude and an amplitude drawn from the high-end of the low-probability stimulus space. For example, with negative step-sizes the response to the stimulating amplitude would be amplified by the residual response to the preceding amplitude.

Are these longer decay times a function of the step-size or an inherent part of the responsiveness of these cells? In order to explore this question further, the time course of the response for stimulus sequences in which both the pre- and post-transition stimuli were drawn from the high-amplitude but low-probability region of the stimulus space was analysed, for both novelty-preferring cells and the remainder of the population (novelty-preferring cells, irrespective of the shape of the step-size function, were grouped together in order to provide a reasonable number of data points). This stimulus sequence was chosen as amplitude levels were high, and would elicit a clear response to the onset of the preceding stimulus, but the step-size changes in amplitude during the transition would be very low, thus one would expect to see response reduction as a function of the small difference in amplitude levels. If latencies are a function of step-size, one would expect to see an effect on response decay of the preceding amplitude. In order for the response to fully develop and decay the spike train following 160ms from stimulus onset were collated. The stimulus sequence is similar to that used in the initial analysis of the response time-course (see Figure 4.24, Figure 4.27 & Figure 4.28) with the `addition' of an extra stimulus at the beginning of the sequence.

The time-course of the responses across conditions, for the majority population of cells with preference for positive step-sizes only, are plotted in Figure 4.29 and in Figure 4.30 for the novelty-preferring neurons. The effect of stimulus context on the response,

across both populations, was twofold:

1. The response to the second high-amplitude, low-probability stimulus was reduced, relative to the response to the first pre-transition stimulus; thus both populations underwent an adaptive and rapid response reduction as a function of the small step-size.
2. The decay latency for the response to reach the steady-state level of the third amplitude in the sequence (of a weaker amplitude than the pre-transition stimulus and most likely drawn from the high-probability region) was longer, relative to the time taken in the previous analysis of a similar stimulus sequence (*please compare the main plot and inset of Figure 4.29*).

Overall responses, for both groups, did not reach a steady-state until ≈ 140ms into the spike train (i.e. 60ms following stimulus onset of the third, lower amplitude). At least, this was the case for the three lowest mean adapting levels (green, blue and red lines and symbols, Figure 4.29 & Figure 4.30); responses to the highest mean adapting condition were relatively flat, with respect to all other conditions and the response elicited in the earlier analysis (see insets, Figure 4.29 & Figure 4.30).

If one compares these decay latencies with that observed for the similar stimulus sequence (see Table 4.1, Table 4.3-4.4, Figure 4.24, Figure 4.27-4.28) one can clearly see the contextual effects of having two strong-amplitudes, from outside the high-probability region presented successively.

For the novelty preferring cells, the time-to decay response, recorded previously, ranged from ≈71-81ms (equivalent to ≈111-121ms in the new analysis) for the non-symmetrical cells to ≈81-101ms (equivalent to ≈121-141ms) for the symmetrical neurons.

For the remainder of the population the previously recorded decay latencies ranged from ≈68-74ms (equivalent to ≈108-114ms in the new analysis)

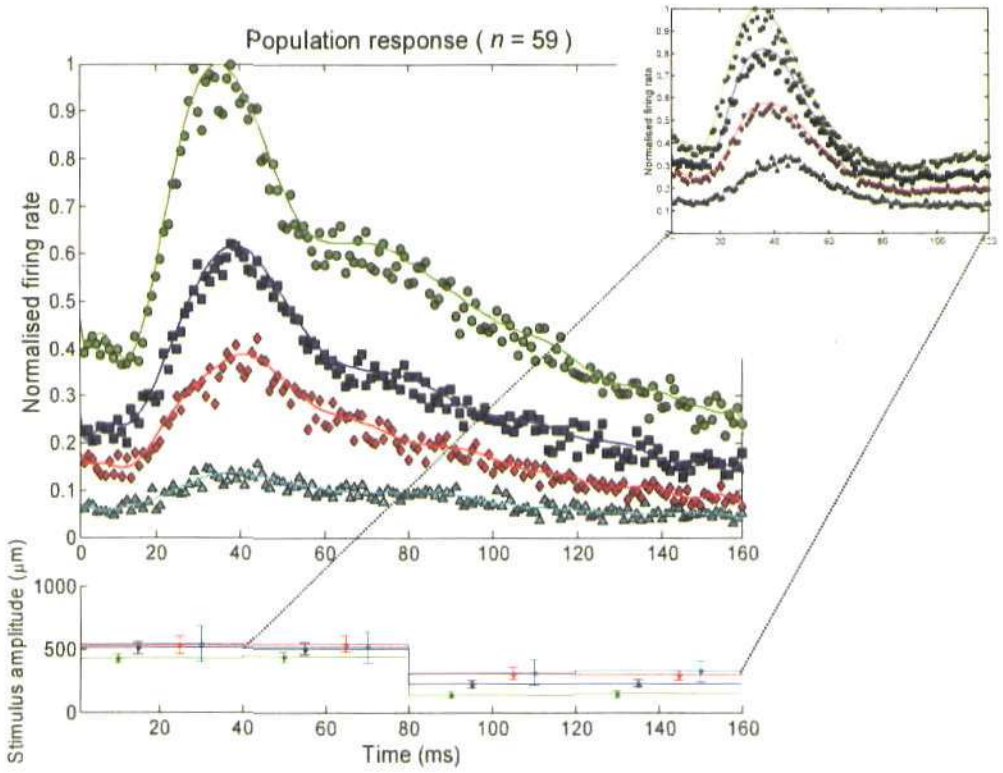

Figure 4.29: Effect of stimulus intensity on the transition response.
Firing rate over time, averaged over all transitions between amplitudes drawn from the high end of the low-probability region. Green circles: 90μm mean adapting level, blue squares: 210μm, red diamonds: 330μm, cyan triangles: 450μm. Colour coordinated solid lines plot the best-fit to the data. Inset is the initial, similar analysis performed using amplitudes drawn from the high end of the low-probability region and amplitudes drawn from the high-probability region, see Figure 4.24

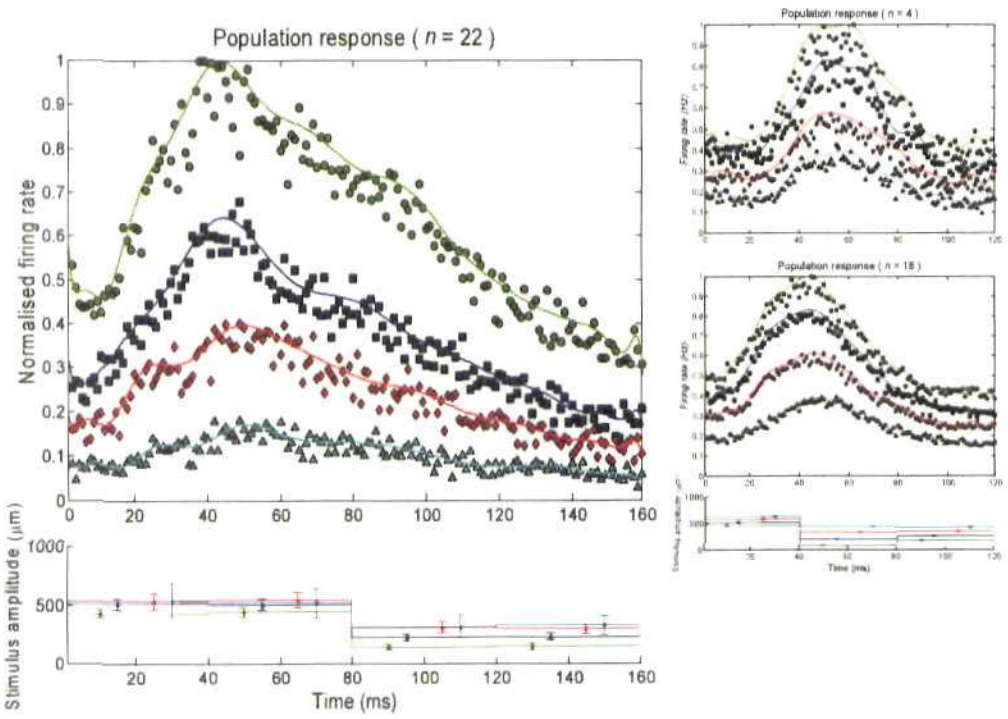

Figure 4.30: Effect of stimulus intensity on the transition response for all cells showing novelty-preference (both symmetric and asymmetric).
Firing rate over time, averaged over all transitions between amplitudes drawn from the high end of the low-probability region. Green lines/circles: 90µm mean adapting level, blue lines/squares: 210µm, red lines/diamonds: 330µm, cyan lines/triangles: 450µm. Colour coordinated solid lines plot the best-fit to the data. Insets are the initial, similar analysis performed using amplitudes drawn from the high end of the low-probability region and amplitudes drawn from the high-probability region for the symmetrical/asymmetrical novelty-type cells; see Figure 4.27 & Figure 4.28.

4.4: Deviation detection

The analysis presented above indicates that barrel cortex neurons are preferentially sensitive to the local, rather than global, stimulus context and respond to how far the stimulus, at any given moment, deviates from the previously presented input. The majority of neurons responded strongly to large step-size changes in stimulus amplitude and could therefore be considered as 'variance detectors'. There is evidence to suggest that V1 neurons respond most strongly to signals that are at least 2 standard deviations away from

the mean input and are therefore acting as deviation detectors (Ringach & Malone 2007, see Chapter 2 for fuller description of the study). It was decided, therefore, to explore the adaptive output of all neurons as a function of **local** stimulus variance, i.e. the squared difference between stimuli presented at each time-bin and the global mean of the stimulus sequence.

The average firing rate (for each high-probability region stimulus, across trials and every presentation within a trial) was calculated, as a function of the difference in whisker-deflection amplitude between the stimulus and the mean input level, for each individual cell and across the population as whole. For every 40ms time-bin (T_{1} / $_{N}$, where N is the total number of stimuli presented during a 30s stimulus-sequence), the spikes elicited (by a given neuron) during T_i were re-sorted into bins according to the squared difference in amplitude between the stimulus (X_i) presented during T_i (given a 20ms latency, see Chapter 3) and the average stimulus presented across the 30s stimulus-sequence, as a whole ($\overline{X}$).

The response to local variance was taken to be the average spike count across all instances in which the variance of the stimulus (the difference between X_i and $\overline{X}$) was equal to $(X - \overline{X})^2$, where X was one of the 25 possible whisker deflection amplitudes (0:720μm, in equal steps of 30μm). On average, the global mean stimulus amplitude (and global variance) for each 30s adapting stimulus sequence was 161μm (variance: 28900μm^2), 252μm (variance: 18496μm^2), 343μm (variance: 13456μm^2) & 433μm (variance: 13689μm^2) for the 90, 210, 330 & 450μm high-probability-region conditions, respectively.

The results of the population analysis, under the condition of adaptation to global-mean, are presented in Figure 4.31, below, with the top plot showing the actual adaptive

response to local variance-from-the-mean and the bottom plot giving the adaptive response as a function of the normalised stimulus (where the input was normalised by dividing $(X_i - \overline{X})^2$ with the inherent global variance of the 30s stimulus-sequence); the abscissa was plotted on a logarithmic scale for ease of viewing.

The responses to the 1st three high-probability-region conditions (90, 210 and 330μm: green lines/circles, blue lines/squares and red line/triangles, respectively, in Figure 4.31) were broadly similar; firing rates tended to remain at roughly 10Hz then began to increase monotonically (at $(X - \overline{X})^2 \approx 10^4 \mu m^2$) up to a maximum firing rate of approximately 30Hz (30.47Hz [±2.2292 Standard Error or SE], 30.60Hz [±2.3758 SE], 29.49Hz [±2.2343 SE] for the 90, 210 and 330μm conditions, respectively). In general responses displayed considerable overlap except at the extremes of the functions, especially with reference to the highest levels of local variance-from-mean, where the functions diverged (this can be seen most clearly in the best-fit to the data, solid lines, Figure 4.31); this is a reflection of the distribution of $(X_i - \overline{X})^2$ values which decreased as the global mean amplitude increased (due to increasing symmetry in low-probability tails of stimulus distribution).

For the 4th condition (cyan lines/diamonds, Figure 4.31), firing-rate responses to levels of $(X_i - \overline{X})^2 < \approx 10^4 \mu m^2$ were fairly constant at approximately 10Hz; for levels of $(X_i - \overline{X})^2$ greater than $10^4 \mu m^2$ the responses oscillated between the base response (of ≈10Hz) and higher firing rates.

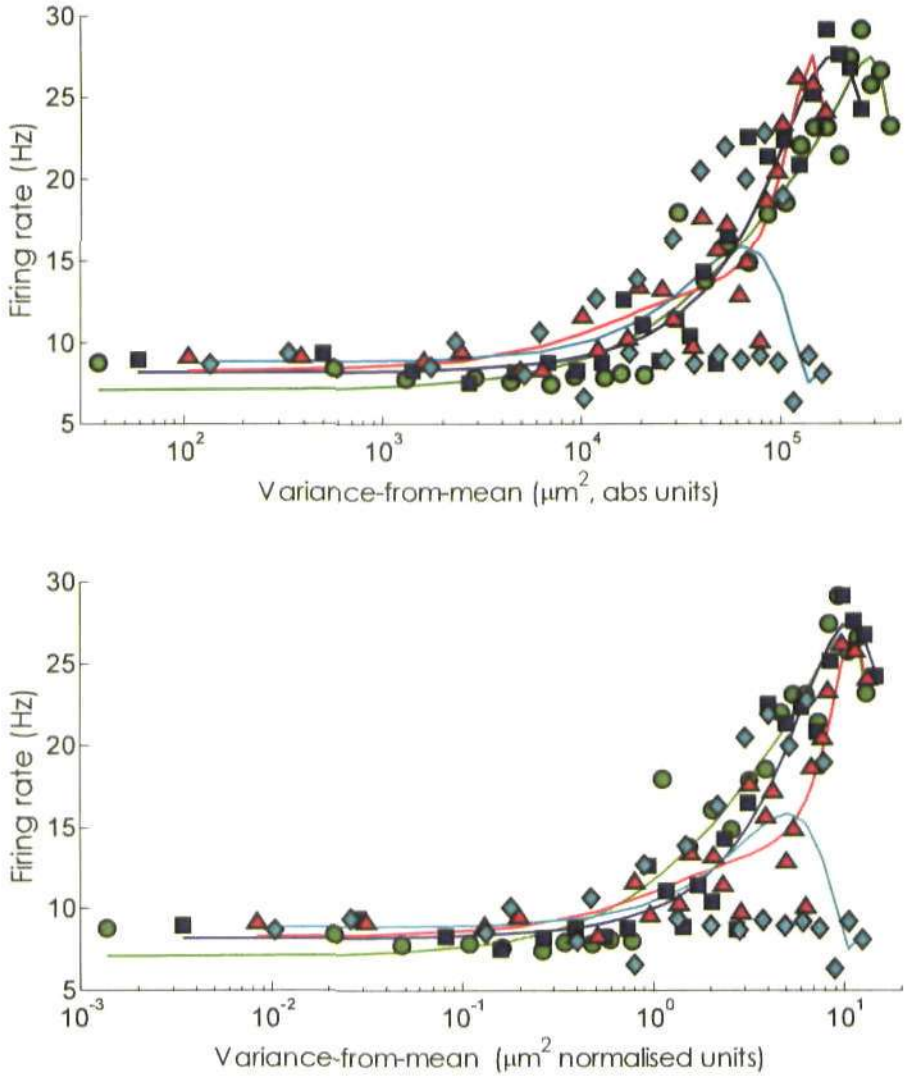

Figure 4.31: Firing rate as a function of local variance-from-the-mean, under conditions of adaptation to global mean whisker deflection amplitude

Top plot: Firing rate as a function of **local**-variance, $(X_i - \overline{X})^2$; plotted on a log-axis for ease of viewing. **Bottom plot**: The functions in the top plot, drawn with local variance normalised, with respect to **global** stimulus variance. Green lines/circles, blue lines/squares, red lines/triangles and cyan lines/diamonds plot the 90, 210, 330 & 450μm high-probability-region conditions, respectively.

The oscillatory behaviour of the response-function could be a reflection of the

balanced stimulus distribution and the stimulus values that set the local variance; the selection probability of amplitude values both below and above the global-mean amplitude had a roughly equal (though low) chance of presentation, relative to the 3 lower global-mean conditions. Firing rates to stimulus amplitudes below global mean were always suppressed relative to higher amplitude stimuli, thus accounting for the apparent oscillation in firing rates. Following normalisation, (achieved by dividing $(X_i - \overline{X})^2$ by the global stimulus variance, Figure 4.31, bottom plot), adaptive neural output, as a function of local variance, did not display invariance; thus suggesting that adaptive responses did not fully scale with stimulus variance.

As the analysis presented earlier in this Chapter has shown, the majority of neurons in the barrel cortex tend to respond preferentially to positive excursions in stimulus amplitude, i.e. are sensitive to the direction of change in the stimulus. Thus it was decided to analyse the data as a function of local standard deviation away from global-mean. The analysis was as for variance-from-the-mean, above, except that deviation was defined as $(X_i - \overline{X})$ rather than $(X_i - \overline{X})^2$.

Responses for all adapting conditions (see Figure 4.32) increased monotonically for local standard deviations of >0μm, and diverged for positive deviations-from-the-mean of > 200μm. This was due to the distribution of the stimulus; deviations ranged from -126.08 to 593.92μm for the 90μm condition (green lines/circles, Figure 4.32), -217.68 to 502.32μm for the 210μm condition (blue lines/squares, Figure 4.32), -310.24 to 409.76μm for the 330μm condition (red lines/triangles, Figure 4.32) and from -401.60 to 318.40μm for the 450μm condition (cyan lines/diamonds, Figure 4.32).

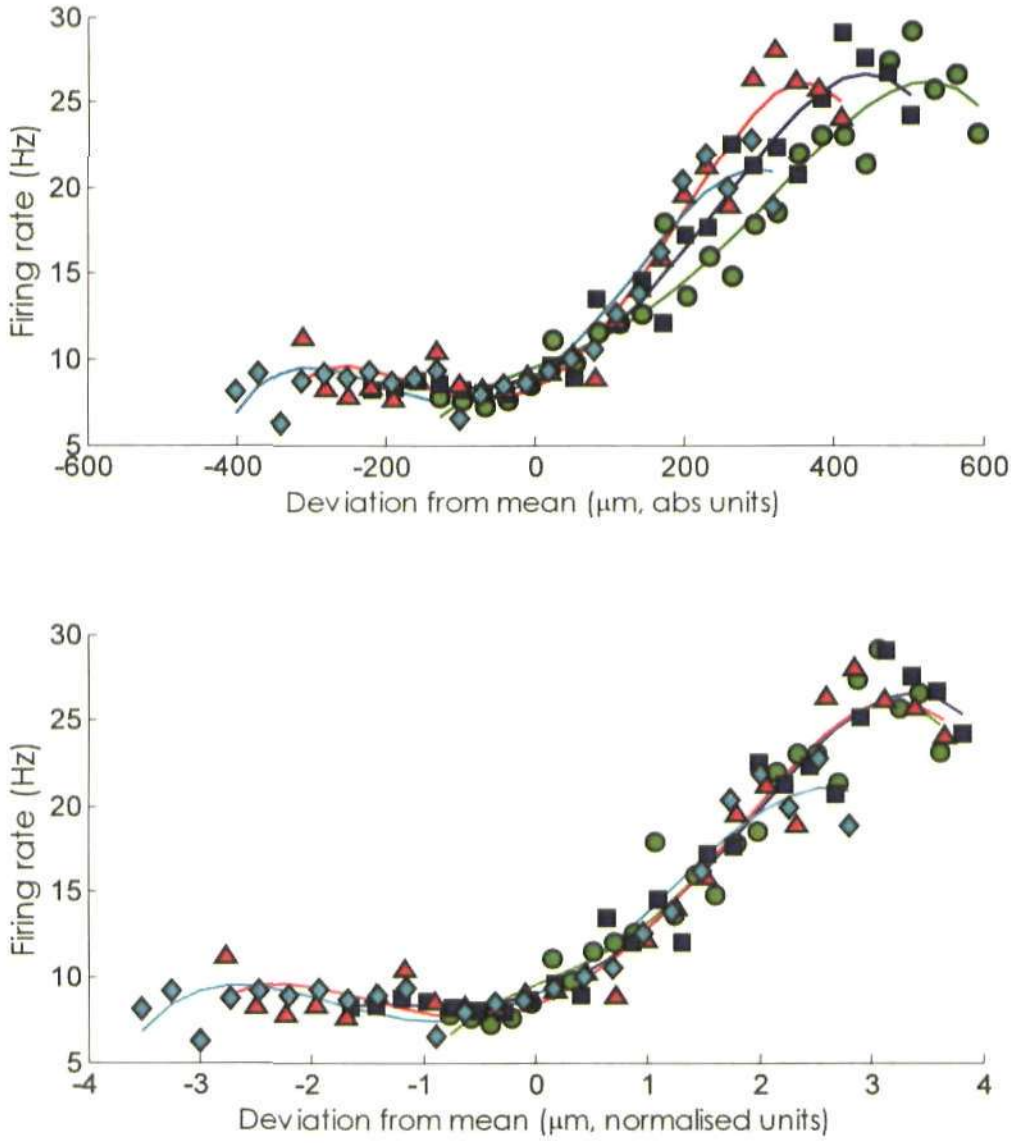

Figure 4.32: Firing rate as function of local deviations-from-the-mean under conditions of adaptation to global-mean.

Top plot: Firing rate as a function of **local**-deviation, $(X_i - \overline{X})$. **Bottom plot**: The functions in the top plot, drawn with local deviation normalised, with respect to standard deviation of the **global** stimulus. Green lines/circles, blue lines/squares, red lines/triangles and cyan lines/diamonds plot the 90, 210, 330 & 450μm high-probability-region conditions, respectively.

However, once inputs were normalised by the global standard deviation (for each 30s stimulus sequence) responses were scaled such that the functions displayed considerable degree of overlap; this was most strongly observed for the three lower global mean adapting conditions.

The response to the standard deviation for those cells displaying novelty preference ($n = 22$; both symmetrical and asymmetrical novelty-preferring cells, see Section 4.3.2) was also similar to that outlined, for the population response, above, except that the responses also increased slightly for deviations <0μm; this can be seen most clearly in the best-fit to the data (solid lines, Figure 4.33), for all conditions, in the actual response (top plot, Figure 4.33) and for the 330 and 450μm high-probability conditions in the normalised response (red and cyan lines; bottom plot, Figure 4.33). Thus these cells clearly display some sensitivity to deviation, irrespective of the direction of change; however, the novelty response as a function of deviation-from-the-mean is much weaker than the response to novelty observed as a function of step-size changes in amplitude (please see Figure 4.12, bottom row).

4.4.1: Adaptation to variance

The data derived under conditions of adaptation to increased levels of global variance was analysed in a similar manner to that outlined above (see Figure 4.34 and Figure 4.35, below).

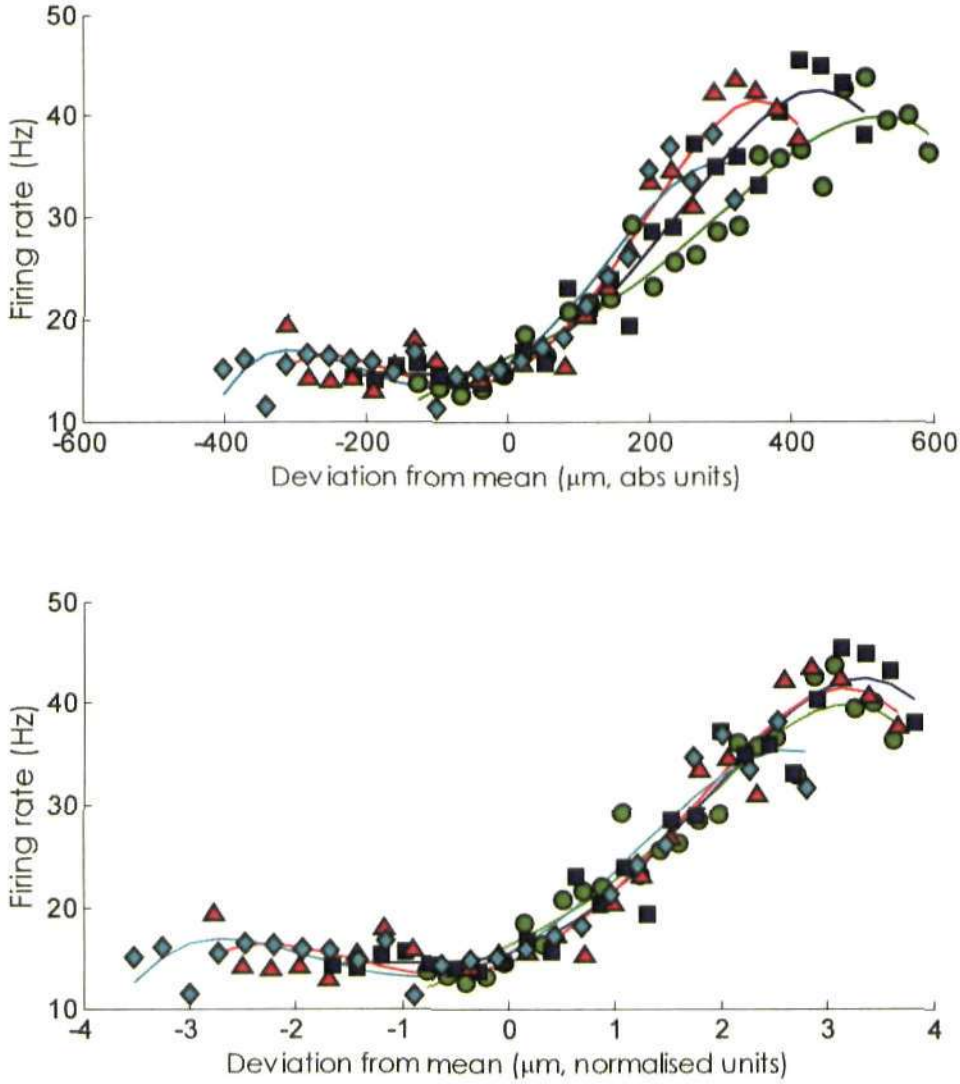

Figure 4.33: Detection of local deviations-from-the-mean for novelty-preferring cells
Top plot: Firing rate as a function of **local**-deviation, $(X_i - \overline{X})$, for all cells ($n = 22$) in the population that displayed novelty-preference to step-size changes in stimulus amplitude. **Bottom plot**: The functions in the top plot, drawn with local-deviation normalised, with respect to standard deviation of the **global** stimulus. Green lines/circles, blue lines/squares, red lines/triangles and cyan lines/diamonds plot the 90, 210, 330 & 450μm high-probability-region conditions, respectively.

In terms of variance-from-the-mean (Figure 4.34; red lines/circles, green lines/squares and black lines/diamonds for the low-, mid- and high-variance conditions, respectively), firing rate, as a function of local variance, showed a general trend to increase from values of $(X_i - \overline{X})^2 > 1.2*10^4$ to a maximum level at $(X_i - \overline{X})^2$ values of $\approx 1.65*10^5$; maximum firing decreased with every increase in global-variance. The shape of the best-fit to the data (solid lines, Figure 4.34, top plot) was generally preserved but tended to be scaled downwards with each successive increase in global-variance. Normalising the input (Figure 4.34, bottom plot), with respect to global-variance, did not result in a perfect scaling up of the response-functions; instead the values of $(X_i - \overline{X})^2$ that elicited maximum firing rate decreased inversely with global-variance.

Under this adaptation regime, the central whisker-deflection amplitude of the high-probability region was fixed, thus the overall symmetry of the stimulus distribution did not change as global-variance increased (in that the reduction in span of the low-probability tails was matched either side of the high-probability region). As one can recall from Chapter 3, increasing global variance resulted in a decrease in the gain of slope of the stimulus-response function and therefore a decrease in the firing rate-response, these two factors (levels of symmetry and reduced firing rates) could account for the preservation and downward scaling of the response, as a function of local-variance.

Increasing the width of the high-probability region had the affect of increasing global-mean slightly, (272 [global variance: 14,640μm^2], 279 [19,600μm^2] and 290μm [269,001μm^2] for the low-, mid- and high-variance conditions, respectively), thus the distribution of values of $(X_i - \overline{X})^2$ decreased in range with each successive increase in

global variance (from 73984 to 200704μm^2, 77841 to 194481μm^2 and 84100 to 184900μm^2 for the low-, mid- and high-variance conditions, respectively); this reduction in range can be observed in the normalised response (Figure 4.34, bottom plot), where local-variance levels that elicited maximum firing rates were reduced.

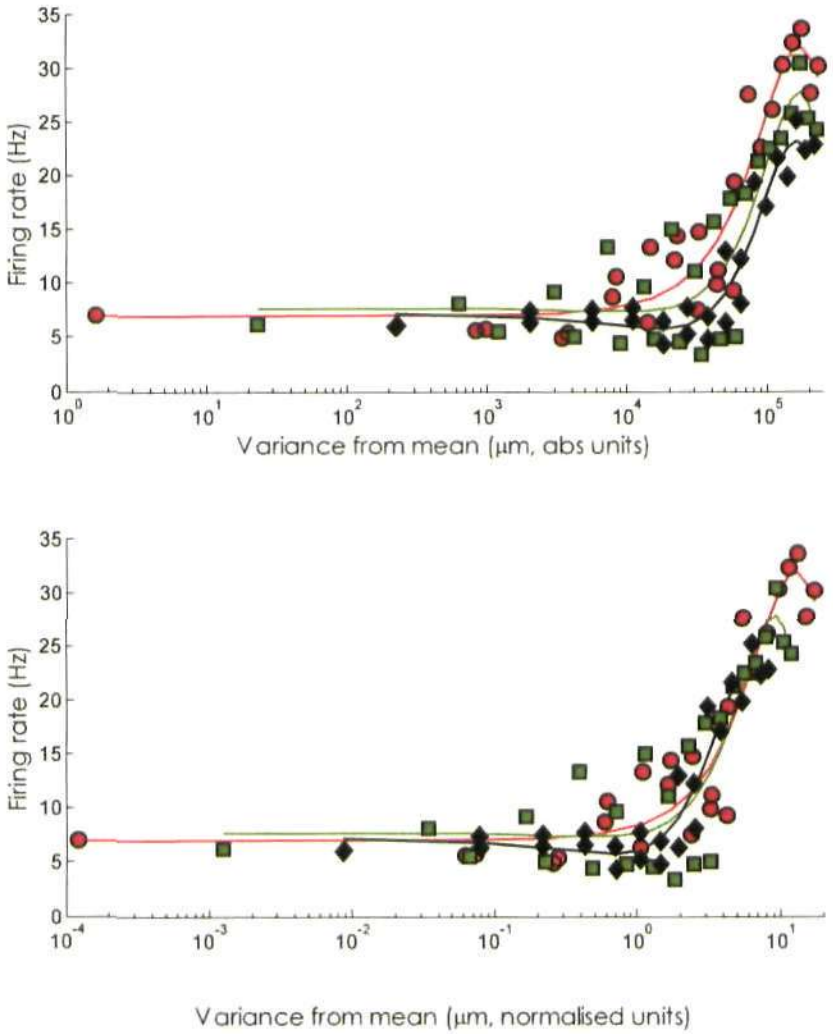

Figure 4.34: local variance under condition of adaptation to global variance
Top plot: Firing rate as a function of **local**-variance, $(X_i - \bar{X})^2$; plotted on a log-axis for ease of viewing. **Bottom plot**: The functions in the top plot, drawn with local-variance normalised, with respect to **global** stimulus-variance. Red lines/circles, green lines/squares, black lines/diamonds plot the low-, mid-& high-variance conditions, respectively.

The relationship between firing rate and local-variance displayed considerable variability within each adapting condition. It was therefore decided to examine deviation-from-the-mean (Figure 4.35, below).

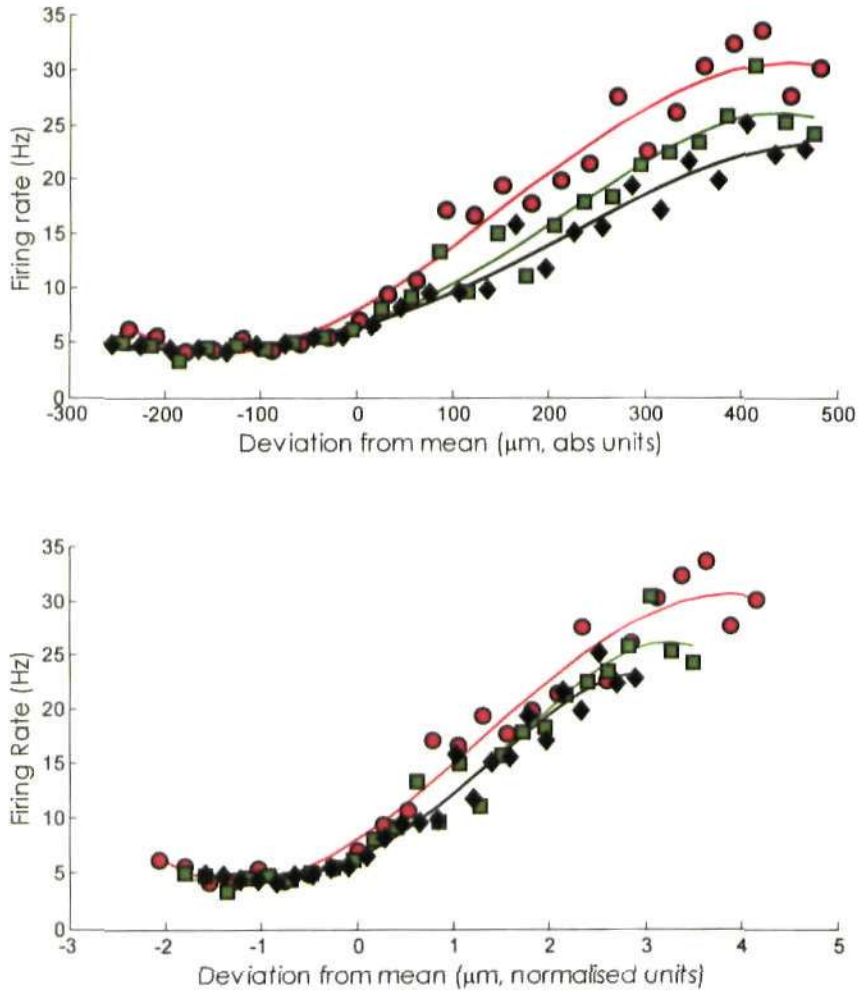

Figure 4.35: Response to deviation-from-mean under conditions of adaptation to global variance

Top plot: Firing rate as a function of **local**-deviation, $(X_i - \overline{X})$. **Bottom plot**: The functions in the top plot, drawn with local-deviation normalised, with respect to the standard deviation of the **global** stimulus. Red lines/circles, green lines/squares, black lines/diamonds plot the low-, mid-& high-variance conditions, respectively.

The relationship between output and local deviation (Figure 4.35) was much more clearly defined than that observed for local variance (see Figure 4.34, above). There was a clear preference across the population for positive deviations away from global mean. The response functions tended to decrease in slope-gain and maximum firing rate levels in an inverse manner to stimulus variance. Normalising the firing-rate functions resulted in a considerable degree of overlap for the mid- to high-variance conditions (green and black lines/symbols, respectively; Figure 4.35, bottom plot), except for at the positive-deviation tail of the functions. Under the low-variance condition (red line/circles; Figure 4.35, bottom plot), there was a clear divergence in the firing rate as a function of deviation-from-the-mean from the two-higher variance conditions. The decrease in the slope of the function from the low- to mid-variance was greater than the decrease between the mid- to high-variance conditions; this could be a reflection of the sharp increase in the width of the high-probability region from a span of three amplitudes to seven, in the change in stimulus configuration from the low- to mid-variance conditions.

4.5: Discussion

There are several key findings from the above analysis:

1. Across the population, the step-size function was roughly invariant, with respect to the global mean of the adapting stimulus

2. However, for those cells adapted only to different levels of global stimulus variance, the step-size function was not invariant.

3. Amongst both populations (neurons adapted to either mean amplitude or adapted to variance only), some cells showed a preference for step-sizes away from 0μm, irrespective of the direction of amplitude change and were classed as novelty-preferring cells.

4. The majority of novelty-preferring cells had a response bias for positive step-changes; however a small proportion of these neurons displayed symmetrical functions.

5. Novelty-preferring cells that displayed symmetry had longer time-to-peak and response-decay latencies than the remainder of the population.

6. Novelty-preferring cells (with symmetry) that were only adapted to stimulus variance had step-size functions that showed some dependence on the width of high-probability region in the stimulus space.

7. As well as being sensitive to local stimulus history, all neurons displayed sensitivity to local deviations from global mean, with a preference for positive deviations.

8. Under conditions of adaptation to global mean, responses, as a function of deviation-from-global-mean, showed scaling, when plotted against normalised input levels.

9. Under conditions of adaptation to global variance, normalised deviation-from-the – mean functions were not invariant, suggesting that the decrease in gain, observed as adapting stimulus-variance increased, were not scaled as a function of local deviation.

The presence of these novelty-preferring cells leads one instantly to the question of whether these are a distinct cell type, in terms of their responsiveness at least (neither morphology nor electrophysiology can addressed here). The clearest difference between the two groups, of novelty-preferring neurons and cells which respond only to positive step-sizes, can be found in the response-decay latency. The influence of response decay is most apparent during negative amplitude transitions, where the preceding stimulus is always greater than the stimulating amplitude. The steady-state response (as measured by counting all spikes within 20-60ms of stimulus onset), if one is present, is thus overshadowed by the decomposition of the preceding amplitude response. From this one could argue that the

sensitivity to negative, as well as positive, step-sizes is simply the result of an artificial inflation of the response. However, this could well be the mechanism by which the cell 'detects' novelty. The speed and shape of the decay response is heavily influenced, not just by the immediate stimulus context, but by the stimuli presented 80ms prior to the current epoch (see Figure 4.29 & Figure 4.30). Of course, from the spike-triggered average data (Figure 4.8) it could be argued that interactions up to 320ms preceding the current stimulus epoch can influence firing rate. Thus a cell's responsiveness to stimulus context obviously extends well beyond the immediate changes in amplitude levels.

Unfortunately, from the present data, it is not possible to state, unequivocally, whether these neurons represent a distinct population of cells that encode step-size differences through adjustments in time-to-decay latency, or whether novelty-preference is an artefact of the longer response-decay latencies displayed by this group of cells.

The unifying factor between all the groups of cells identified, however, is the invariance of the response to step-size, as a function of the global mean.

The invariance of the step-size function, following adaptation to the mean, suggests that the lateral displacement of the stimulus-response function acts to preserve the relationship between firing-rate and step-size and not necessarily to encode the stimulus distribution as a whole (as predicted by Information theory and the efficient-coding hypothesis). Thus one can argue that adaptation serves to encode relative changes in intensity, with a temporal resolution on the order of tens of milliseconds.

For the general population, and across mean-adaptation levels, the highest firing-rates were found only for those positive step-sizes in which the stimulating amplitude was drawn from the high-amplitude, low-probability tail of the stimulus distribution (see Figure

4.14, bottom row). Thus the relationship between step-size and firing-rate was maintained by the lateral displacement of the stimulus-response function; as the lowest firing rates were always found for amplitudes within and below the high-probability region.

In the absence of any adaptive shift in stimulus-response function, one would observe, with each successive increase in the high-probability region level, a firing-rate increase for step-sizes of 0μm to ±120μm, as if one recalls, step-sizes ranging from 0μm to ±120μm are most likely to be derived from interactions between amplitudes within the high-probability region (Figure 4.14, top row).

The amount of stimulus-response function displacement must also play a role in maintaining the step-size function. If one looks closely at the population stimulus-response function (SRF, Figure 4.1, above), one can see that the stimulus amplitude that elicits the maximum population response (≈600μm) is fairly constant across all adapting conditions. Thus the stimulus-response function is displaced such that the dynamic region of the response always covers the stimulus-space that spans from the high-probability region until ≈600μm; stimulus amplitudes of ±600μm represent the extent of the firing-rate increase seen in the population step-size function.

If one recalls from Chapter 3, the maximum response of the stimulus-response function was reduced for each successive increase in the global mean adapting level (see Chapter 3). The average step-size difference under each adapting condition also decreased as the mean adapting amplitude increased. The number of possible stimuli located within the high-end tail of the low-probability stimulus space decreases as the centre of the high-probability region shifts to higher amplitudes. For example, the average step-size difference for transitions between the high-amplitude, low-probability stimuli and amplitudes within

the high-probability region decreased with the adapting mean level (see Figure 4.24, Figure 4.27 & Figure 4.28 and Table 4.5). Thus the reduction in maximum response in the stimulus-response function can also be accounted for by the relationship between firing-rate and step-size.

Table 4.5: Average step-size (in μm) under each mean adapting level, for transitions between amplitudes drawn from high-end of the low probability tail (LPR-high) and the high-probability region (HPR)

	90μm	210μm	330μm	450μm
LPR (high) to HPR	365.46 ± 20.14	302.07 ± 17.27	239.70 ± 15.64	179.18 ± 13.53

Whilst there is a strong case for stating that the displacement of the stimulus-response function acts to maintain the relationship between output and step-size, the same argument does not apply to the decrease in neural gain, observed following adaptation to increased levels of stimulus variance. The step-size function was not scaled with respect to the levels of stimulus variance; thus giving further credence to the notion that adaptation to global variance and global mean are underpinned by separate mechanisms.

Of course, the population step-size function did show some dependence on global variance levels, the slope of the function tended to decrease, under conditions of increased variance, arguably as a reflection of the increased range of interactions between stimuli within the high-probability region. As such, neurons may not be encoding novelty (in the strictest sense) through their step-size functions, but rather the range of diversity in highly probable stimulus amplitude interactions.

Overall, the clear result from this analysis is that local stimulus context plays a major role in determining the output of a cell. This level of interaction has been found

across cortical systems and species, for example in the fly visual system (Fairhall et al. 2001), and more recently in the barrel cortex (Maravall et al. 2007). Whiskers were deflected along a single dimension, with the position randomly assigned using a Gaussian distribution and refreshed every 5ms; the variance of the distribution was varied every 5s in a switching stimulus paradigm. Thus barrel cortex cells were adapted to two levels of variance only (high and low), with the majority of cells displaying levels of adaptation to whisker-deflection velocity, in proportion to stimulus variance.

As stated earlier (see Section: 4.3.1) it could be argued that plotting the response as a function of the difference in amplitude, from one time-step to the next, is a form of velocity measure, in that one is describing the relationship between neural output and the direction and distance of change in intensity levels. Thus, if one takes the argument that the step-size function is comparable to a velocity-response function, there is clearly a contradiction between the cellular response to stimulus velocity (Maravall et al. 2007) and the adaptive response to stimulus intensity, as the adaptive response to increased levels of stimulus variance did not scale.

However, an adaptive rescaling, such that responses were roughly invariant, was only observed following adaptation to different levels of the stimulus mean. As mentioned, within the study exploring adaptation to whisker deflection velocity (Maravall et al. 2007) the stimulus mean and variance were co-dependent, thus the suggestion here is that the rescaling observed in response to increased variance, in the velocity feature of the stimulus, may have been influenced by levels of stimulus mean rather than only global variance.

Overall, it would appear that adaptation acts to rescale responses so as to maintain information transmission, about the local changes in stimulus amplitude, across the

different global mean levels, encountered during adaptation process. For the majority of neurons within the population, most showed distinct response bias for large, positive step-sizes. As we have seen, the majority of these step-size responses are elicited to strong amplitudes, generally located outside the high-probability region. Thus these cells can be said to be responding predominately to large deviations in the input signal, essentially in a manner similar to that observed in the visual cortex (Ringach & Malone 2007).

The analysis presented within this Chapter demonstrates that barrel cortex neurons are sensitive to positive deviations away from global mean on the 40ms scale. Ringach and Malone (2007) suggested that deviation detection sets the operating point of the visual cortex, such that spiking threshold is, on average, one standard deviation from the mean of any ongoing activity, irrespective of contrast levels; thereby enhancing the detection of signals in background noise over maximising information transmission

It would appear that barrel cortex neurons, under conditions of adaptation to global mean, are also adapting to maintain response sensitivity to positive deviations from the mean of the input. However, it should be noted that neurons show a base firing-rate level to all standard deviations (normalised units) which increases monotonically from deviations greater than 0μm; thus the threshold for firing-rates to exceed the base-level is lower than the spike-threshold observed in V1 (Ringach and Malone 2007) and close to global-mean. This suggests that the operating regime serves to enhance detection of all signals that exceed the average intensity of the input and that adaptation to global mean is underpinned by a mechanism of asymmetric deviation detection, or background suppression, in which signals of lower intensity than global-mean are suppressed.

The same degree of invariance in the deviation-response functions was not observed

following adaptation to global variance, even though the general response-trends were similar under both adaptation regimes (i.e. responses increased monotonically from a base-line level for deviations greater than the threshold, which was roughly equal to the mean of the stimulus distribution). This again lends credence to the notion that adaptation to variance may be underpinned by a different mechanism, such that when the signal becomes increasingly noisy it may be optimal to encode the stimulus distribution and maximise information transmission.

In conclusion, the experimental adapting regime was designed to mimic natural conditions, more closely, in that stimulus amplitude fluctuated rapidly about a mean level and rate functions were calculated on the fly. Many adaptation experiments tend to present static stimuli and assess responsiveness following stimulus offset. If these findings are applicable to a real-world setting, then this suggests a change in coding strategy as the stimulus environment becomes 'noisier'. It may be that under normal conditions cells are primed to detect novel/deviant stimuli of stronger intensity than the global mean, however in a noisy environment it may be more functionally advantageous to have cells primed for detecting the overall distribution of interactions between stimulus intensity.

Chapter 5: Mechanisms of adaptation: a modelling study

5.1: Introduction

The previous two chapters have shown that adaptation to the global variance and mean of an adapting stimulus are distinct processes, which underpin computation not only of the stimulus distribution but relative, short term changes in stimulus amplitude too.

A useful question to pose at this juncture, then, is what are the mechanisms that underpin both forms of adaptation? In order to address this issue, computer simulations of various models of gain control were developed and compared to the experimental data. Simulations were performed using an implementation of the conductance-based, integrate and fire cell developed by Alain Destexhe (Destexhe 1996).

It was found that whilst spike-rate adaptation could generate the lateral shift associated with adaptation to global mean (as suggested by a recent modelling study (Garcia-Lazaro et al. 2007) and other experimental work in the visual cortex (e.g. Sanchez-Vives et al. 2000) it did not elicit the expected, invariant, tuning of the model cell to local changes in stimulus amplitude (see Chapter 4). In order to generate the appropriate step-size functions, a neural model with spike-rate adaptation was augmented through the inclusion of a tonic inhibitory conductance (Murphy & Miller 2003), whose strength was determined by the global mean adapting amplitude.

The mechanism underpinning adaptation to variance is more elusive. Several models have been developed that rely on extrinsic mechanisms to generate gain modulation, but only to steady state stimuli; thus none have directly addressed adaptation to stimulus statistics.

The key feature of all the models reviewed was that neural gain (i.e. the slope of the stimulus-response function) was dependent upon the relative balance of inhibition and excitation impinging upon the target cell, either in terms of modulatory background firing rates (Chance et al. 2002) or afferent activity (e.g. Rothman, Cathala, Steube, & Silver 2009, Murphy & Miller 2003 & Chelaru & Dragoi 2008). As such two models were developed, one based on background firing rates (as proposed by Frances Chance [Chance et al. 2002]) and the other based primarily on a recent model of asymmetric synaptic depression (Chelaru & Dragoi 2008) within V1, in which the decrease in neural gain, following adaptation to variance, arose through a change in the relative levels of afferent inhibition and excitation.

5.1.1: Spike frequency adaptation

In order to explore the intrinsic mechanisms that underpinned the adaptive displacement of the stimulus-response function in response to increased stimulus mean, Jan Schnupp and colleagues (Garcia-Lazaro et al. 2007) built realistic neural models of four reconstructed cells (taken from the model NEURON database, Mainen & Sejnowski 1996), using the NEURON programming language (Hines & Carnvale 1997). The cells were adapted through the simulated injection of a continuous current, whose amplitude was refreshed every 40ms and drawn from a highly skewed distribution (see Figure 5.1, below), as per the experimental stimulus outlined in Chapter 3.

Of the four model cells only two (the layer 3 pyramid cell model, Figure 5.1B, and the layer 4 stellate cell model, Figure 5.1D) displayed an adaptive lateral displacement of the stimulus-response function, in response to increasing the global mean of the stimulus.

The adaptive shift, in the stimulus-response function of the layer 3 pyramid cell model, was abolished (see Figure 5.2A and B, below) when the conductance of the calcium-dependent, potassium based current, gK_{Ca}, was reduced to below $0.5pS/\mu m^2$. Increasing gK_{Ca} above the base level of $3pS/\mu m^2$ (at which the responses plotted in Figure 5.1, were generated, see Figure 5.2C), heightened the degree of rate-function displacement, especially for the final high-probability-region tested (Figure 5.2C & D, cyan lines), but the level of displacement was not as pronounced as that observed when gK_{Ca} was increased from 0.5-3 $pS/\mu m^2$.

Thus, there is a strong indication that gK_{Ca} is the source of adaptation to mean stimulus amplitude, in the model neurons at least, but that the amount of lateral shift in the stimulus-response function is roughly independent of the strength of the current (as increasing gK_{Ca} doesn't lead to a proportional shift in neural firing threshold). This is in line with the experimental data as, if one recalls, the population stimulus-response function (stimulus-response function) for each high-probability region was displaced to lie outside of the high-probability region but the maximum of the curve was centred at $\approx 600\mu m$, irrespective of the global mean of the stimulus (see Figure 4.1, Chapter 4).

The gK_{Ca} generates an after-hyperpolarisation, AHP, in the membrane potential, post spike generation. As discussed in Chapter 2, the AHP is believed to underpin spike frequency adaptation, whereby, under constant, continuous stimulation, the inter-spike interval, ISI, increases (e.g. Madison & Nicoll 1984).

Spike frequency adaptation is most commonly associated with regular spiking cells (Connors & Gutnick 1990) and, interestingly, one of the model cells that displayed adaptation was a reconstructed regular-spiking, layer-4 spiny cell (Figure 5.1D).

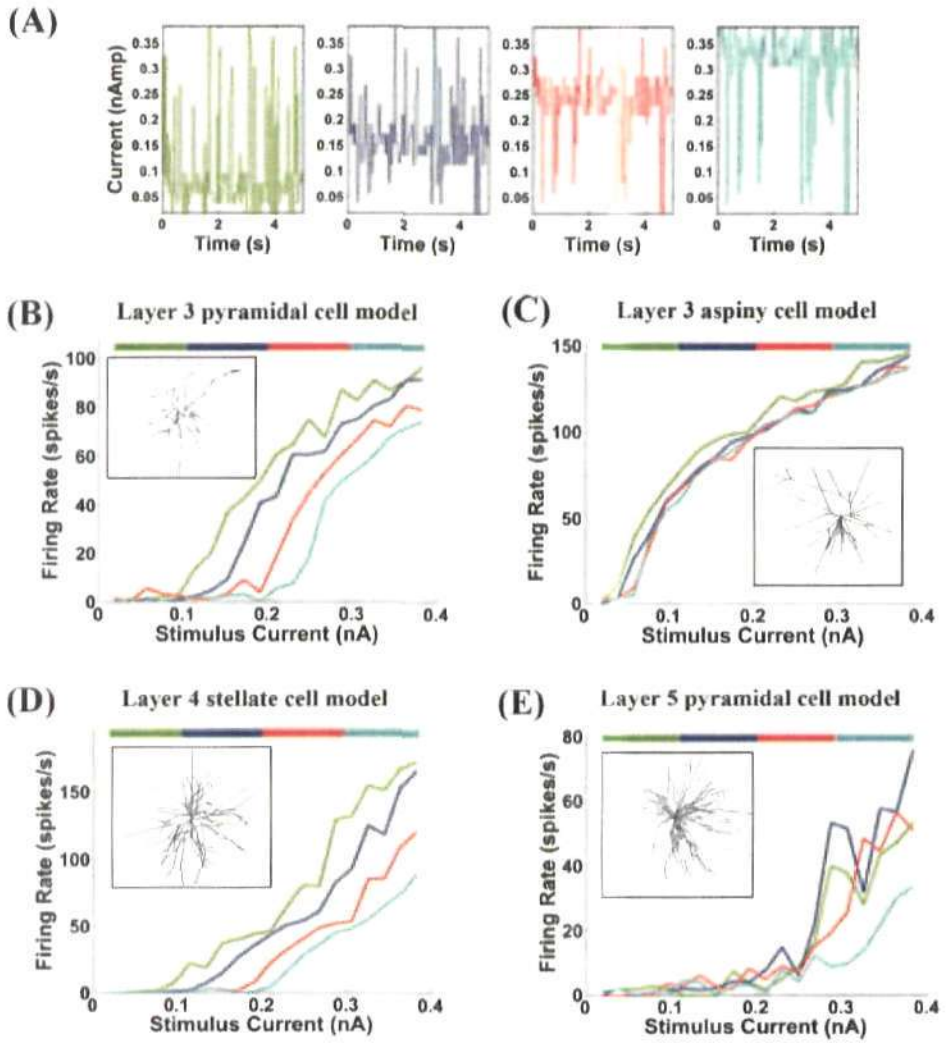

Figure 5.1: Adaptation to global mean in four reconstructed model neurons
A: 5s sample of the adapting stimulus; from left to right the high-probability regions (high-probability region) spanned 0.019 to 0.096nA (green), 0.115 to 0.191nA (blue), 0.210 to 0.287nA (red), and 0.306 to 0.383nA (cyan). **B:** Reconstructed NEURON model of a layer 3 pyramidal cell. **C:** Layer 3 aspiny cell. **D:** A layer 4 stellate cell. **E:** Layer 5 pyramidal cell model. Only B and D displayed adaptation. The insets in B to E illustrate the morphologies of each reconstructed cell. Taken from Garcia-Lazaro et al. (2007)

However, AHP is not restricted to this class of cell and can influence spike or burst generation in non-regular spiking type cells; the other model that displayed adaptation was the burst firing, Layer 3, pyramidal cell (Figure 5.1B).

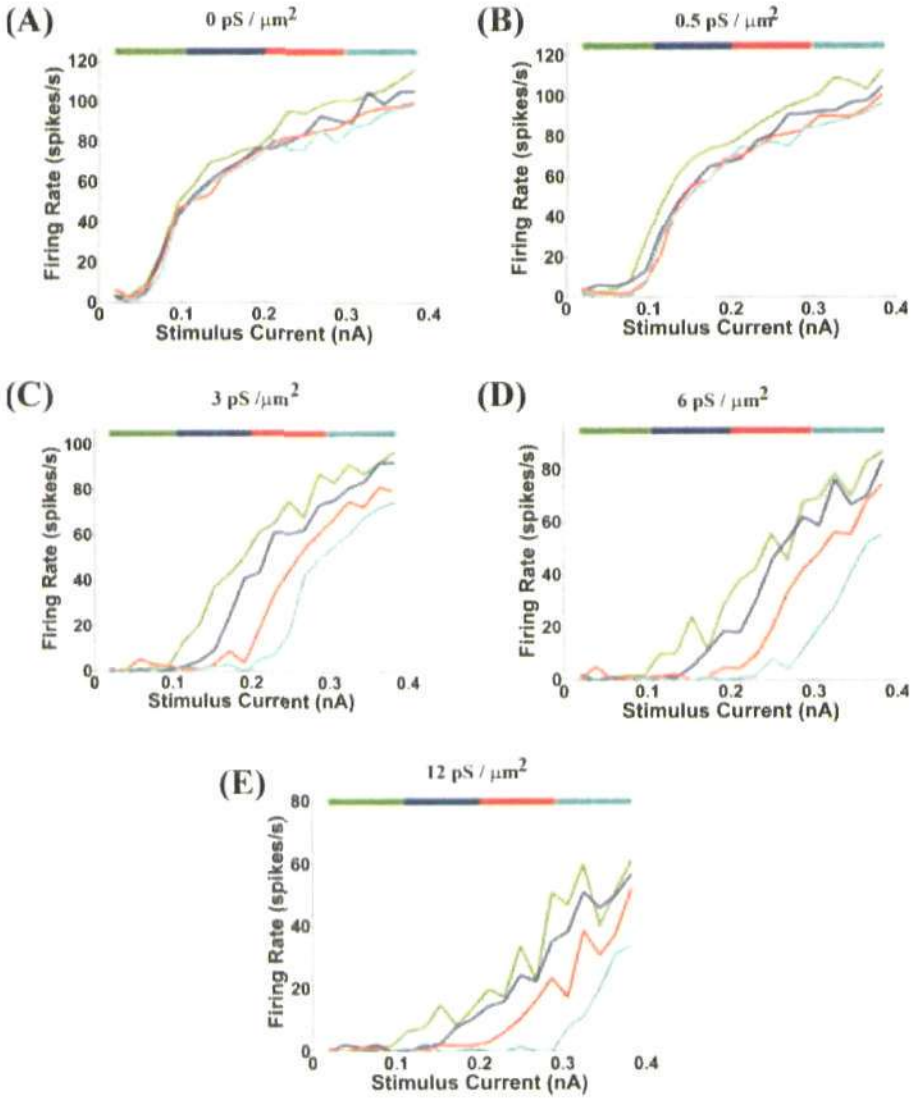

Figure 5.2: Spike-rate adaptation in a model neuron
Adaptation in the layer 3 pyramidal cell model was dependent on the strength of the gK_{Ca}. **A:** gK_{Ca} = 0pS/μm². **B:** 0.5pS/μm². **C:** 3pS/μm². **D:** 6pS/μm². **E:** 12pS/μm². Taken from (Garcia-Lazaro et al. 2007)

5.1.2: Adaptation to variance

As discussed, two out of the four model cells tested displayed adaptation to global mean; however neither of these two model neurons adapted to increases in the variance of the stimulus (see Figure 5.3). Thus, it can be argued, from the results of the modelling

study above at least, that the AHP current does not generate adaptation to variance.

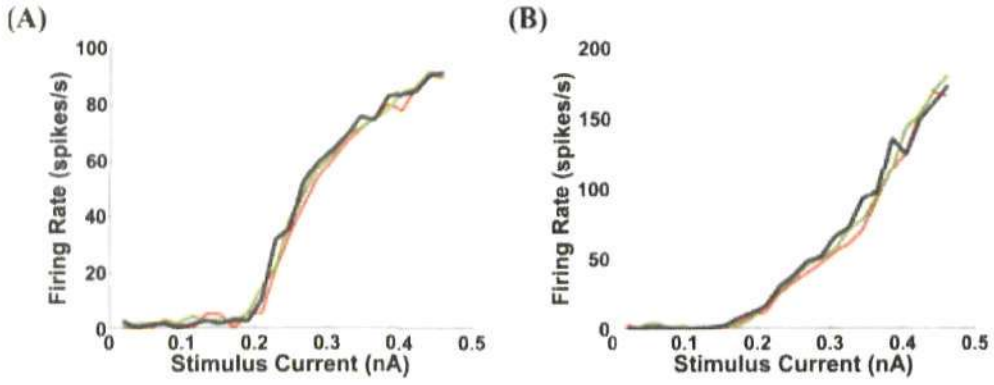

Figure 5.3: Adaptation to variance in two reconstructed neural models that displayed adaptation to mean
Stimulus-response functions, in response to adaptation with stimuli of increasing variance **A:** The layer 3 pyramidal cell. **B:** The layer 4 stellate cell (see Figure 5.1 B and D, respectively). Taken from Garcia-Lazaro et al. (2007)

In fact some of the most compelling evidence to date suggests that gain modulation, i.e. changes in neural output that affect the slope of the stimulus-response function, is dependent, to a lesser or greater degree, on the interplay between excitation and inhibition, across the network as a whole. As such, the focus of this chapter will be on extrinsic mechanisms, namely:

- Levels of background noise (Chance et al. 2002)
- A combination of the levels of background inhibition and synaptic depression (Murphy & Miller 2003; Rothman et al. 2009)
- The overall balance of synaptic depression acting on excitatory and inhibitory inputs (Chelaru & Dragoi 2008)

Gain modulation from background synaptic input

In vivo, neurons are subjected to continuous barrages of either driving (i.e. afferent excitation) or noisy synaptic input. Stochastic inputs can be both excitatory and inhibitory and are generally considered as background noise, in that whilst the level of noise affects response variability it does not, strictly, drive the cell to respond. However it has been demonstrated, *in vitro*, that, within rat somatosensory cortex, the amount of background noise can influence neural-gain (Chance et al. 2002). The details of the experiment are reviewed briefly below.

Using a dynamic clamp technique, in a slice preparation, it was found that the gain of target neurons decreased, relative to a control condition (see Figure 5.4) when firing rates of both background excitation and inhibition were increased by the same amount; gain modulation was quantified as a change in the slope of the f-I curve (firing rate verses driving current). Excitatory and inhibitory pre-synaptic spikes were generated using inhomogeneous Poisson processes, with basic firing rates of 7000Hz and 3000Hz (to represent the summed input of many synapses) for excitation and inhibition, respectively. For each presynaptic action potential (AP) a unitary conductance was added to running tally of conductances. Conductances were calculated using a difference-of-exponentials equation and were introduced to cell *in vitro* as synaptic currents, using the dynamic clamp method, along with a constant driving current. Both inhibitory and excitatory background conductances were balanced, such that their net synaptic current was kept approximately at zero. By increasing the firing rates of both inhibition and excitation in a balanced manner (e.g. from the control condition of 7000 & 3000Hz [1X condition], to 14000Hz & 6000Hz [2X] for excitation and inhibition, respectively) there was an increase only in input noise and not net synaptic current. The increased input noise resulted in a decrease in neural gain

(see Figure 5.4).

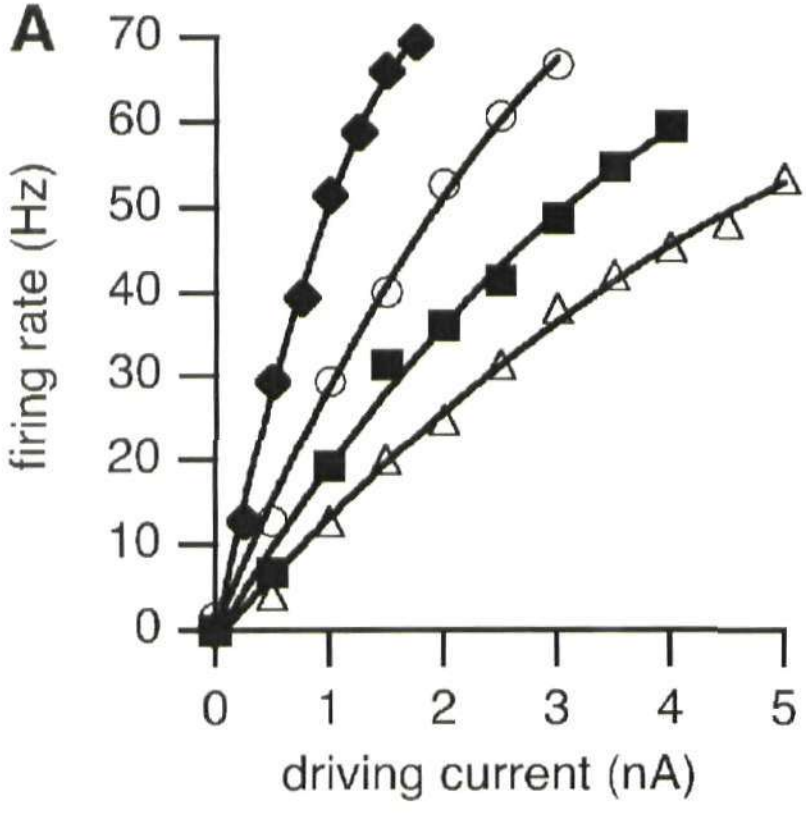

Figure 5.4: Gain modulation in rat somatosensory cortex, in vitro
Increasing background input (by X amount) resulted in a decrease in the neural gain of the neuron tested. Diamonds: 0X condition; circles: 1X; squares: 2X; triangles: 3X. Taken from Chance et al. (2002)

Inhibition-mediated gain change

The background noise model of gain modulation (Chance et al. (2002), see section above) assumes that neural gain is dependent on stochastic, background firing rates, however an alternative approach is to assume that gain-control is instead mediated via tonic inhibition. A modelling study by Murphy & Miller (Murphy & Miller 2003) explored the relationship between the introduction of constant, background synaptic input and gain

modulation of a model V1 cell. The study concluded that it was possible to elicit multiplicative scaling in a model cell (i.e. an increase or decrease in the slope of the neural-response function such that the rate-function could be scaled up or down to fit the base-level response, see Figure 5.5), following the introduction of either a tonic excitatory or inhibitory conductance into the model. However, gain modulation was only observed when the relationship between the external stimulus parameter and input current was non-linear (e.g. if the relationship between stimulus and driving current was sigmoidal, as per the contrast-response observed in the visual cortex) and there was also a non-linear relationship between membrane potential and firing rate (e.g. a spike-threshold mechanism).

The modelling study simulated the iontophoretic application of pharmacological agents, which increased either modulatory excitation or inhibition, by introducing a constant conductance for the relevant receptor. Modulatory input was small compared to driving current but sufficient to generate a change in gain, when the other non-linearities were present. The simulated introduction of a constant conductance resulted in multiplicative scaling of the model cell's contrast-response function (see Figure 5.5A&B); with the direction of gain change determined by whether the conductance was excitatory (i.e. modulated by either α-amino-3-hydroxyl-5-methyl-4-isoxazole-propionate (AMPA) or N-methyl-D-aspartic acid (NMDA) conductance) or inhibitory (gamma-aminobutyric acid-A or B ($GABA_A$ or $GABA_B$) conductance). The introduction of GABA-mediated conductances resulted in a decrease in gain, with the effect most pronounced for the simulated $GABA_B$ conductance.

A potential source of non-linearity, within sensory cortices in general, is synaptic depression. As discussed in Chapter 2, the evidence that thalamocortical synaptic

depression underpins some of the more high-level features of contrast-gain control in the visual cortex, such as cross-orientation suppression and contrast-invariance of the orientation tuning curve, is mixed (e.g. Carandini, Heeger & Senn 2002; Freeman, Durand, Kiper & Carandini 2002; Boudreau & Ferster 2005, Li, Thompson, Duong, Peterson & Freeman 2006).

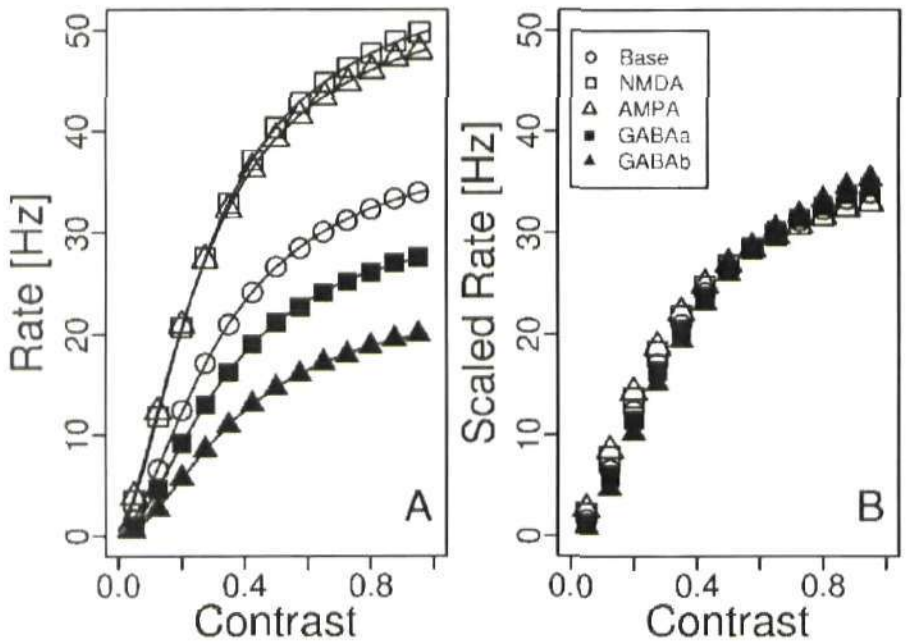

Figure 5.5: Tonic inhibition/excitation mediated gain modulation in a model cell
A: Plot of contrast versus firing rate for a model neuron under conditions of simulated iontophoretic application of NMDA (open squares), AMPA (open triangles), $GABA_A$ (closed squares), or $GABA_B$ (closed triangles); control condition was represented by open circles. Solid lines are fits of the data to a hyperbolic ratio function. **B:** Curves from A scaled to optimally (least squares) fit the Base curve. Taken from Murphy & Miller (2003)

However, in terms of simple features, such as saturation of the neural response function, synaptic depression remains a strong candidate mechanism (Adorjan et al., 1999, Kayser, Preibe & Miller 2001); as when presynaptic firing rates become too fast, to allow for a full recovery from depression between spiking inputs, post-synaptic responses tend to saturate (Abbott et al., 1997; Tsodyks & Markram 1997).

Recently, an experimental (*in vitro*, see Figure 5.6, Top row) and modelling study (Figure 5.6, Bottom row) has found that, within the cerebellum at least, synaptic depression enhances and alters inhibition-mediated gain-control (Rothman et al. 2009); gain is multiplicative (i.e. affects slope of the stimulus response function only) in the presence of synaptic depression and additive (affects stimulus-response function threshold) in the absence of depression.

Symmetry of excitatory and inhibitory synaptic depression

Synaptic depression is not restricted solely to excitatory inputs; in fact shifting the balance between levels of excitatory and inhibitory synaptic depression can be a powerful mechanism for controlling overall levels of excitability within the cortical network (see the following text). In general, cortical circuits have recurrent, as well as feedforward and feedback, connections. Weak feedforward signals can be amplified by recurrent excitation; however, to prevent runaway activation, recurrent inhibition is believed to act as a form of brake (Galarreta & Hestrin 1998).

Synaptic depression has been shown to have both a fast (order of 100s ms) and slow (seconds) component (Varela et al. 1997). Both the fast (Varela et al. 1999) and slow components (Galarreta & Hestrin 1998) are sensitive to frequency of stimulation and asymmetric, in that excitatory synapses depress more strongly than inhibitory synapses. It has been generally assumed that asymmetric synaptic depression acts to shift the dynamics of the network towards inhibition (Galarreta & Hestrin 1998, Varela et al. 1999), however it has recently been argued (Chelaru & Dragoi 2008) that if pre-synaptic inhibition is the major contributor to the cortical circuit, suppression of the excitatory drive acting on inhibition (resulting from strong depression of excitatory synapses) could shift the bias of

the system towards excitation, thus responses would be facilitated as a consequence of asymmetric synaptic depression; under the same conditions, balanced synaptic depression would result in a decrease in neural-gain (see Figure 5.7).

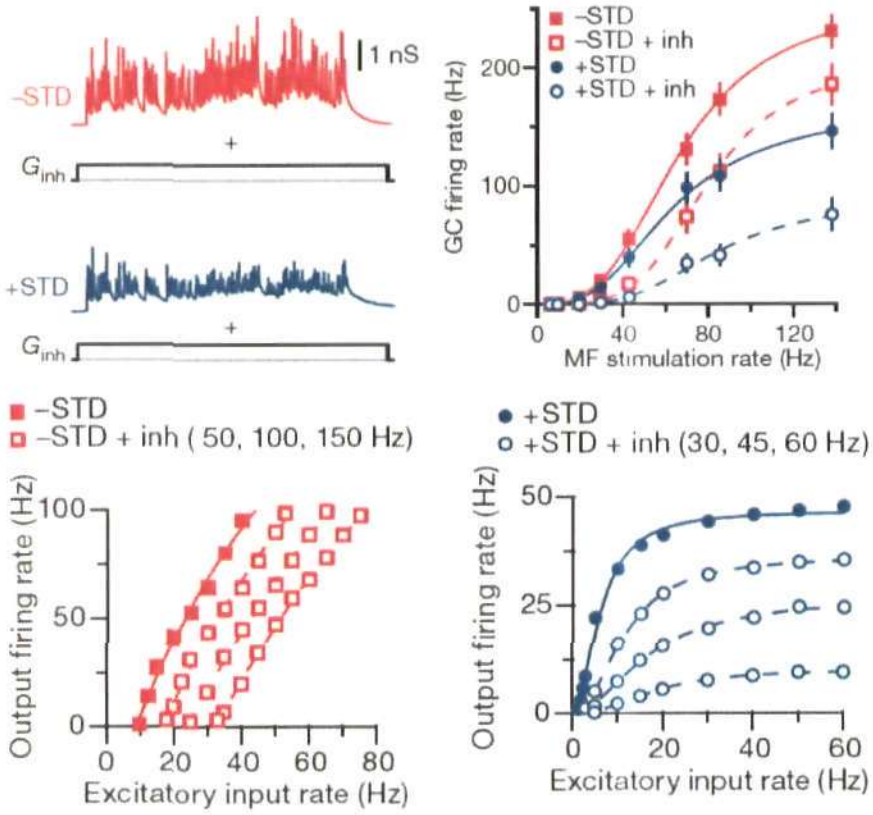

Figure 5.6: Tonic inhibition-mediated gain modulation in the cerebellum
Top left: Sample excitatory conductance with (blue) and without (red) synaptic depression (STD) used to stimulate a cerebellar granule cell, during a dynamic clamp experiment; black lines plot the tonic inhibitory conductance (G_{inh}). **Top right**: Average response of a granular cell, during the dynamic clamp experiment, to stimulation with and without STD (blue and red symbols, respectively) and tonic inhibition (open symbols/dashed lines and filled symbols/solid lines, respectively); lines plot the best-fit to the data. **Bottom left**: Input/output relationship of a model cell without STD; solid symbols/lines gives the control response, empty symbols/dashed lines plot the response as a function of inhibitory input rate (moving from left to right). **Bottom right**: Input/output relationship of a model cell with STD; solid symbols/lines gives the control response, empty symbols/dashed lines plot the response as a function of inhibitory input rate (moving from top to bottom).Taken from Rothman et al. (2009).

Gain control through adjustments in the balance between synaptic inhibition and excitation has also been found under conditions of spike-timing dependent plasticity (STDP), a well explored cortical and hippocampal phenomenon in which synapses are strengthened (long-term potentiation) when a pre-synaptic spike precedes a post-synaptic action potential and weakened (long-term depression) by the reverse sequence of spiking events (e.g. Bliss & Lomo 1973; Markram, Lübke, Frotscher & Sakmann 1997; Zhang, Tao, Holt, Harris & Poo 1998, Bi & Poo 1998, Feldman 2000). Long-term potentiation and depression has been observed in both excitatory and inhibitory synapses via a di-synaptic mechanism; in which inhibition can either be potentiated or depressed through the strengthening or weakening of excitatory inputs at the first synapse in the circuit (Lamsa, Heeroma & Kullman 2005).

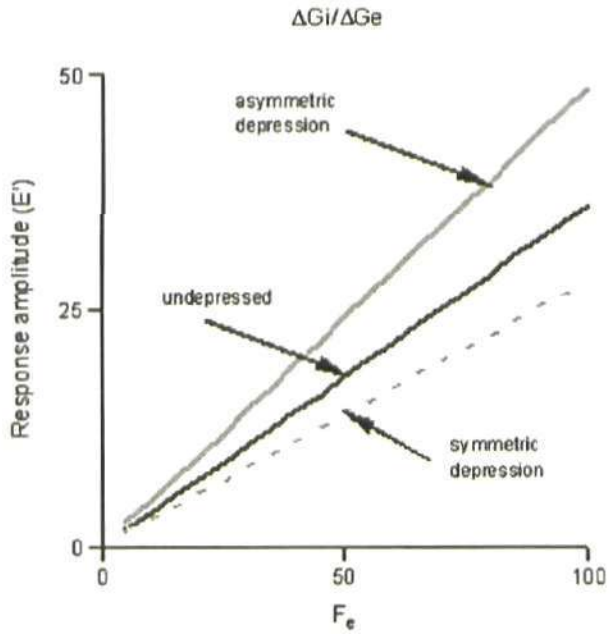

Figure 5.7: Asymmetric and symmetric synaptic depression.
The response of an arbitrary, model, cell (see Section 3.3, below) to feedforward excitatory input (F_e) under conditions of no depression (black line), balanced depression (grey, dashed line) and asymmetrical depression (grey, solid line).Taken from Chelaru & Dragoi, (2008).

A recent experimental and modelling study (Carvalho & Buonomano 2009) looking at STDP in the hippocampus found that the full range of adaptive gain changes (i.e. lateral shifts and/or an increase/decrease in the slope of the neural rate function) were elicited (and simulated in a computational model) by modulating the levels of long-term plasticity in both excitatory and inhibitory inputs (see Figure 5.8).

Lateral displacement of the input/output function arose due to either excitatory long-term potentiation (dark-green line; Figure 5.8, Top) or long-term depression (light-green line; Figure 5.8, Top) alone (see also dashed, black lines in Figure 5.8, Bottom).

Inhibitory long-term depression induced a reduction in stimulus threshold, whilst increasing gain (red line; Figure 5.8, Top); whilst inhibitory long-term potentiation induced the opposite effect (magenta line; Figure 5.8, Top).

The balanced potentiation of both excitation and inhibition resulted in a decrease in neural gain only, i.e. threshold remained constant; balanced long-term depression induced the opposite effect (see dashed, grey lines in Figure 5.8, Bottom).

Whilst there is no suggestion that STDP is the mechanism underpinning adaptation to stimulus statistics, the effects of STDP on neural output has parallels with the other mechanisms highlighted in this chapter. Long-term depression of excitation causes a rightward displacement in the neural function, much as spike-rate adaptation does; both long-term depression and spike-rate adaptation work by increasing the stimulus strength required to elicit a given response (e.g. the half maximal response). A change in neural gain was achieved through adjusting the non-linear relationship between excitation and inhibition, as a function of stimulus intensity. This is of course, similar to the inhibition-mediated gain control model (Murphy & Miller 2003) highlighted earlier, but is much

closer, in its mode of action, to the model of asymmetric depression (Chelaru & Dragoi 2008) outlined above. For all models, however, decreased neural gain arises due to increases in both excitation and inhibition.

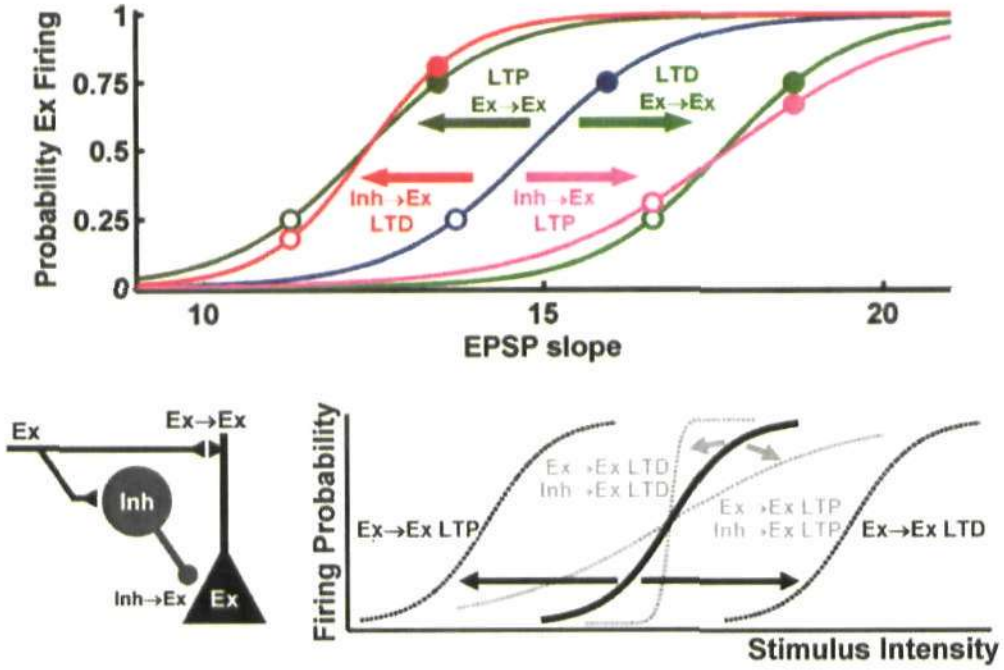

Figure 5.8: Long term depression (LTD) and potentiation (LTP) of excitation and inhibition

The response of an arbitrary, model, cell to LTP and LTD of excitatory and inhibitory inputs. **Top:** Response as a function of spiking probability verses slope of the post-synaptic potential; blue line plots the control response. LTP of excitation (dark-green line) would result in a leftward shift in stimulus-response function of the model cell whilst LTD of excitation (light-green) would elicit a rightward shift. LTP of (red line) and LTD (magenta line) would results in a decrease and increase in gain, respectively with lateral displacement in the opposite direction of excitation. **Bottom:** Balanced LTP/LTD of both excitation and inhibition (grey dashed lines) result in a decrease and increase in gain (i.e. the slope of firing probability plotted as a function of stimulus intensity), respectively, whilst threshold is unchanged; black, dashed lines plot LTP/LTD of excitation only; the diagram to the left is a hypothetical di-synaptic unit. Taken from Carvalho & Buonomano (2009)

5.1.3: Relating models of gain modulation to adaptation to variance

In general, the models described above are focused on the steady-state response to isolated stimuli, rather than the adaptive response to time-varying stimuli. The question addressed below is how can these models be related to the adaptation data, which has been the focus of this thesis?

Modulating levels of background noise

If one refers back to the results highlighted in Figure 5.4, it is clear that the model of gain modulation via increased background noise only accounts for the *in vitro* condition and represents the steady-state response to pulses of current injection. How then can this model of gain modulation be reconciled with the decrease in gain, observed *in vivo*, during adaptation to stimulus variance (Garcia-Lazaro et al. (2007) and Chapter 3, this thesis)? For the model to represent accurately the adaptive mechanisms present *in vivo*, two important assumptions must be made: firstly, that background firing rates are dependent, to some degree, on the stimulus itself and, secondly, that background input rates only increase in line with the variance of the stimulus.

The caveat that background rates must be dependent on the input is not without justification; the somatotopology of the barrel cortex is such that neighbouring cells, within a column, are generally tuned to the same principle whisker (Armstrong-James & Fox 1987; Simons 1978), especially within the barrel regions of layer 4 (Woolsey & der Loos 1970). If driving input can be considered as that which is derived directly from thalamic input (or in the case of layer 2/3 cells, from feedforward afferents from layer 4 or layer 5, Shepherd, Stepanyants, Bureau, Chklovskii & Svoboda 2005), then noisy, background inputs arise from the remainder of the synaptic input: connections of inter-columnar, inter-

areal and, significantly, local and/or recurrent origin. For example, layer 4 stellate cells synapse with at least 30% of surrounding neurons, within the barrel region of the cortex (Feldmeyer, Egger, Lubke & Sakmann 1999), whilst layer 3 pyramid cells have been found to connect extensively, within the cortical column, both intra- and inter-laminarly (Feldmeyer, Lubke & Sakmann 2006). Whilst recent evidence suggests that there is a large degree of heterogeneity, within columns, in the strength of principle whisker selectivity (Sato, Gray, Mainen & Svoboda 2007) it nevertheless remains true that stimulation of a neuron's principle whisker will elicit responses in both the neuron and its neighbouring cells. Thus it can be argued that the levels of background input can be influenced by the stimulus.

The second assumption, that background firing rates increase only with stimulus variance and not, as one might intuitively expect, with stimulus mean, is also valid when one considers how adaptation affects the output of a cell. Consider that adaptation to global mean tends to shift the linear portion of the stimulus-response function to regions in the stimulus space that are generally of larger amplitude than that covered by the high-probability region. Therefore, responses to the stimuli within the high-probability region are generally suppressed and, importantly, similar across all mean adapting levels. Thus, as the high-probability region stimuli are presented most frequently, it can be argued that overall background firing rates wouldn't increase in proportion to any increase in the mean of the adapting stimulus, especially if adaptation occurred rapidly.

In the case of adaptation to variance, as the width of the high-probability region increases there is a downward scaling of the stimulus-response function, thus intuitively one might expect to see a decrease in background firing rates. However, as there is no or

little lateral displacement of the neural rate function (as the adapting stimulus variance increases), there is less suppression of the response to the strongest stimuli within the high-probability region, relative to lower variance conditions. Thus one can argue that, overall, background firing rates would increase with stimulus variance.

Tonic inhibition

In terms of inhibition-mediated gain modulation the observed decrease in neural-gain, following adaptation to increased levels of variance, could be generated by a process of increasing tonic inhibition in proportion to the amount of noise within the adapting stimulus. The excitatory drive acting on the target cell would also have to be transformed from a linear to non-linear signal, either sub-cortically or during thalamocortical transmission.

As stated, one potential source of non-linearity in the cortex is synaptic depression. Use-dependent synaptic depression has been observed in both thalamocortical and cortico-cortical connections within the barrel cortex (e.g. Petersen 2002) and modulatory tonic inhibition is present in all layers of the somatosensory cortex (but strongest in Layer V [Yamada, Furukawa, Ueno, Yamamoto & Fukuda 2007; Kyriazi, Carvell, Brumberg & Simons *1998*]) where it is mediated mainly via $GABA_A$-type receptors (Salin & Prince 1996; Kyriazi, Carvell, Brumberg & Simons 1996). Thus the conditions necessary to generate gain modulation, within the constraints of either approach, are present in the barrel cortex.

Balance of synaptic depression

Adaptation to variance could be achieved through a process of adjusting the overall

influence of inhibition on the cortical network. This could be achieved by manipulating the relative levels of inhibitory and excitatory depression in a manner that is dependent on stimulus noise; i.e. excitatory and inhibitory depression would become progressively more balanced as stimulus variance increased.

The nature of the adapting stimulus itself means that some relative relief from depression should occur as stimulus variance increases. As the width of the high-probability region widens, the number of possible inputs both above and below the central stimulus increases thereby increasing the probability of presentation of stimuli that do not elicit an afferent response and allowing for the synaptic resource to be replenished. Thus, if one assumes that excitatory depression levels are more sensitive to stimulus noise, one could argue that as variance increases, net excitatory depression is weakened, driving it closer to inhibitory levels of depression.

What evidence is there that the balance between excitatory and inhibitory synaptic depression is adjusted, within the barrel cortex, in a stimulus-dependent manner?

It is well established that the strength and time course of depression is sensitive to stimulus frequency (e.g. Varela et al. 1997, Varela et al. 1999, Galarreta & Hestrin, 1998), however, it is less clear how stimulus frequency affects the balance between afferent excitation and inhibition. Within layer IV of the rat barrel cortex, periodic whisker deflection, has been shown to adapt incoming excitation and inhibition to an equal degree (Higley & Contreras 2006). However, for longer periods of stimulation, inhibition tends to adapt much more strongly than excitation, resulting in disinhibition within the circuit (Heiss et al. 2008). Whilst none of the two studies (Higley & Contreras, 2006; Heiss et al. 2008) addressed the mechanisms of adaptation directly, one could argue that differences in

the strength of inhibition reflect a change in the balance of synaptic depression, i.e. for long periods of stimulation excitatory depression becomes stronger, thus reducing the overall levels of inhibition; of course, sustained stimulation resembles most strongly the low-variance stimulus regime.

Key features of the models

From the review above it would seem that the crucial factor underpinning gain control is the balance between inhibition and excitation. However, the tonic inhibition model also relies heavily on a transformation of the relationship between stimulus and driving input to elicit a change in either neural threshold or gain. If synaptic plasticity is the source of non-linearity, then it would have to be 'switched on and off' depending on the adapting regime.

Of course, there is evidence to suggest that thalamocortical depression, both in the barrel (Castro-Alamancos 2004a; Castro-Alamancos 2004b) and visual cortex (Boudreau & Ferster 2005), is highly saturated due levels of tonic activity in the thalamus. Therefore, it is possible that under conditions of adaptation to global mean and low-variance, thalamocortical transmission of afferent excitatory input would be highly suppressed (due to the sustained presentation of a small number of stimulus intensities). In these circumstances increased tonic inhibition could lead to lateral displacement of the neural-rate function. For higher-variance adapting conditions, there would be a corresponding release from saturation in levels of thalamocortical depression. However, viewed in isolation, release from depression in afferent excitation could be expected to increase gain, as firing rates would be increased. Conversely, when viewed within the constraints of the

asymmetric model of depression, outlined above (Chelaru & Dragoi 2008), release from synaptic depression is actually a necessary condition for generating gain modulation.

In conclusion, the mechanism of adjusting the relative balance between the influence of excitation and inhibition (Chelaru & Dragoi 2008) offers the most flexibility of all the models outlined here. However, as there is a clear mechanism by which changes in the balance of synaptic depression can be achieved, it is argued that this represents a biologically valid model.

As such a model of asymmetric depression was developed, based largely on the model developed by Chelaru and colleagues (Chelaru & Dragoi 2008), and tested under conditions of adaptation to variance and mean. This was compared to a model of gain modulation through increasing the firing rates of background noise as suggested by Frances Chance and colleagues (Chance et al. 2002).

5.2: Model cell: pulse based model

It was decided to develop a simple neural model, i.e. a single-compartment Hodgkin-Huxley type model, to assess the validity of the putative mechanisms of gain modulation outlined in the introduction above.

As the primary interest was in exploring extrinsic cellular mechanisms it was felt that a single-compartmental model would be more useful than re-implementing the NEURON model outlined in Section: 5.1. above (Garcia-Lazaro et al. 2007).

NEURON models have the advantage of being biologically realistic; however their complexity can be a disadvantage when one is interested in general mechanisms only as one must untangle the affects of dendritic structure on the passive properties of the cell.

The equations governing the Hodgkin-Huxley (HH) equations (Hodgkin & Huxley 1990), for a single-compartmental neural model are given below:

$$\begin{aligned} CmV &= -g_{leak}(V - E_{leak}) - g_{Na} m^3 h(V - E_{Na}) - g_k n^4 (V - E_k) \\ \dot{m} &= -\alpha_m(V)(1 - m) - \beta_m(V)m \\ \dot{h} &= -\alpha_h(V)(1 - h) - \beta_h(V)h \\ \dot{n} &= -\alpha_n(V)(1 - n) - \beta_n(V)n \end{aligned}$$

5.1

Where Cm is the membrane capacitance, V is the membrane potential, g_{leak}, g_{Na} and g_K are the leak, sodium and potassium conductances, respectively, E_{leak}, E_{Na} & E_k are the corresponding reversal potentials and m, h and n are the gating-particle variables, governed by the rate constants α and β

One issue with the HH model is that it is computationally expensive; the gating-particle variables must be differentiated at each time step. Integrate and fire neurons (I & F) offer the simplest model of a spiking cell, however they generally ignore the complexity of possible currents/conductance's acting on the cell's membrane potential. As a compromise between both models, Alain Destexhe developed a conductance-based integrate and fire neuron (Destexhe 1997) that assumed that gating dynamics, during firing, can be approximated as 'pulses' and as such don't need to be differentiated at every step. The values of the voltage-dependent parameters (α and β, see equation 5.1) were fixed at their value when the membrane potential was 20mV or-70mV, depending on the firing/non firing state of cell (based on the Traub-Miles model [Traub & Miles 1991]). A comparison between the HH model and the pulse-based (PB) model is shown in Figure 5.9. Both the PB and HH models respond in almost identical manner to a current injection of 20nA; however the PB model is the most efficient in terms of computation power, as there are fewer calculations per time step.

5.2.1: Single-compartment, pulse-based, Integrate & Fire, neural model with ionic conductances

The pulse-based model (PB model) uses the Hodgkin-Huxley (HH) equations as its basis. The gating-particle variables, m, h and n depend on the evolution, over time, of the voltage-dependent rate constants: α_m, β_m, α_h, β_h, α_n, β_n. These exhibit sharp transitions during spiking and as such can be approximated by determining their value at two extremes of membrane voltage (20 and -70 mV).

The rate constants are thus assigned as constants, whose values depends on whether the cell is firing or not (see Table 5.1, below).

Table 5.1: Values of the voltage dependent rate constants under firing/not firing conditions
Using the Traub-Miles formulation (Traub & Miles 1991)

Rate constant (mV)	Firing	Not firing
$\alpha_m(20)$	$22ms^{-1}$	0
$\beta_m(-70)$	0	$13ms^{-1}$
$\alpha_h(-70)$	0	$0.5ms^{-1}$
$\beta_h(20)$	$4ms^{-1}$	0
$\alpha_n(20)$	$2.2ms^{-1}$	0
$\beta_n(-70)$	0	$0.76ms^{-1}$

Within the PB model, action potentials are of a fixed duration (0.6ms), thus, once the membrane potential crosses a given threshold, a spike or pulse is generated.

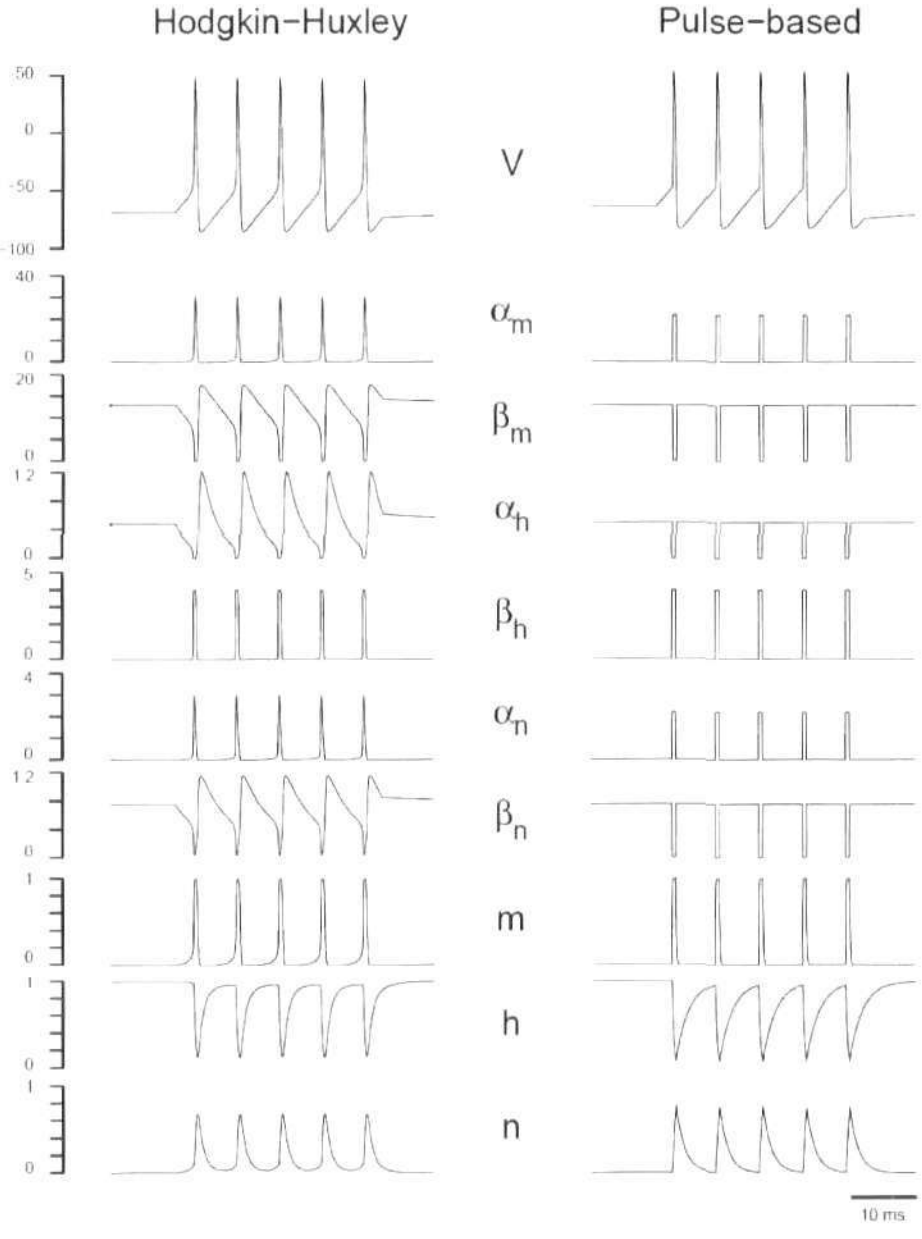

Figure 5.9: Comparison of the HH model with the pulse-based (PB) model
For parameters, please see the text. Taken from Destexhe (1997).

During firing, m, n, and h are given by:

$$\begin{aligned} m(t) &= 1 + (m_0 - 1)\exp[-\alpha_m(t - t_0)] \\ h(t) &= h_0 \exp[-\beta_h(t - t_0)] \\ n(t) &= 1 + (n_0 - 1)\exp[-\alpha_n(t - t_0)] \end{aligned}$$

5.2

Where t_0 is the time at which the pulse started.

When the cell is not firing, m, n, and h are given by:

$$\begin{aligned} m(t) &= m_0 \exp[-\beta_m(t - t_0)] \\ h(t) &= 1 + (h_0 - 1)\exp[-\alpha_h(t - t_0)] \\ n(t) &= n_0 \exp[-\beta_n(t - t_0)] \end{aligned}$$

5.3

Where t_0 is the time at which the last pulse ended.

The membrane potential at each time step is calculated as:

$$V(t+1) = V(t) + \mathrm{dt}(I_{inj} - g_{leak}(V - E_{leak}) - \bar{g}_{Na} m^3 h(V - E_{Na}) - \bar{g}_k n^4 (V - E_k)/cm)$$

5.4

The PB cell model, was implemented in MatLab (release 2009a), using the same parameters as in Destexhe (1997), see Table 5.2. A comparison between the implemented PB model and a HH model is plotted in Figure 5.10. For the HH model, the rate constants were derived (at each time step) using the Traub-Miles formulas (Traub & Miles 1991), as derived by Ermentrout (Ermentrout, Pascal & Gutkin 2001), see Equation 5.5:

$$\begin{aligned} \alpha m(V) &= 0.32(V + 54)/(1 - \exp(-(V + 54)/4)) \\ \beta m(V) &= 0.28(V + 27)/\exp((V + 27)/5) - 1) \\ \alpha h(V) &= 0.128\exp(-(V + 27)/18) \\ \beta h(V) &= 4/(12 + \exp(-(V + 27)/5)) \\ \alpha n(V) &= 0.032(V + 52)/(1 - \exp(-(V + 52)/5)) \\ \beta n(V) &= 0.5 * \exp(-(V + 57)/40) \end{aligned}$$

5.5

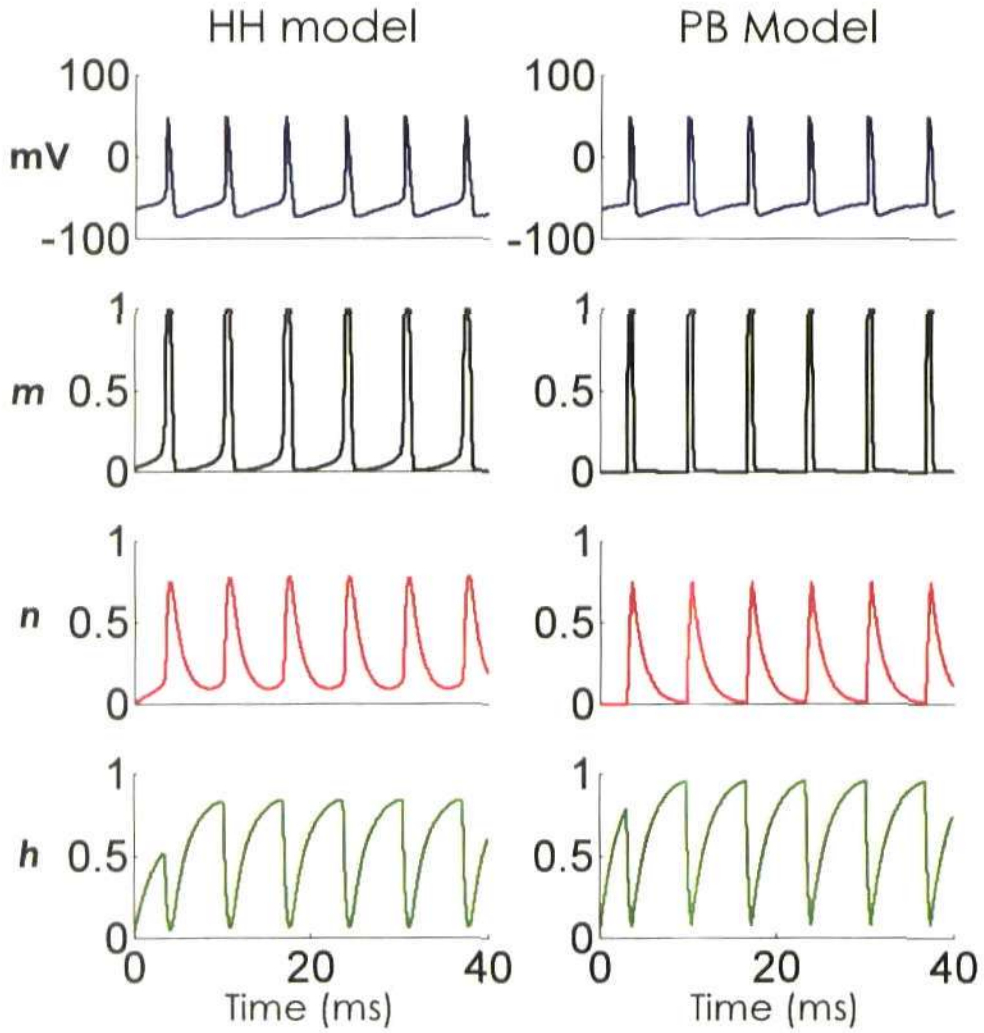

Figure 5.10: Comparison of the pulse-based and Hodgkin Huxley model implementations
Both cell models were stimulated with a depolarising 2nA current. The parameters for the models are in Table 5.2

Both implementations of the HH model and the PB model compare well with each other (and compare favourably with the implementation of Destexhe (1997), see Figure 5.9 & Figure 5.10, for comparison), thus it was decided to use this implementation of the PB model to test the various models of adaptation.

Table 5.2: List of pulse-based model cell parameter values (taken from Destexhe 1997)

Parameter	Value
Absolute Threshold	-54.3mV
E_{leak}	-70mV
E_{Na}	50mV
E_K	-77mV
Cm	$1\mu F/cm^2$
g_{leak}	$0.3mS/cm^2$
g_{Na}	$120mS/cm^2$
g_K	$36mS/cm^2$
Δt	0.01ms
Spike duration	0.60ms
Refractory period	2.00ms

5.2.2: Adapting firing threshold

As stated within the introduction, the $I_{K(Ca)}$ current, is believed to underpin, in part at least, spike frequency adaptation and has been shown to be responsible for adaptive shifts in the stimulus-response function in two reconstructed model cells (see Section 5.1.1). However a simpler method of generating spike frequency adaptation is to introduce a dynamic firing threshold (Liu & Wang 2001). Each time a spike is fired the adaptive threshold, θ, is increased by a constant weight b and then decays back to the absolute threshold, θ_0, with time constant, τ_{Adapt}.

Under this regime both membrane potential and threshold are artificially inflated above typical physiological levels (e.g. between -50-40mV). Under a regime of $I_{K(Ca)}$ current-based spike-rate adaptation, membrane threshold is constant; the $I_{K(Ca)}$ current increases hyperpolarisation, thereby affecting the temporal dynamics of the membrane potential are affected (as it takes relatively longer to reach threshold from resting potential

levels).

For dynamic changes in membrane threshold to occur *in vivo*, the sensitivity of the voltage-gated sodium and potassium membrane channels would have to adjust every time a spike was fired. The $I_{K(Ca)}$ current represents the most biologically valid mechanism of spike-rate adaptation. However, the introduction of a hyperpolarising $I_{K(Ca)}$ current into the model cell would impact on maintaining the balance between background inhibition and excitation. Thus it was decided to incorporate the less realistic, but computationally more efficient, dynamic firing threshold mechanism into the model cell.

The threshold at each time step, $\theta(t)$, is governed by Equations 5.6-5.8.

$$\frac{d\theta}{dt} = \frac{\theta(t) - \theta_0}{\tau_{Adapt}} \qquad \textbf{5.6}$$

If a spike is fired at time-step = t, then the threshold, θ, at time-step = $t+1$ is given by:

$$\theta(t+1) = \theta(t)(1 - (dt/\tau_{Adapt})) + ((\theta_0/\tau_{Adapt})dt) + b \qquad \textbf{5.7}$$

Otherwise:

$$\theta(t+1) = \theta(t)(1 - (dt/\tau_{Adapt})) + ((\theta_0/\tau_{Adapt})dt) \qquad \textbf{5.8}$$

The threshold and output of the model cell, over time, for various levels of the weight constant, b, are plotted in Figure 5.11.

Including a dynamic threshold, within the PB model, generated spike frequency adaptation (the inter-spike interval is increased with each successive spike), the strength of

which was determined by the parameter *b*.

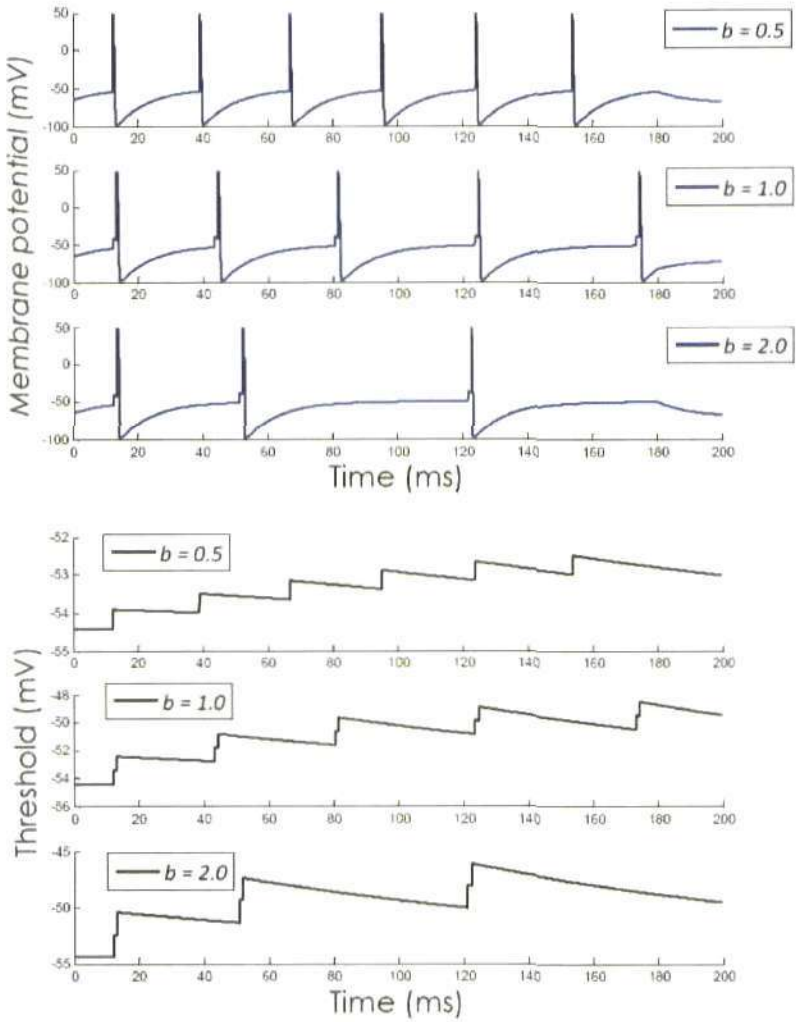

Figure 5.11: Spike-rate adaptation in the model cell
Showing effect of the parameter *b* on firing threshold and membrane potential of a model cell, in response to a simulated current step injection of 2nA; duration: 180ms. **Top:** The dynamic threshold of the cell as a function of *b*. **Bottom:** Membrane potential of the model cell.

5.2.3: Stimulus

The stimulus for the model cells were always one of the following (unless explicitly stated otherwise in the text):

- Current steps - cells were injected with simulated current steps ranging from 0 to 0.4nA (in either 0.1 or 0.05 steps). In general, each current step was presented for 500ms; the response, for a given current step, was taken to be the sum of the spike count during the current pulse, divided by the pulse duration.

- Adaptation to global mean - the stimulus distribution ranged from 0 to 0.4011nA (in 0.0191nA steps). The high-probability regions (high-probability region) were 0.0955nA wide (i.e. spanning 5 amplitudes) and were centred on 0.0573nA, 0.1337nA, 0.2101nA & 0.265nA. As per the experimental paradigm, the model cell was stimulated for 30s, under each high-probability region condition, and amplitudes were refreshed every 40ms. An example of the stimulus, under each high-probability region condition is plotted in Figure 5.12).

- Adaptation to variance - the stimulus paradigm was as for adaptation to the mean except the distribution ranged from 0 to 0.456nA (in 0.019nA steps) and the high-probability region were all centred on 0.152nA and ranged in size from 0.1330-0.1710nA (spanning 3 amplitudes, low-variance condition), 0.0950-0.2090nA (spanning 7 amplitudes, mid-variance) & 0.0570-0.2470nA (spanning 11 amplitudes, high-variance). A sample stimulus, of 500ms duration, is plotted in Figure 5.13.

In some simulations the stimulus was not an injected current but rather an excitatory conductance; under these circumstances, the driving current was converted into a firing rate (Hz) and fed into a Poisson spike generator, from which the conductance was calculated. The conversion from current injection into firing rate was achieved by multiplying the stimulus (in nA) with a conversion constant, c (in this case $c = 40$); c was determined by finding which value resulted in an average synaptic current (as a function of

the conductance and membrane voltage) that was as close as possible to the original driving current.

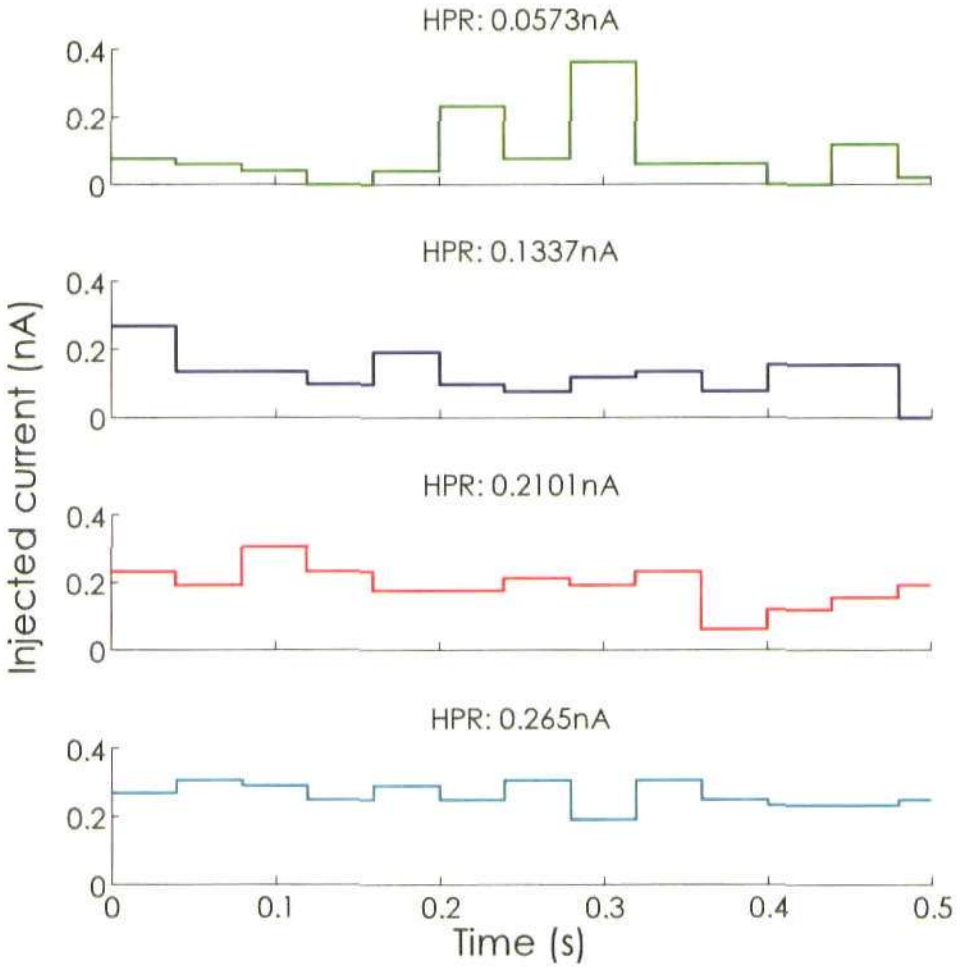

Figure 5.12: Sample of the stimulus to the model cell: Adaptation to mean
Please see text for details

Synaptic conductances were calculated using a difference-of-exponentials function, with a rise time of 0.1ms; decay time was 10ms for the inhibitory conductance and 5ms for excitatory.

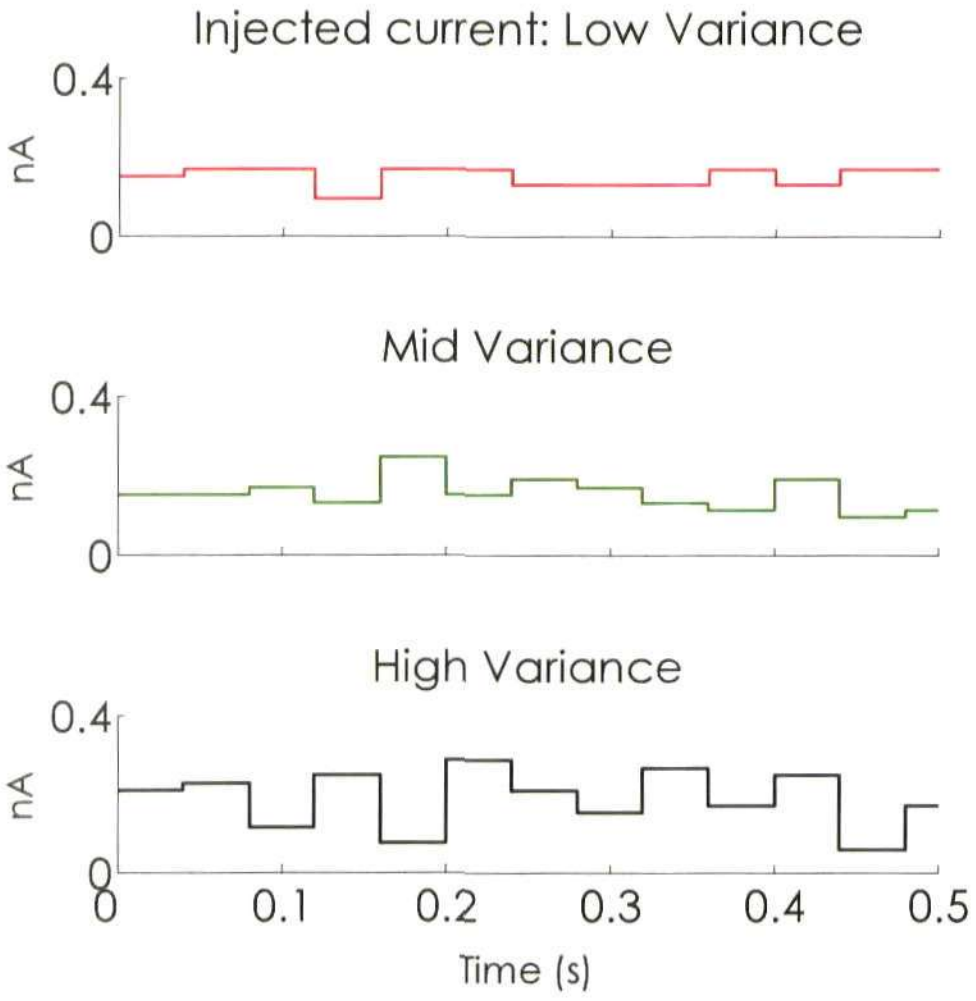

Figure 5.13: Sample of the stimulus to the model cell: Adaptation to variance
Please see text for details

Peak conductance was set to 0.4 and 1.6 times the resting conductance for excitatory and inhibitory conductances, respectively, and the time step, Δt, was 0.01ms; the parameters were taken from Chance et al. (2002). Each pre-synaptic spike generated a unitary conductance (calculated for a 100ms following spike generation), which was added to the running total of the synaptic conductance. The excitatory and inhibitory reversal potentials were 0mV and -90mV, respectively.

5.2.4: Tuning of the model cell

The average tuning of the model cell to injected current steps of 500ms duration is plotted in Figure 5.14, below. The response increased monotonically, with stimulus current.

5.2.5: Synaptic depression as a non-linearity

As suggested earlier, synaptic depression may generate the necessary conditions for gain modulation, thus a description of synaptic depression, within the model, is given below. In order to replicate this in the model cell, the driving current was replaced by an excitatory conductance that underwent synaptic depression. Pre-synaptic spikes were generated by an inhomogeneous Poisson process, with the firing rate determined by the input current (the driving current was multiplied by the constant, c ($c = 40$), to convert to Hz).

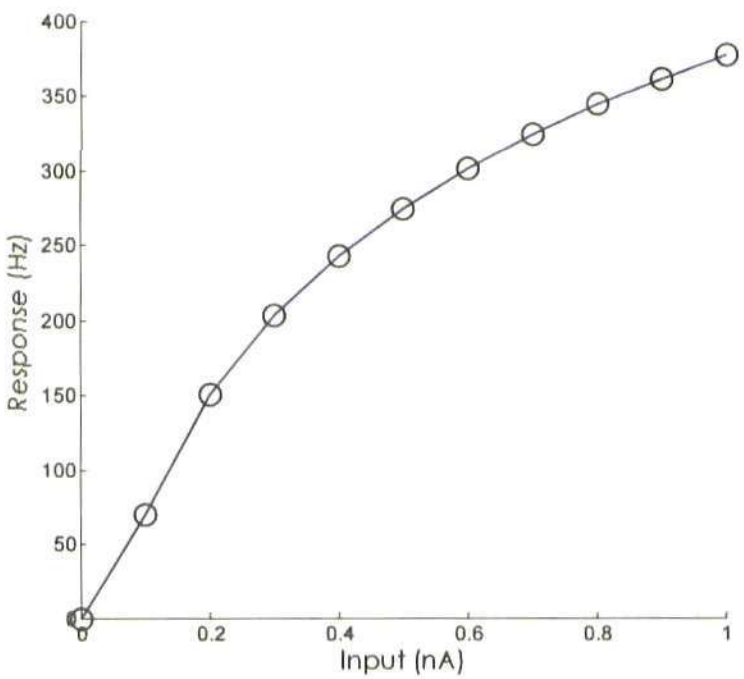

Figure 5.14: Tuning curve of the model cell to injected current steps of amplitudes 0nA-1nA

The parameters of the synaptic conductance were as for Section: 5.2.3 above, however, the conductance for each spike was multiplied by the fraction of synaptic resource available, U_{SR}, at the time of firing. For each action potential, U_{SR} was multiplied by a given depression fraction D and rose exponentially, in between spikes, back to the resting level of unity with a time constant of 300ms (Tsodyks & Markram 1997). Under a condition of no synaptic depression, D was equal to unity; otherwise D was set to be within the range of:

$$0.3 \leq D < 1 \qquad \textbf{5.9}$$

The closer D was to unity, the weaker the depression acting on the conductance.

An example of the excitatory conductance under the depressed and non-depressed condition is plotted in Figure 5.15, along with the available synaptic resource at each time step; the input rate was 30Hz (equivalent to a current injection of 0.75nA).

The average excitatory synaptic conductance (over 100 trials), at each time step, as a function of stimulus amplitude (0 to 1nA, in 0.1nA steps) and under both conditions (depression and no depression) is plotted in Figure 5.16. Without synaptic depression the average driving conductance was approximately a linear function of the input rate. However, in the presence of synaptic depression, the mean, excitatory, conductance diverged from the undepressed conductance almost immediately; the average depressing conductance also saturated rapidly.

5.2.6: Integrate & Fire units

Synaptic depression is highly dependent on the variation in the stimulus at each time-step, as this will affect recovery of the synaptic resource. Thus, in order to examine how the time varying stimulus affected output, simulations were also run, in which the

stimulus was used to drive a small network of Integrate & Fire (I & F) units with Poisson spiking statistics. A circuit was developed with 4 recurrently connected excitatory I & F neurons that had reciprocal connections to a single inhibitory I & F unit, that was also recurrently connected to itself; all five units provided afferent input to a target neuron that was used to measure the response to the stimulus.

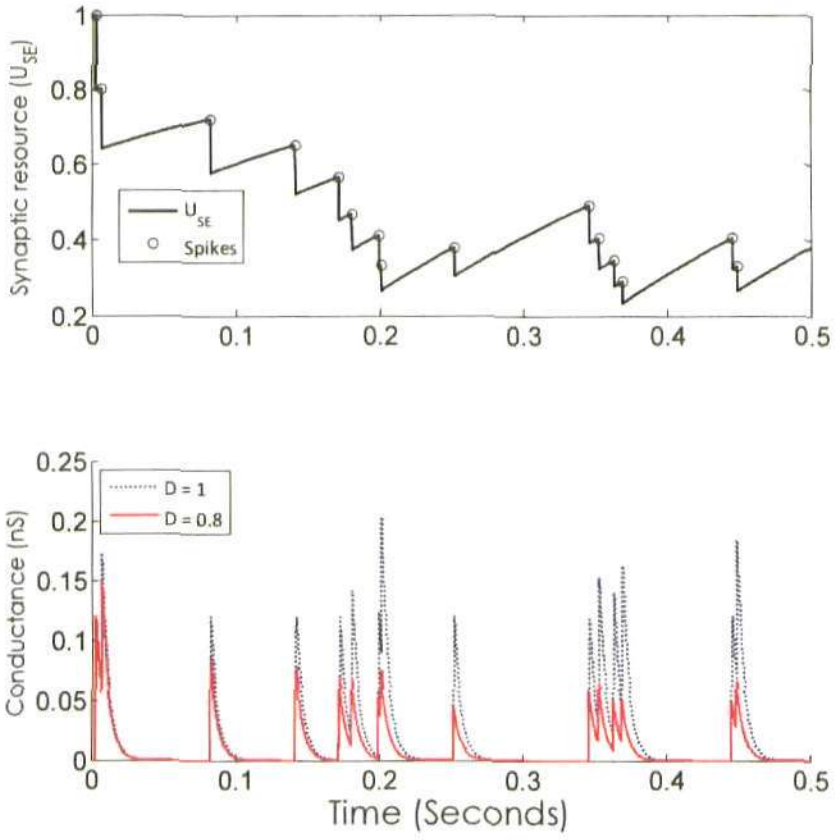

Figure 5.15: Comparison of a depressed and undepressed synaptic conductance, generated in response to the same spike train
Top: The amount of synaptic resource (solid lines) available at each time step, in the presence of synaptic depression (red lines, bottom figure), the circles indicate the spike time and the amount of synaptic resource available for each spike-generated conductance. **Bottom:** The actual synaptic conductance, in response to a given spike train, in the depressed (red, solid line) and undepressed (blue, dashed line) case, see text for details of parameters and spike-train generation (please see Section: 5.3.1 & 5.2.5)

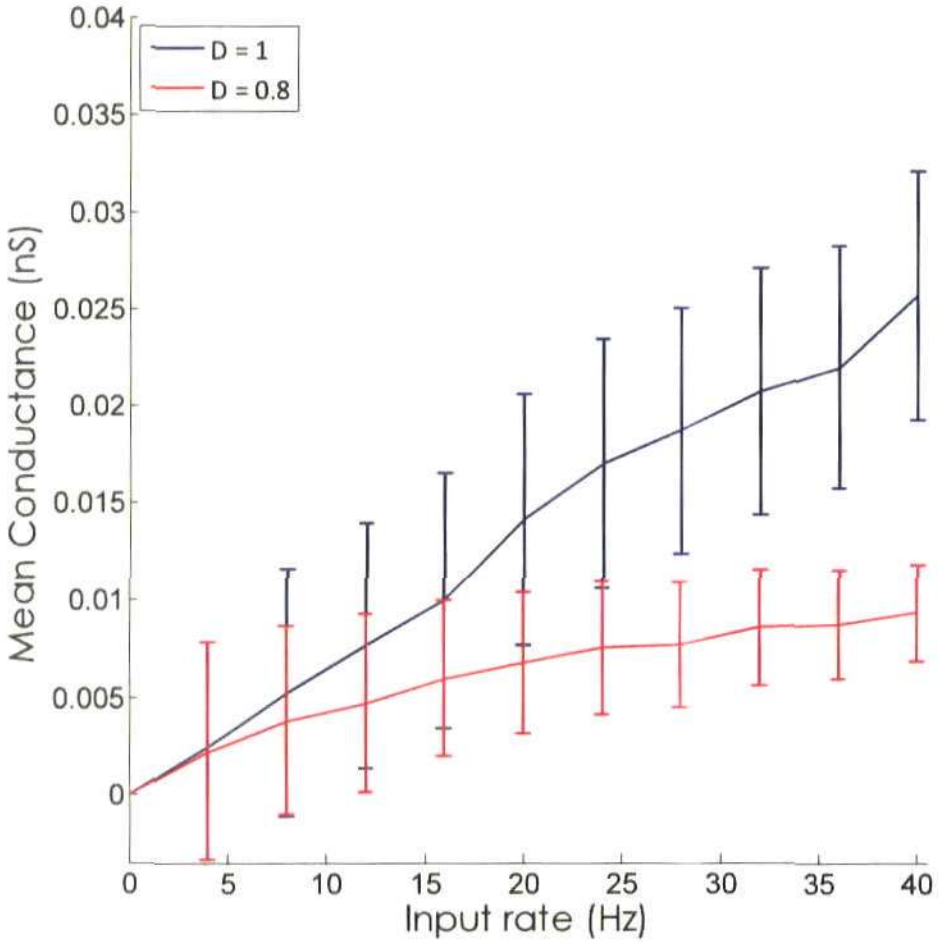

Figure 5.16: Comparison of the mean depressed and undepressed synaptic conductance, as a function of input rate
The average conductance was calculated, for each stimulus presentation and across ten trials, for each stimulus amplitude (0 - 1nA, converted into Hz), by taking the average of the steady-state (>200ms) conductance used to drive the model cell under both depressed (red line) and undepressed (blue line) conditions; the error bars give the standard error, across the ten trials.

The parameters for the I & F units were as follows (parameters for peak conductances were derived from Chelaru & Dragoi [2008]):

- Membrane resistance (Rm): 10000 Ohms;
- Membrane time constant (τ_m): 10ms;

- Leak reversal potential (E_L): -75mV;
- Reset voltage (V_{reset}): -80mV
- Threshold: -54mV
- Refractory period: 2ms;
- Peak conductance of recurrent excitatory conductance:0.1nS
- Peak conductance of recurrent inhibitory conductance:0.6nS
- Peak conductance of recurrent excitatory to inhibition conductance:0.44nS
- Peak conductance of recurrent inhibition to excitatory conductance:0.49nS
- Peak conductance of feedforward inhibition/excitation on target cell was 0.375/0.350nS respectively.

All other parameters, for generating synaptic conductances, were as for Section 5.2.3 above.

The membrane potential at each time step (t), for the time-dependent current (I), was calculated as:

$$V_{inf} = E_L + Rm(I(t) - (\sum_{j=1}^{n} g_{me}(t-1)(V(t-1) - V_E)) - (g_m(t-1)(V(t-1) - V_I)))$$
$$V(t) = V_{inf} + (V(t-1) - V_{inf})\exp(-dt/\tau_m) \quad \textbf{5.10}$$

Where m is the polarity of the I & F unit, i & e represent the inhibitory/excitatory recurrent and reciprocal conductances, n is the number of excitatory units.

5.3: Replicating models of gain modulation in a pulse-based neural model.

As stated earlier, the aim of Chapter 5 was to explore which of the various models

of gain modulation, highlighted in the introduction, can best account for the decrease in neural-gain observed following adaptation to stimulus variance. The first stage of analysis was to replicate the general results from each model, using the conductance-based neural cell model described above. The first model replicated was the background noise model, as proposed by Chance and colleagues (Chance et al. 2002).

5.3.1: Increased background noise decreases neural gain

In order to replicate the results from Chance et al. (2002), simulated current pulses were injected into the model cell (see Figure 5.17).

. The stimulus was a square-wave current pulse (duration: 500ms) and increased in increments of 0.05nA with each presentation (from a starting value of 0nA and a final value of 0.4nA). In addition to the driving current, background excitatory and in inhibitory conductances were injected into the model cell. Pre-synaptic spikes were generated using an inhomogeneous Poisson process, with the base firing rate (1X condition), for both inhibition and excitation, set to 50Hz. The conductances were calculated using a difference-of-exponentials function, with a rise time of 0.1ms; decay time was 10ms for the inhibitory conductance and 5ms for excitatory. For parameters see Section 2.3, above. The background conductance elicited average sub-threshold membrane fluctuations with a standard deviation of ±5mV. An example of the conductances under the 1-3X conditions is given in Figure 5.17.

The input/output (I/O) curve was taken to be the average number of spikes fired once the model cell had reached the steady state (taken to be from 200ms into stimulus presentation) as a function of the amplitude of the current-step. The I/O curve of the model cell under the 1-3X conditions are plotted in Figure 5.18. As can be seen, increasing

background firing rates resulted in a decrease in gain, thus the model cell fully replicated the *in vitro* data (Chance et al. 2002).

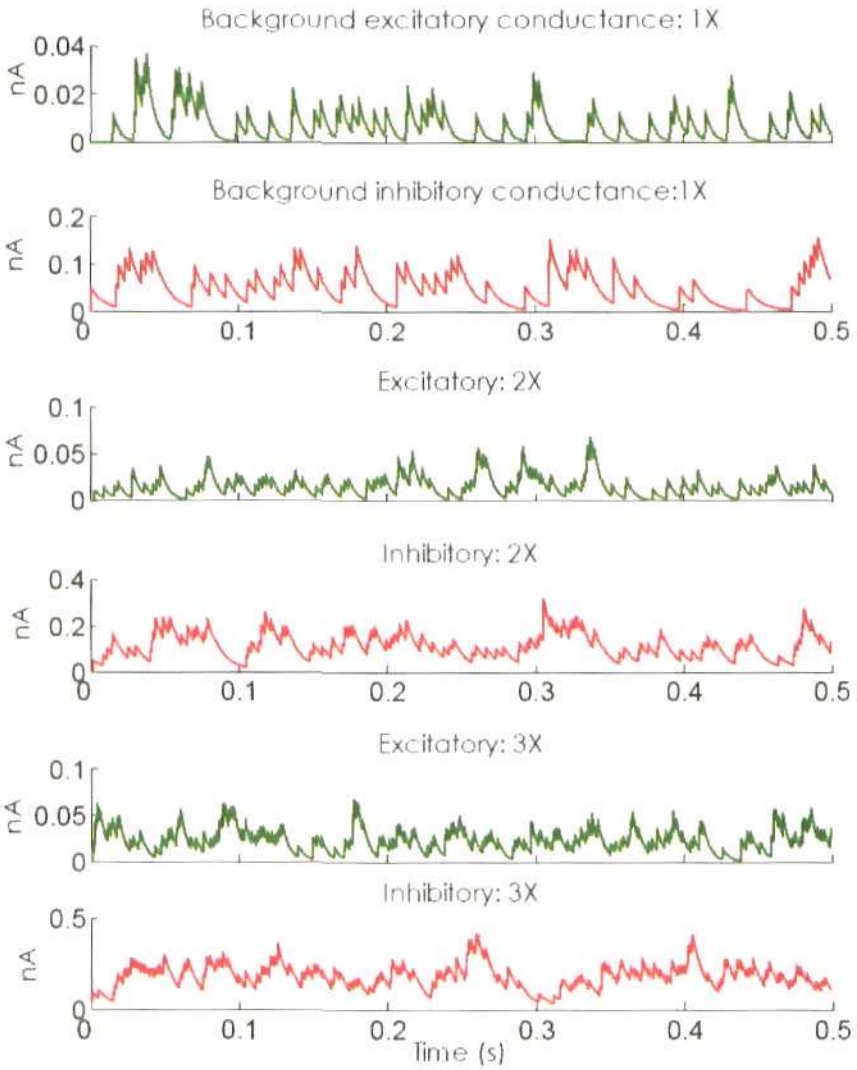

Figure 5.17: Example of background conductance at 50Hz
A sample of the background conductance levels under the 1-3X conditions (see text for details) are plotted from top (1X) to bottom (3X) axes; red plots the inhibitory and green the excitatory conductances

On average the half-maximal amplitudes, or threshold (S_{50}), were 0.3504nA, 0.4092nA & 0.4362nA, whilst slope, at the S_{50}, was 0.0407Hz/nA^2, 0.0292Hz/nA^2 &

0.0281Hz/nA^2 for the 1X, 2X and 3X conditions, respectively.

5.3.2: Inhibition-mediated changes in gain

The scaling of the stimulus-response function through increasing levels of tonic inhibition (Murphy & Miller 2003) was replicated in the model cell.

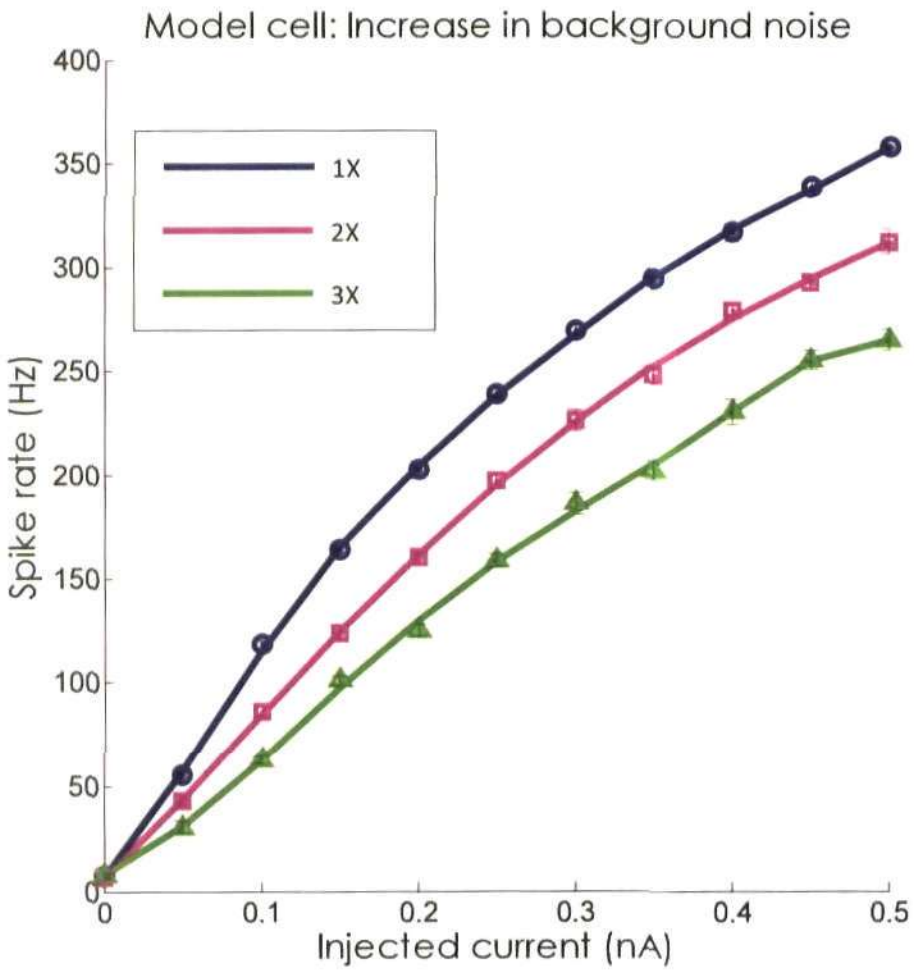

Figure 5.18: Increased background firing rates elicit decrease in neural gain
Background firing rates (1X (50Hz) –blue lines/symbol, 2X (100Hz) – magenta, 3X (150Hz) – green) were increased in a balanced manner and neural gain decreased. Solid lines plot the best-fit to the data (symbols).

The background conductances were as for the model above in the 1X configuration

(i.e. 50Hz, parameters as in Section 5.3.1, above). Tonic inhibitory conductances were mimicked by introducing a constant conductance of 0.25-0.75nS (the reversal potential for tonic inhibition was -80mV). The stimulus was the amplitude of the injected current in nA (from 0-1nA, in steps of 0.05nA). Under the non-linear condition the actual injected current into the model cell was a sigmoidal function of the stimulus:

$$s_{Inj} = I_{MAX}\left(s^n / s^n + s_{50}^n\right) \qquad \textbf{5.11}$$

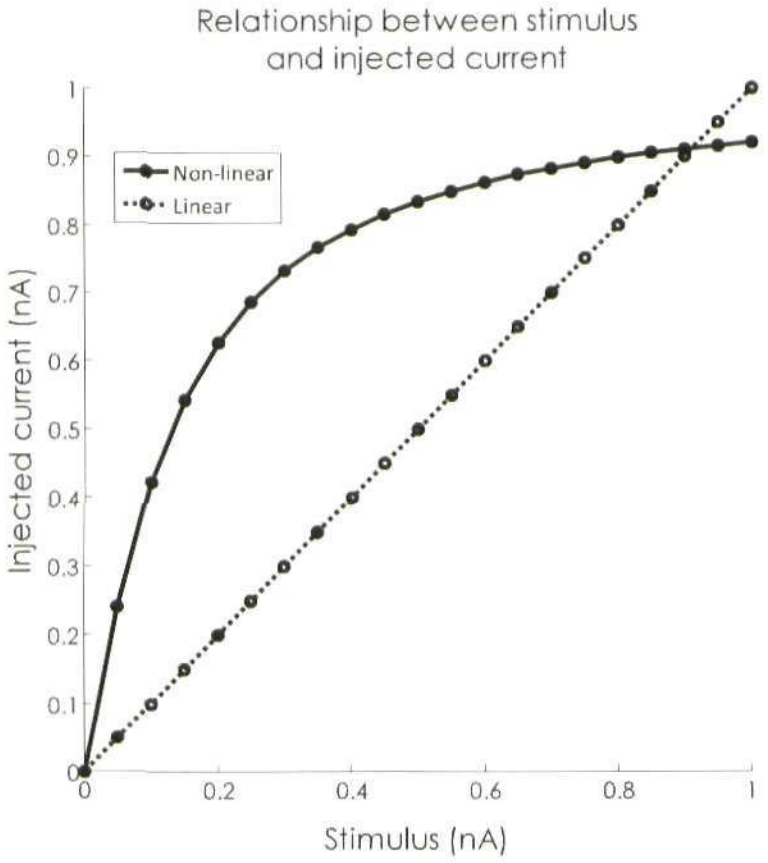

Figure 5.19: Relationship between stimulus and injected current, under linear/non-linear conditions
Solid line/circles gives the actual injected current under the non-linear condition (calculated by equation 5.11) and dashed line/empty circles plots the linear condition.

Where s_{Inj} is the injected current, I_{MAX} is the maximum current injection and was set

to 1nA, s is the stimulus, n governs the slope of the function and was set to 1.2 and S_{50} is the half-maximal constant and was set to 13% of I_{MAX} (the equations and parameter values were derived from Murphy & Miller [2003]). The relationship between the stimulus and the current injected into the model cell, under both linear and non-linear conditions, is plotted in Figure 5.19.

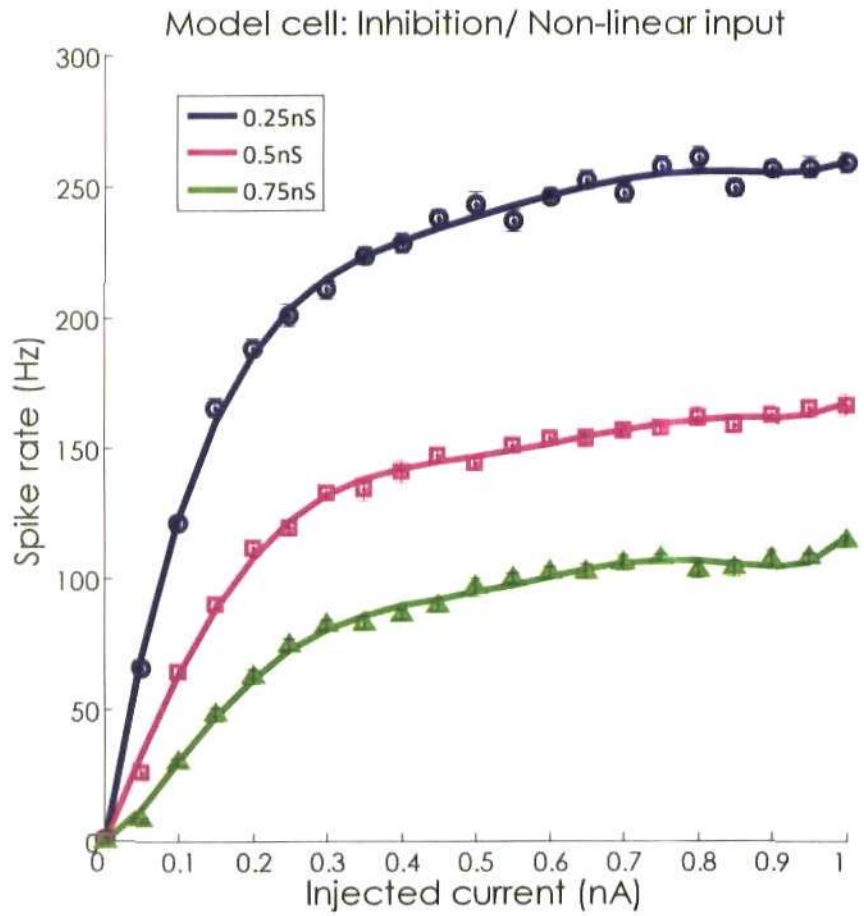

Figure 5.20: Response of the model cell to non-linear stimulus, with tonic inhibition
Blue lines/circles plot the response in the presence of tonic inhibition of 0.25nS, magenta line/squares plot the 0.5nS condition and green line/triangles plots the response for the 0.75nS condition. Solid lines give the best-fit to the data (symbols).

The mean output of the model cell (averaged over 100 trials) under conditions of

increasing tonic inhibition for both non-linear and linear input are plotted in Figure 5.20 & Figure 5.21, respectively.

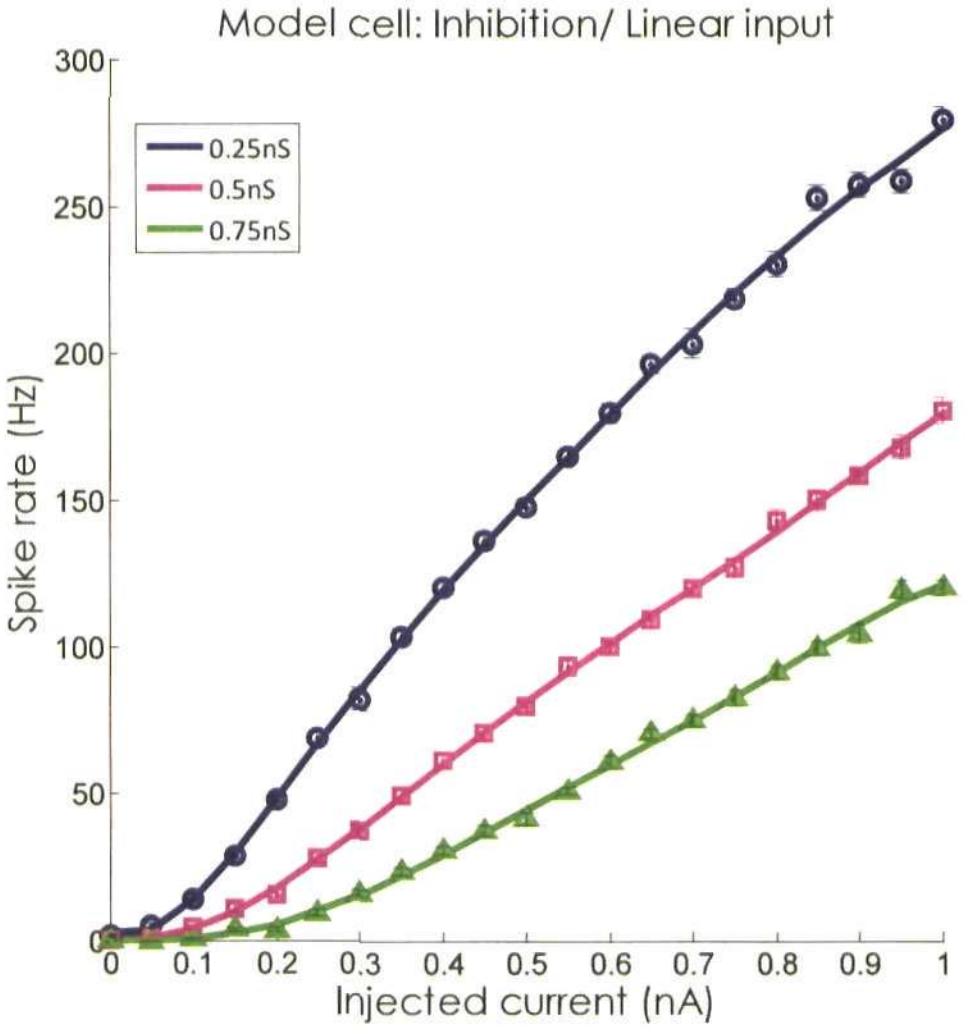

Figure 5.21: Response of the model cell to a linear stimulus, with tonic inhibition
Blue lines/circles plot the response in the presence of tonic inhibition of 0.25nS, magenta line/squares plot the 0.5nS condition and green line/triangles plots the response for the 0.75nS condition. Solid lines plot the best-fit to the data (symbols).

Table 5.3: Amplitude and slope of model stimulus-response function, for linear and non-linear input, with tonic inhibition

	Non-linear			*Linear*		
Inhibition (nS)	*0.25*	*0.5*	*0.75*	*0.0*	*0.25*	*0.5*
S_{50} *(nA)*	0.1130	0.1464	0.1868	0.4636	0.5386	0.6154
(±SE)	(0.003)	(0.003)	(0.003)	(0.007)	(0.007)	(0.009)
Slope (Hz/nA)	0.0391	0.0242	0.0144	0.0149	0.0104	0.0077
(±SE)	(0.017)	(0.009)	(0.005)	(0.004)	(0.004)	(0.004)

Overall, the gain of the model cell, under both linear and non-linear conditions, decreased as tonic inhibition was strengthened. However, under the linear condition (Figure 5.21) the model neural rate-function also underwent a lateral displacement, as evinced by the increase in stimulus amplitude, at the half-maximal point, from 0.46nA-0.62nA, following an increase in tonic inhibition from 0.25nS to 0.75nS, respectively. This is similar to the adaptive response observed during adaptation to stimulus mean

5.3.3: Balance of synaptic depression

In order to explore how asymmetric synaptic depression affects response gain, a recent modelling study (Chelaru & Dragoi 2008) developed a simple population model in which two populations of excitatory (E) and inhibitory (I) neurons were recurrently connected (Figure 5.22, populations within dashed box). Both populations received feedforward input (F_E, F_I, respectively), with the inhibitory population receiving feedforward input that was proportional to excitatory inputs (governed by the constant k, see equations 5.12-5.13).

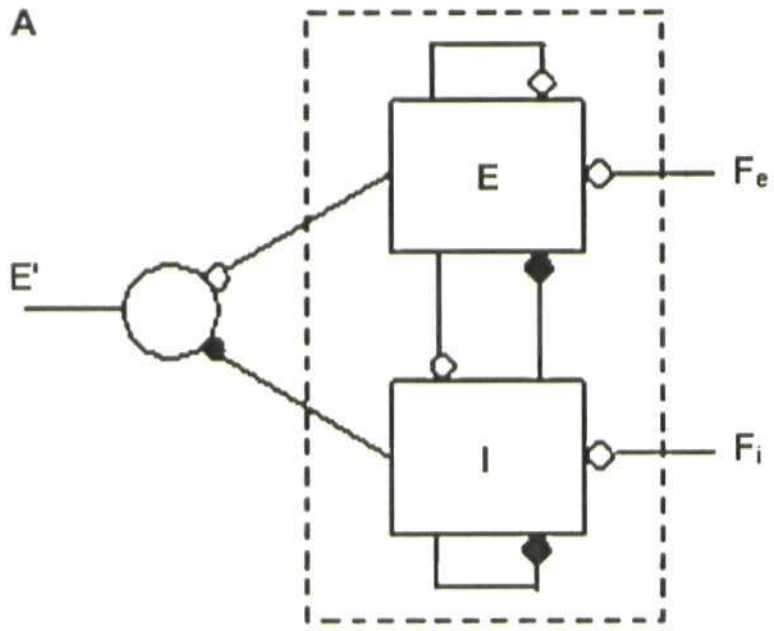

Figure 5.22: Schematic diagram of a recurrently connected population of excitatory and inhibitory cells.
Dashed line and connectors with empty symbols represent the recurrent population and connections, respectively, whilst filled connector symbols represent afferent, feedforward connections. Adapted from Chelaru & Dragoi (2008).

The output of each population is the product of the gain and the feedforward input:

$$
\begin{aligned}
E &= G_e F_e \\
I &= G_i F_i \\
F_i &= kF \\
G_e &= \frac{1 + W_{ii} - kW_{ei}}{D} = \frac{1 + a}{D} \\
G_i &= \frac{1 + W_{ie} - k(W_{ee} - 1)}{kD} = \frac{b + k}{kD} \\
a &= W_{ii} - kW \\
b &= W_{ie} - kW \\
D &= W_{ei} W_{ie} - (W_{ee} - 1)(W_{ii} + 1)
\end{aligned}
$$

5.12

Where E and I are the responses of the excitatory and inhibitory populations, respectively,

to the feedforward input F; G is the gain of the population, e and i indicate the polarity of the population (with the 1st letter indicating the polarity of the post-synaptic population) and W gives the synaptic weight of the recurrent connections.

In the presence of synaptic depression the responses of the recurrent populations are given by (where the strength of depression between excitatory and inhibitory connections is determined by the constants α and β, respectively):

$$
\begin{aligned}
&0.3 < \alpha < \beta < 1 \\
&E = G_e F_e \\
&I = G_i F_i \\
&F_i = kF \\
&G_e = \alpha \frac{1 + \beta W_{ii} - k\beta W_{ei}}{D(\alpha, \beta)} = \alpha \frac{1 + \beta a}{D(\alpha, \beta)} \\
&G_i = \alpha \frac{\alpha W_{ie} - k(\alpha W_{ee} - 1)}{kD(\alpha, \beta)} = \alpha \frac{\alpha b + k}{kD(\alpha, \beta)} \\
&a = W_{ii} - kW_{ei} \\
&b = W_{ie} - kW_{ee} \\
&D(\alpha, \beta) = \alpha\beta W_{ei} W_{ie} - (\alpha W_{ee} - 1)(\beta W_{ii} + 1) \\
&G_e > 0;\ G_i > 0;\ D(\alpha, \beta) > 0
\end{aligned}
\qquad \textbf{5.13}
$$

Where:

$\alpha \neq \beta$ = Asymmetric depression

$\alpha = \beta$ = Symmetric depression

$\alpha = \beta = 1$ = No depression

The feedforward inputs, F_e and F_i, were also depressed by the same amount, α. Both recurrent populations provided input to an output cell, E' (equation 5.13 taken from Chelaru & Dragoi 2008).

In order to implement the population dynamics (described in equations 5.12-5.13)

in the model cell, the feedforward input, F_e, was set equal to the stimulating current and the driving input acting on the model cell was the weighted difference of the excitatory population response, F_eG_e, minus the inhibitory response, F_iG_i:

$$E' = W_e E - W_i I \qquad \textbf{5.14}$$

Using the parameters from Chelaru & Dragoi (2008):

$$W_e = 3;\ W_i = 3{:}75;\ W_{ee} = 3;\ W_{ii} = 6;\ W_{ei} = 4{:}8;\ W_{ie} = 4{:}4;\ k = 0{:}45;\ F_i = kF_e$$

Figure 5.23 plots the relationship between the output of the equations and the afferent input to the model cell as a function of stimulus amplitude, for different values of alpha. Within the model, the important factors for determining the net population response (see equations 5.12, 5.13 & 5.14) is the overall levels of the population excitation and inhibition (dotted and dashed lines in Figure 5.23, respectively) and the relative gain of the responses. For example, in the asymmetric case in which excitation depressed more strongly than inhibition (green lines/circles, Figure 5.23), excitation is relatively lower than in the non-depressed case, however inhibition is also relatively weaker and has a lower gain, thus the overall output is stronger than the undepressed case (and the other two depressed conditions).

The average output of the model cell (across 10 trails, stochasticity is introduced via balanced excitatory and inhibitory conductances, as for Section 3.1, above), under different ratios of excitatory to inhibitory depression, is plotted in Figure 5.24. One can see that the neural gain of the model cell is strongest under the asymmetric depression condition (excitatory synapses depress more than inhibitory synapses) and decreases as depression becomes more balanced; only the balanced condition elicits a response that is weaker than the no depression case.

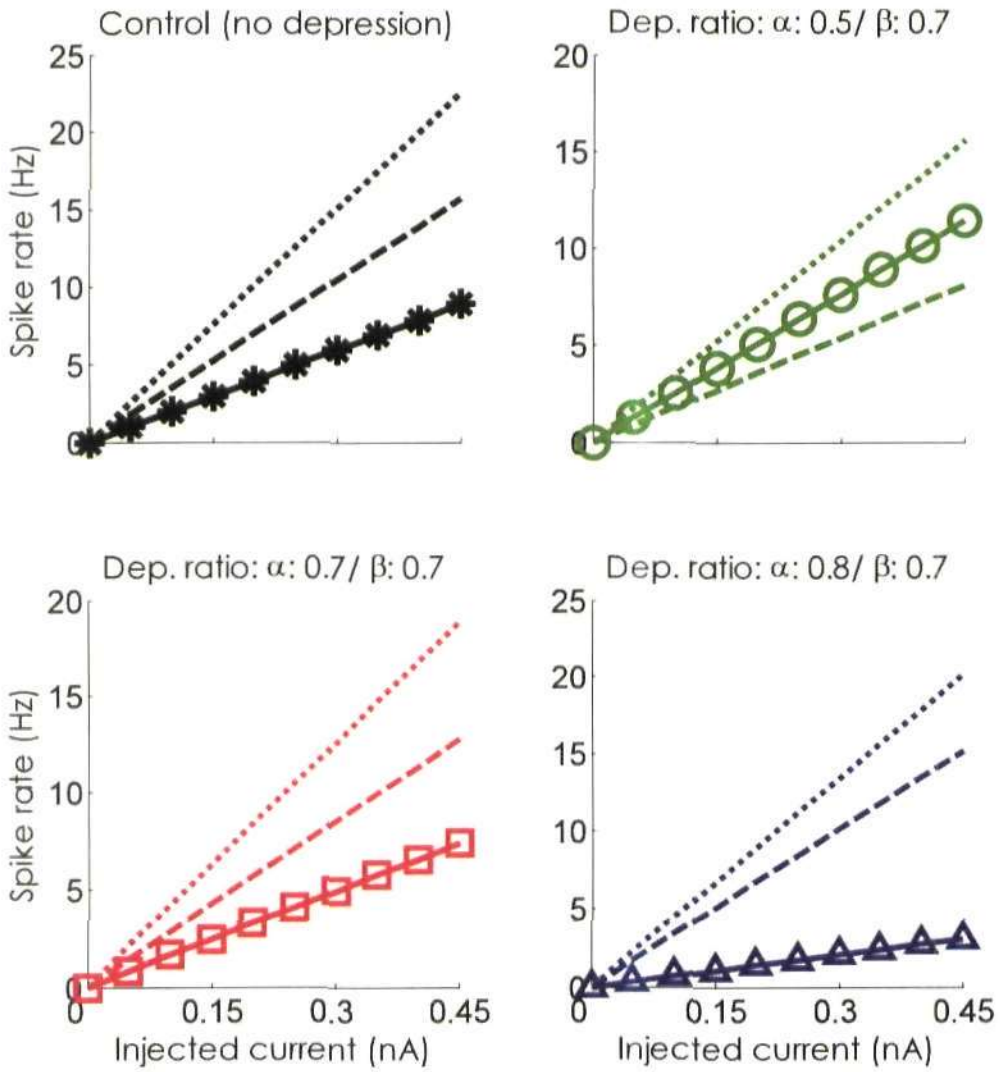

Figure 5.23: Comparison of population responses to different ratios of synaptic depression

The response as a function of current strength (see equations 5.12-5.14 for the excitatory *population (dotted lines), inhibitory population (dashed lines) and the weighted difference* of both populations (solid lines and symbols) is plotted under conditions of no depression (black lines, star symbol), balanced depression(red lines, squares) and two cases of asymmetric depression (excitatory depression > inhibitory: green lines/circles; excitatory depression < inhibitory: blue lines/triangles)

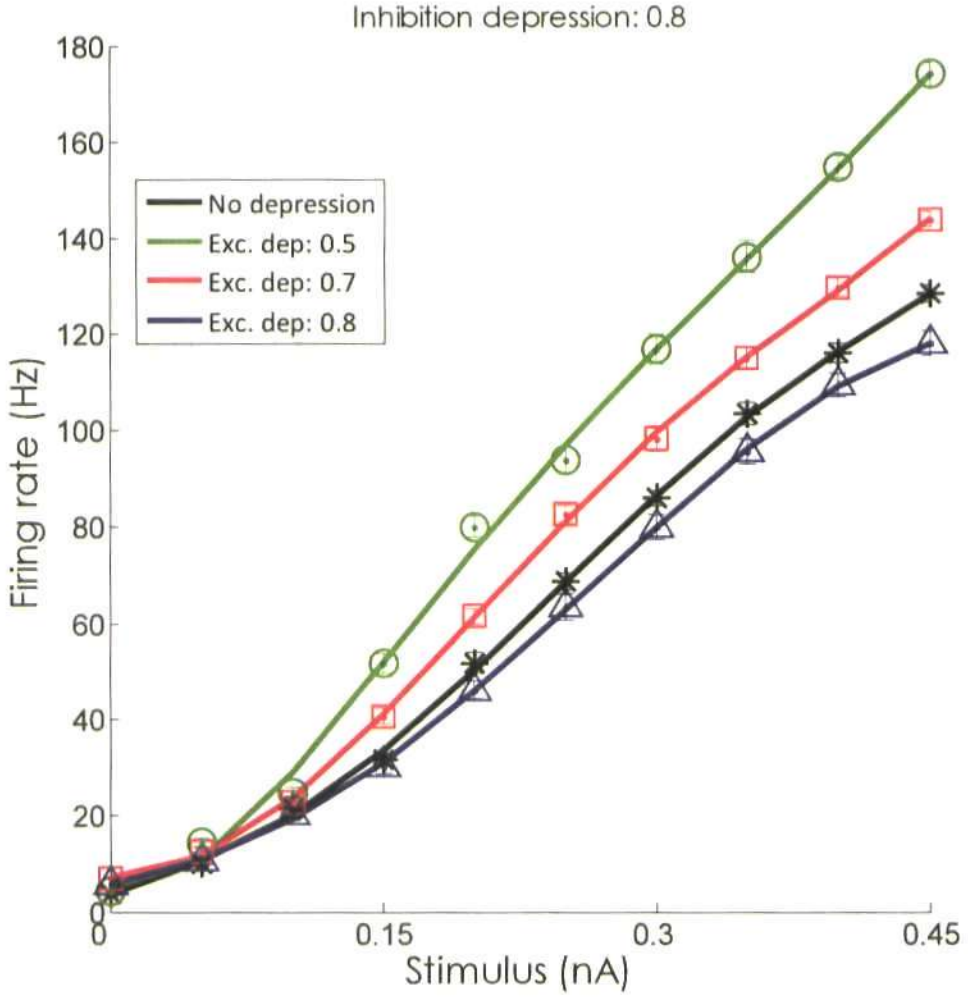

Figure 5.24: Output of the model cell under different ratios of synaptic depression
Steady state firing rate response of the model cell as a function of initial stimulus amplitude (nA) under conditions of no depression (black lines, star symbol), balanced depression(blue lines, triangles) and two cases of asymmetric depression (excitatory depression > >inhibitory: green lines/circles; excitatory depression >inhibitory: red lines/squares). The actual stimuli were current steps, presented for 500ms, whose amplitudes were derived from the weighted difference of the population excitatory and inhibitory steady-state responses (see equations 5.12 & 5.13) to the initial stimulus amplitude.

5.4: Modelling adaptation to stimulus statistics

As outlined in the introduction, the aim of the final part of the analysis was to

relate the models discussed in the sections above to the actual adaptation data. In general, the previous model descriptions focused on the steady-state response, to a single stimulus, presented in isolation. Here, the interest was in how the statistics of a dynamic, adapting stimulus affected the responsiveness of the model neuron and which of the experimental models best described the spike-count data.

Input to the model cell was a time varying stimulus whose amplitude was refreshed every 40ms and drawn from a highly skewed distribution either in the form of direct current injection or by converting the current into firing rate with which to drive a Poisson spike generating mechanism.

5.4.1: Adaptation to stimulus statistics: Spike-rate adaptation

A reminder of the experimental data is presented in Figure 5.25, which plots the population firing-rate response as a function of stimulus amplitude (Figure 5.25, top left), step-size (Figure 5.25, top right) and deviation from the mean (Figure 5.25, bottom).

The pulse-based model cell was tested for adaptation to global mean without (fixed threshold, see Figure 5.26) spike-rate adaptation (SRA) and with (dynamic threshold, see Figure 5.27). To introduce some stochasticity into the simulation, stimulus presentation order was refreshed for each high-probability region trial (thus minimising presentation effects) and a small amount of noise (standard deviation of ±0.011nA) was added to the input current, at each time-step. For each high-probability region trial of the simulation, the model cell was adapted for 30s; resulting in 750 stimulus presentations per trial.

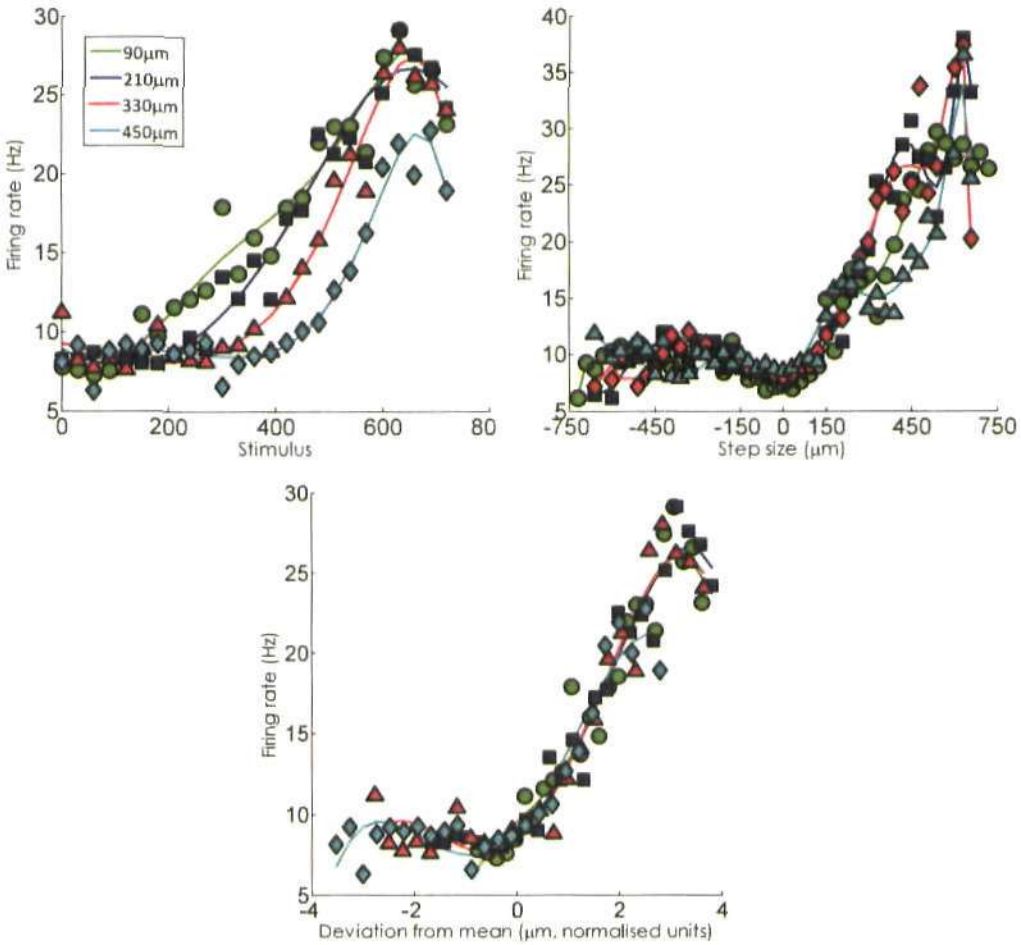

Figure 5.25: Population response to global mean; experimental data.
Top Right: Population stimulus-response function; average response plotted as a function of stimulus amplitude. **Top Left:** Step-size function; average response plotted as a function of step-size change in stimulus amplitude. **Bottom**: Deviation function; average response plotted as function of deviation of the stimulus from global mean. Green lines/circles: 90µm condition, blue lines/squares: 210µm condition, red lines/triangles: 330µm condition, cyan lines/circles: 450µm condition.

In the absence of SRA, the model cell displayed no adaptation to the global mean of the stimulus (Figure 5.26). The non-monotonic nature of the stimulus-response function, for stimuli greater than ≈0.2nA, was due to the limits on firing rate imposed on the model. Each spike-pulse was 0.6ms in length, and the refractory period was set to 2ms, thus limiting maximum firing rate to ≈400Hz.

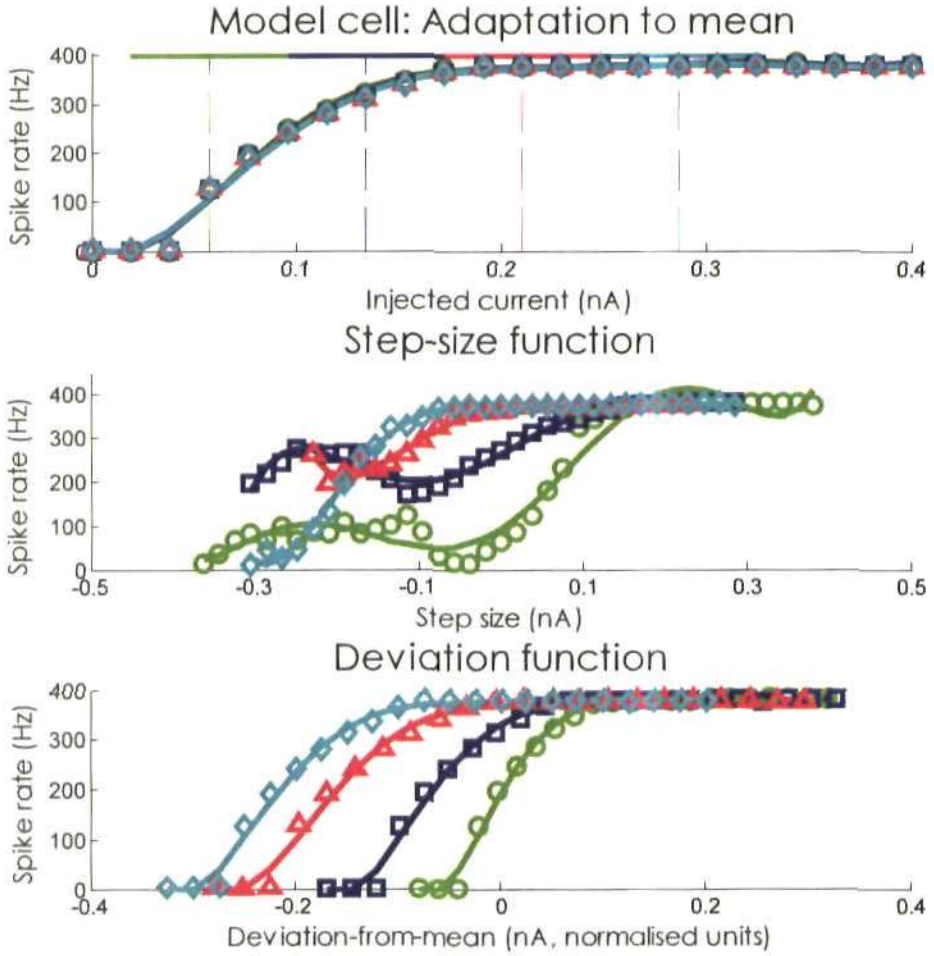

Figure 5.26: The model cell's firing rate, step-size and deviation-from-the-mean functions following adaptation to four global mean levels, with no spike-rate *adaptation*

Top: Stimulus-response function. Coloured, horizontal lines at the upper limits of the axis span the extent of the high-probability region, for each mean adapting level, whilst vertical, dashed lines plot the central amplitude of the corresponding high-probability region. **Middle:** Step-size function. **Bottom**: Deviation-from-mean functions (with input normalised with respect to the standard deviation of the global mean). For all plots, solid lines represent the best-fit to the data (5th-order polynomial) and symbols represent the actual data. Green line/circles: high-probability region centred on 0.0573nA; blue line/squares: 0.1337nA; red line/triangles: 0.2101nA; cyan line/diamonds: 0.265nA. (Threshold was -48mV to allow for comparison when spike-rate adaptation was present, [as the average, dynamic, threshold for the lowest mean amplitude condition, with spike-rate adaptation, was 45mV, see below])

However in the presence of a dynamic threshold the model cell displayed the characteristic shift in the rate-function observed experimentally (Figure 5.27). The average (across 100 trials) half-maximal amplitudes, or threshold (S_{50}), increased from 0.131nA at the lowest mean amplitude (equivalent to the 90μm condition) to 0.191nA at the highest mean amplitude condition (equivalent to 450μm) whilst the slope of the stimulus-response function, at the S_{50}, ranged from 0.039Hz/nA^2 to 0.032 Hz/nA^2 across the four mean levels tested (thus slope decreased slightly as global mean increased); the values are given in Table 5.4.

Incorporating an adaptive threshold into the model generated a shift in stimulus-response function that was dependent on the overall, mean injected current amplitude, and was in agreement with the findings of Garcia-Lazaro and colleagues (Garcia-Lazaro et al. 2007). However, the lateral displacement of the stimulus-response function was not sufficient to generate an invariant step-size or deviation-from-the-mean function (when normalised by dividing the input by the global standard deviation of the stimulus), as observed experimentally (see Figure 5.25, Chapter 4 and Dean et al. 2005). If one examines the model cell's rate-function (see Figure 5.27, top) one can see that the high-probability region of the stimulus space was covered by the linear region of the corresponding stimulus-response function, with the exception of the lowest mean amplitude condition (see Figure 5.27, top; green lines/symbols). Thus the displacement of the stimulus-response, due to spike-rate adaptation, was not sufficient to account for all features of adaptation to global mean.

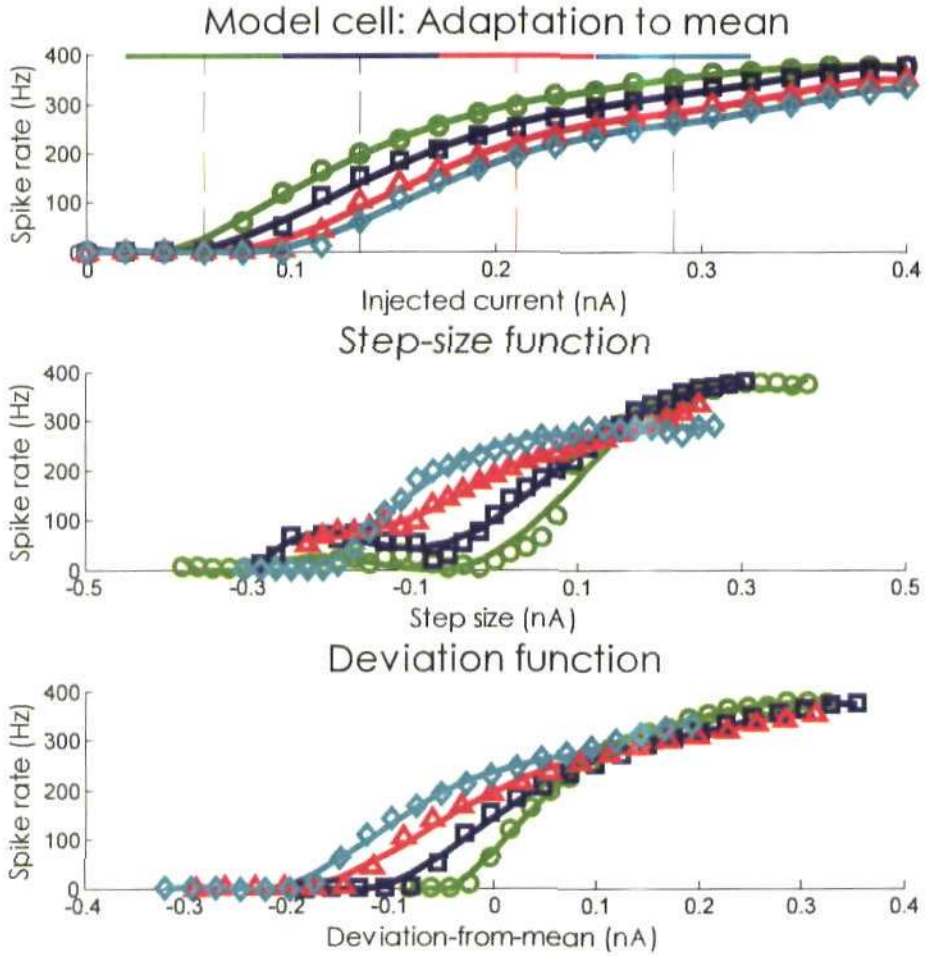

Figure 5.27: The model cell's firing rate, step-size and deviation functions following adaptation to four global mean levels, in the presence of spike-rate adaptation (adaptation constant, b, was set to 1)

Top: The stimulus-response function of the model cell. Coloured, horizontal lines at the upper limits of the axis span the extent of the high-probability region, for each mean adapting level, whilst vertical, dashed lines plot the central amplitude of the corresponding high-probability region. **Middle:** Step-size function of the model cell. **Bottom:** Deviation-from-mean function (normalised). For both plots, solid lines represent the best-fit data (5th-order polynomial) and symbols represent the actual data. Colours/symbols as for Figure 5.26, above

As discussed in Chapter 4, the rough invariance of the step-size function, observed *in vivo*, arises due the displacement of the rate-function onto stimulus intensity levels outside the high-probability region of the stimulus distribution. However, whilst increasing the amount of spike rate-adaptation (by increasing the adaptation constant from 1 to 4) did increase threshold (relative to the original simulation, above, compare Table 5.4 & Table 5.5) it was not sufficient to render the response to both relative amplitude (see Figure 5.28, middle) and normalised deviation-from-the-mean invariant (see Figure 5.28, bottom).

Table 5.4: Slope and current amplitude at the half-maximal point following adaptation to mean

Mean Stimulus (nA)	**0.0573**	**0.1337**	**0.2101**	**0.2650**
S_{50} *(nA)*	0.131	0.160	0.180	0.191
(± SE)	(0.0094)	(0.0053)	(0.0046)	(0.0084)
Slope (Hz/nA)	0.0387	0.0332	0.0330	0.0321
(± SE)	(0.0093)	(0.0024)	(0.0036)	(0.0077)

Table 5.5: Slope and current amplitude at the half-maximal point following adaptation to mean, with SRA (b = 4)

Mean Stimulus (nA)	**0.0573**	**0.1337**	**0.2101**	**0.2650**
S_{50} *(nA)*	0.200	0.217	0.250	0.272
(± SE)	(0.0041)	(0.0018)	(0.0013)	(0.0041)
Slope (Hz/nA)	0.0215	0.0219	0.0203	0.0227
(± SE)	(0.0046)	(0.0042)	(0.0051)	(0.0027)

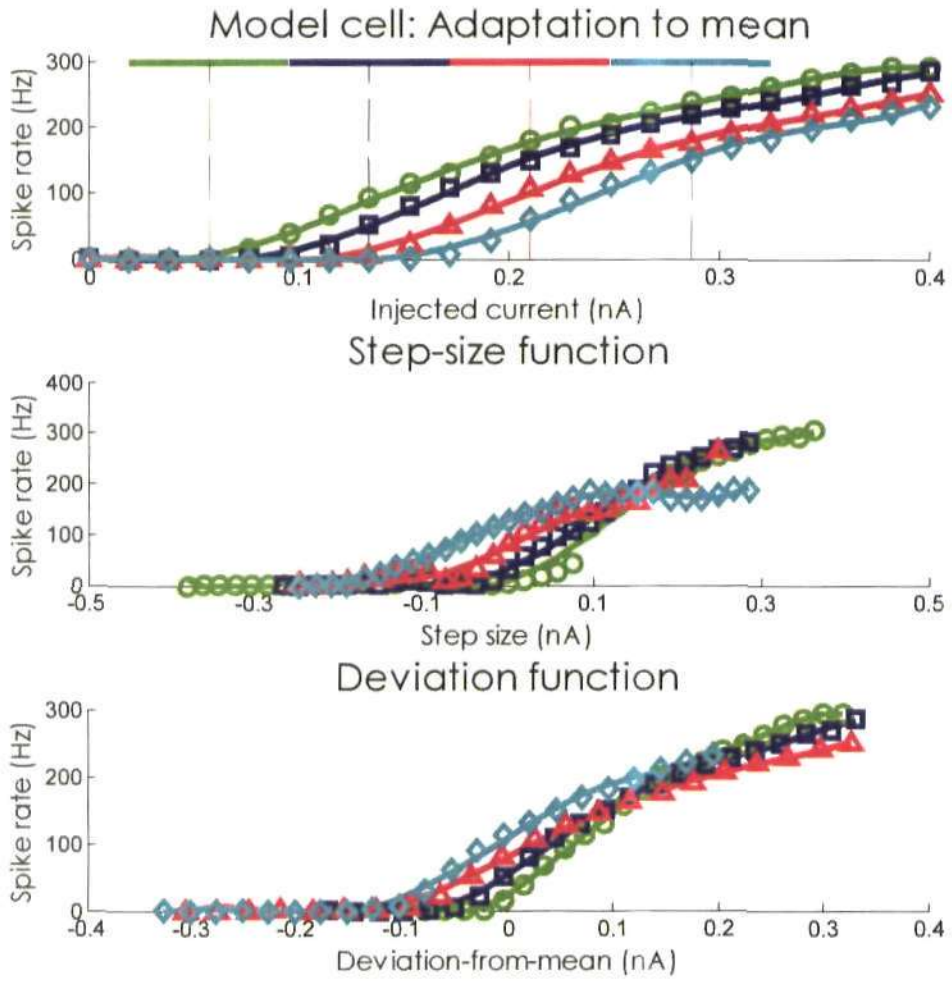

Figure 5.28: The model cell's firing rate, step-size and deviation functions following adaptation to four global mean levels, in the presence of spike-rate adaptation (adaptation constant, b, was set to 4

Left: Stimulus-response function of the model cell. Coloured, horizontal lines at the upper limits of the axis span the extent of the high-probability region, for each mean adapting level, whilst vertical, dashed lines plot the central amplitude of the corresponding high-probability region. **Middle:** Step-size function of the model cell. **Bottom:** Deviation-from-mean function (normalised). For both plots, solid lines represent the best-fit data (5th-order polynomial) and symbols represent the actual data. Colours/symbols as for Figure 5.26, above

Thus it was decided to incorporate tonic inhibition into the model (Murphy & Miller 2003). Tonic inhibition, in this case, would act as form of high-pass filter on the stimulus, as only those stimuli strong enough to counteract afferent inhibition would elicit a response in the model cell, thereby effectively increasing threshold. Tonic inhibition was simulated by introducing a constant conductance into the cell, whose amplitude was determined by the global mean of the stimulus, at each trial (i.e. on average, across 100 trials, the inhibitory conductance (nS) was set to: 0.0920 (±0.0008 Standard Error or SE), 0.1515 (±0.0007SE), 0.2085 (±0.0006SE), 0.2652 (±0.0011SE), for the four mean adapting high-probability region).

The introduction of tonic inhibition into the model cell, in the presence of spike-rate adaptation, resulted in rate-functions that were displaced rightwards, in a stimulus-mean dependent manner (see Figure 5.29, top plot). The stimulus amplitude that elicited the half-maximal response ranged, on average, from 0.169nA (±0.0015SE) to 0.297nA (±0.0008SE), from the lowest to highest mean-stimulus adapting condition; slopes at the half-maximal point ranged from 0.043 (±0.0011SE) to 0.046Hz/nA (±0.0057SE). Both the step-size and normalised-deviation functions displayed similar levels of invariance to that observed experimentally (see Figure 5.29, middle and bottom plots, respectively). The model cell was also simulated under a condition of tonic inhibition but no spike-rate adaptation (see Figure 5.30). Tonic inhibition alone did not elicit the level of stimulus-response function displacement required to generate the invariant step-size and deviation functions plotted in Figure 5.29 above.

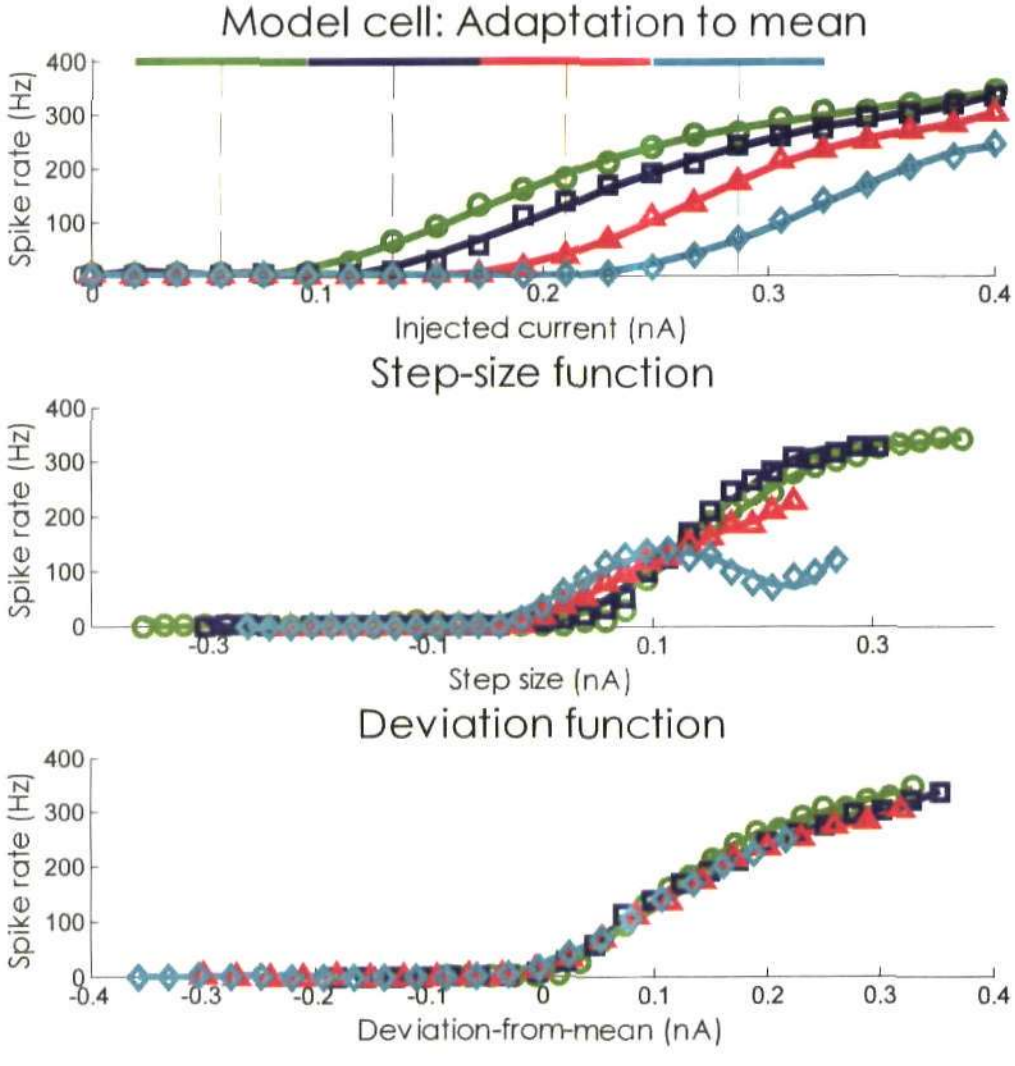

Figure 5.29: The model cell's firing rate and step-size functions following adaptation to four global mean levels, in the presence of spike-rate adaptation and tonic inhibition.

Left: Stimulus-response function of the model cell. Coloured, horizontal lines at the upper limits of the axis span the extent of the high-probability region, for each mean adapting level, whilst vertical, dashed lines plot the central amplitude of the corresponding high-probability region. **Middle:** Step-size function of the model cell. **Bottom:** Deviation-from-mean function (normalised). For both plots, solid lines represent the best-fit data (5th-order polynomial) and symbols represent the actual data. Colours/symbols as for Figure 5.26, above.

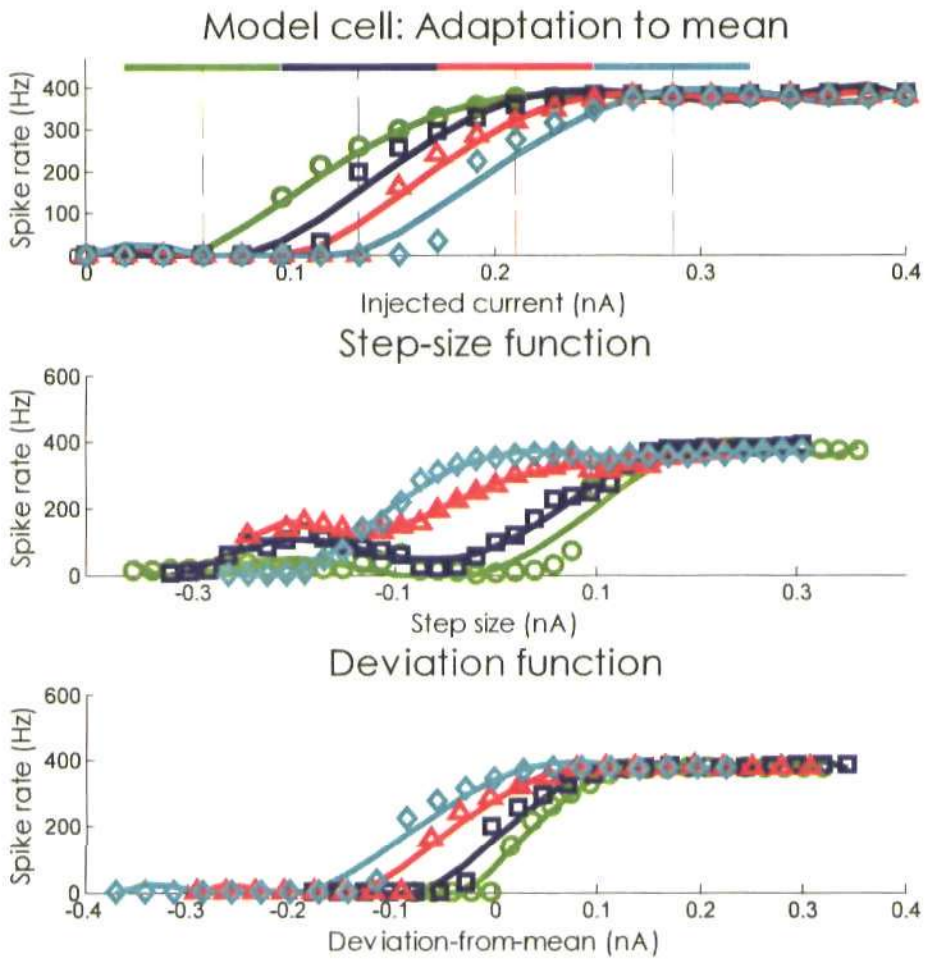

Figure 5.30: The model cell's firing rate and step-size functions following adaptation to four global mean levels, in the presence of tonic inhibition, alone.
Left: Stimulus-response function of the model cell. Coloured, horizontal lines at the upper limits of the axis span the extent of the high-probability region, for each mean adapting level, whilst vertical, dashed lines plot the central amplitude of the corresponding high-probability region. **Middle:** Step-size function of the model cell. **Bottom:** Deviation-from-mean function (normalised). For both plots, solid lines represent the best-fit data (5th-order polynomial) and symbols represent the actual data. Colours/symbols as for Figure 5.26, above

Thus, within the model cell at least, a combination of global-mean dependent tonic

inhibition and spike-rate adaptation was necessary to replicate the experimental data. As mentioned above, tonic inhibition acts by effectively increasing the spike-threshold. As stimulus-mean (and therefore tonic inhibition) increased, there was a concomitant increase in the intensity of stimulus required to elicit a spike. Spike-rate adaptation was therefore only affecting the dynamic membrane spike-threshold for increasingly higher stimulus amplitudes, thereby resulting in a reduction in firing rate and further displacement of the stimulus-response function.

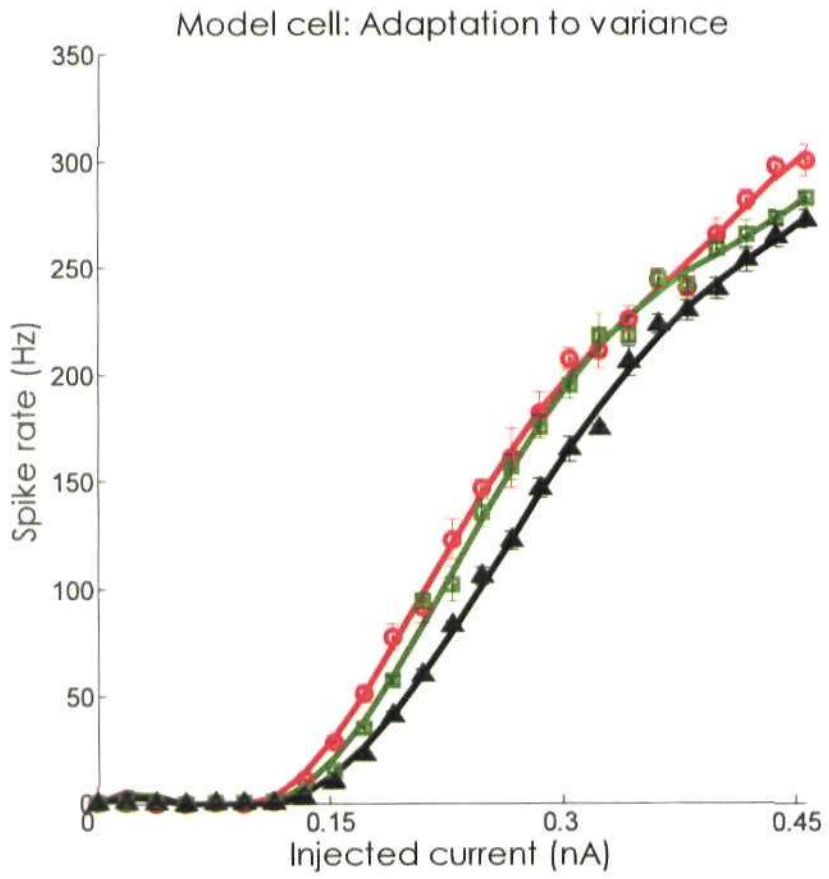

Figure 5.31: Adaptation to variance, in the model cell, with SRA and tonic inhibition Solid lines represent the best-fit data (5th-order polynomial) and symbols represent the actual data. Red line/circles represent the response to the low variance stimulus, green line/squares plots the mid-variance response and black line/triangle is the high-variance condition.

A further question to ask was whether the combination of spike-rate adaptation and tonic inhibition could generate the gain modulation observed in response to increased stimulus variance, but this was found not to be the case (see Figure 5.31), instead the stimulus response function was displaced slightly. Tonic inhibition increased in line with the mean of the stimulus and had an average value of 0.1548nS, for the low-variance condition, 0.1595nS for the mid- and 0.1669nS for the high-variance condition.

Fisher Information

The Fisher Information functions of the stimulus-response functions plotted in Figure 5.29 were calculated and plotted in Figure 5.32.

Fisher Information for the first 3 adapting conditions (green, blue and red lines, Figure 5.32) corresponded well with the experimental data in that the Fisher functions tended to peak at stimulus amplitudes that were of greater intensity than the corresponding high probability regions of the stimulus space. For the final adapting condition (cyan line, Figure 5.32), the peak of the Fisher Information was located within the high-probability region. Clearly the degree of displacement of the stimulus-response function under this condition was not as strong as that observed experimentally, and could account for why the step-size function does not match up more closely with step-size functions for the lower, global-mean adapting conditions (see middle plot, Figure 5.29).

5.4.2: Adaptation to variance can be simulated by increasing background input in proportion with levels of stimulus variance

In order to replicate the experimental data it was decided to simulate the background-noise model of gain control as proposed by Frances Chance and colleagues (Chance et al. 2002).

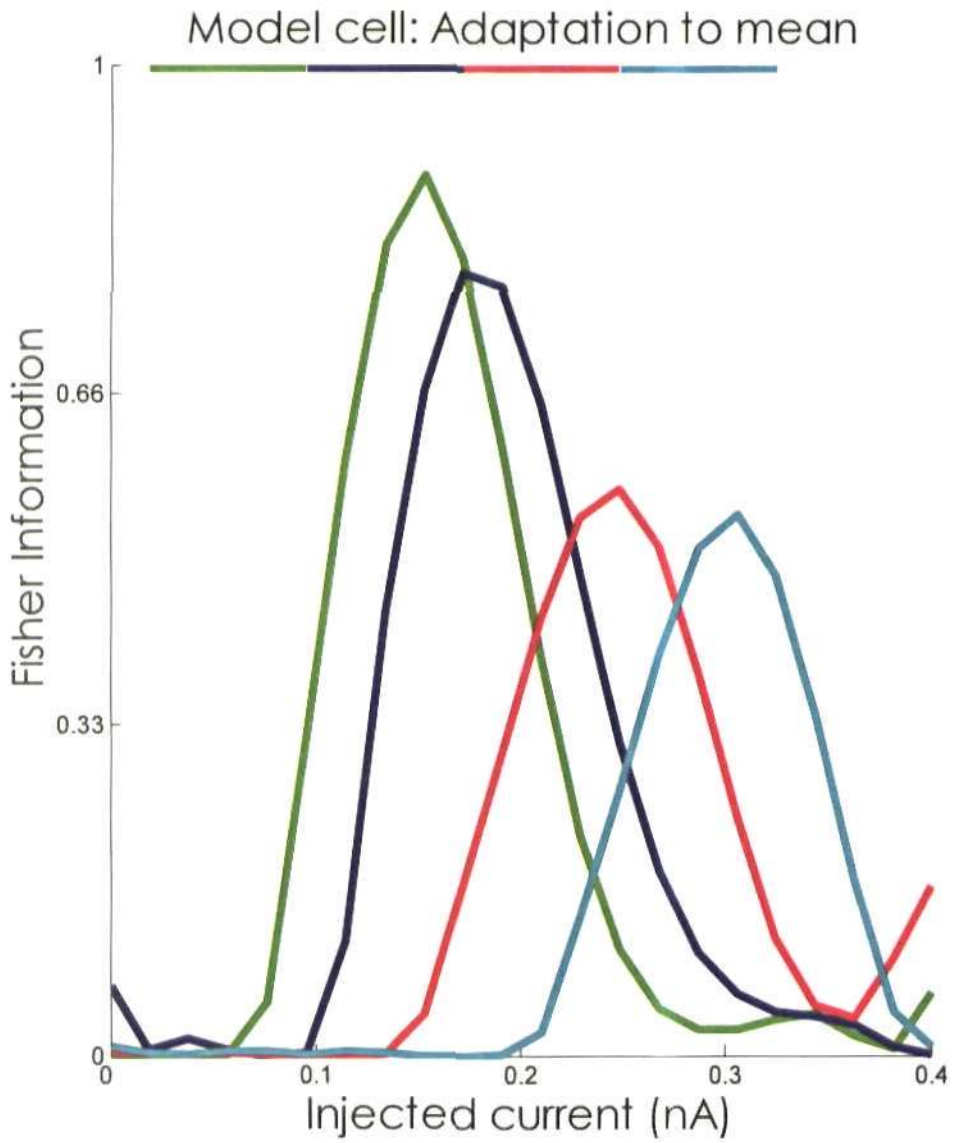

Figure 5.32: Fisher Information functions from the stimulus-response function of the adapted model cell

Coloured, horizontal lines at the upper limits of the axis span the extent of the high-probability region, for each mean adapting level. Solid lines plot Fisher Information (I) for each stimulus (s) as a function of the firing-rate function of the model cell ($f(s)$), assuming a Gaussian noise distribution: $I(s) = f'(s)^2 / \sigma^2$. Green line/circles: high-probability region centred on 0.0573nA; blue line/squares: 0.1337nA; red line/triangles: 0.2101nA; cyan line/diamonds: 0.265nA.

The experimental population response to increased stimulus-variance is re-plotted in *Figure 5.33*, below.

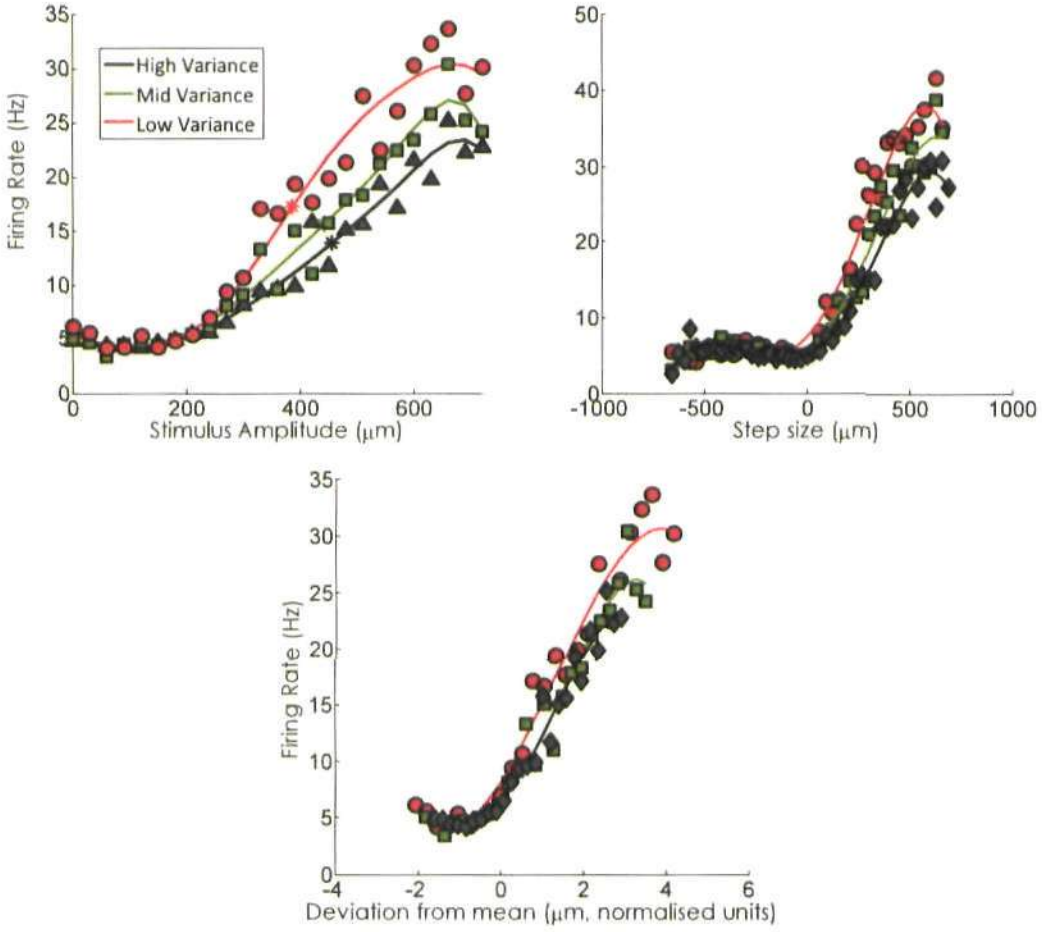

Figure 5.33: Population responses from adaptation to variance experiments
Top Left: Population stimulus-response function (SRF) in response to increasing levels of adapting stimulus variance; the slope, or gain, of the SRF decreased with every increase in global variance. **Top Right:** Population step-size function; increasing variance resulted in slight lateral displacement of the function. **Bottom:** Population, normalised, deviation function; the response as a function of local deviation-from-the-mean was not fully invariant (especially with reference to the tails of the functions), for all global stimulus-variance levels. Red lines/circles plot the low-variance condition, green lines/squares plot mid-variance and black lines/diamonds plot the high-variance condition.

As mentioned earlier (see Section 5.1.3), the simplest method of relating the background-noise model to the adaptation data is to assume that as variance increases the background firing rates also increase. This was simulated by incorporating random, background, excitatory and inhibitory modulatory inputs into the pulse-based neural model whose firing rates increased in proportion to stimulus variance (see Figure 5.34 for a schematic description of the model).

The driving input into the model cell was as for Section 5.2.3, above, in that the stimulus was an injected, constant current, with an amplitude that was refreshed every 40ms and drawn from a non-uniform distribution, centred on a high-probability region. For each successive increase in the width of the high-probability region the firing rates for both stochastic inhibition and excitation were also increased, such that for the low-variance condition, background firing rates were fixed at the base level of 50Hz (1X condition), but for the mid- to high-variance stimulus configurations, background firing rates were doubled (2X) and tripled (3X), respectively.

The average results across 100 trials are plotted in Figure 5.35, below. The I/O curves show a decrease in gain for each increase in high-probability region width and the step-size function of the model cell is not invariant with respect to stimulus variance; both these results are in line with the experimental data (see Chapter 3-4 in this thesis).

The average value (across 100 trials) of the stimulus-response function slope at the half-maximal current amplitude was 0.140nA (±0.011SE), 0.091nA (±0.008SE) and 0.084nA (±0.004SE) for the low, mid and high variance conditions respectively; the average value of the half-maximal amplitude, or threshold, was 0.255 (±0.007SE), 0.286 (±0.006SE) and 0.290Hz/nA (±0.006SE).

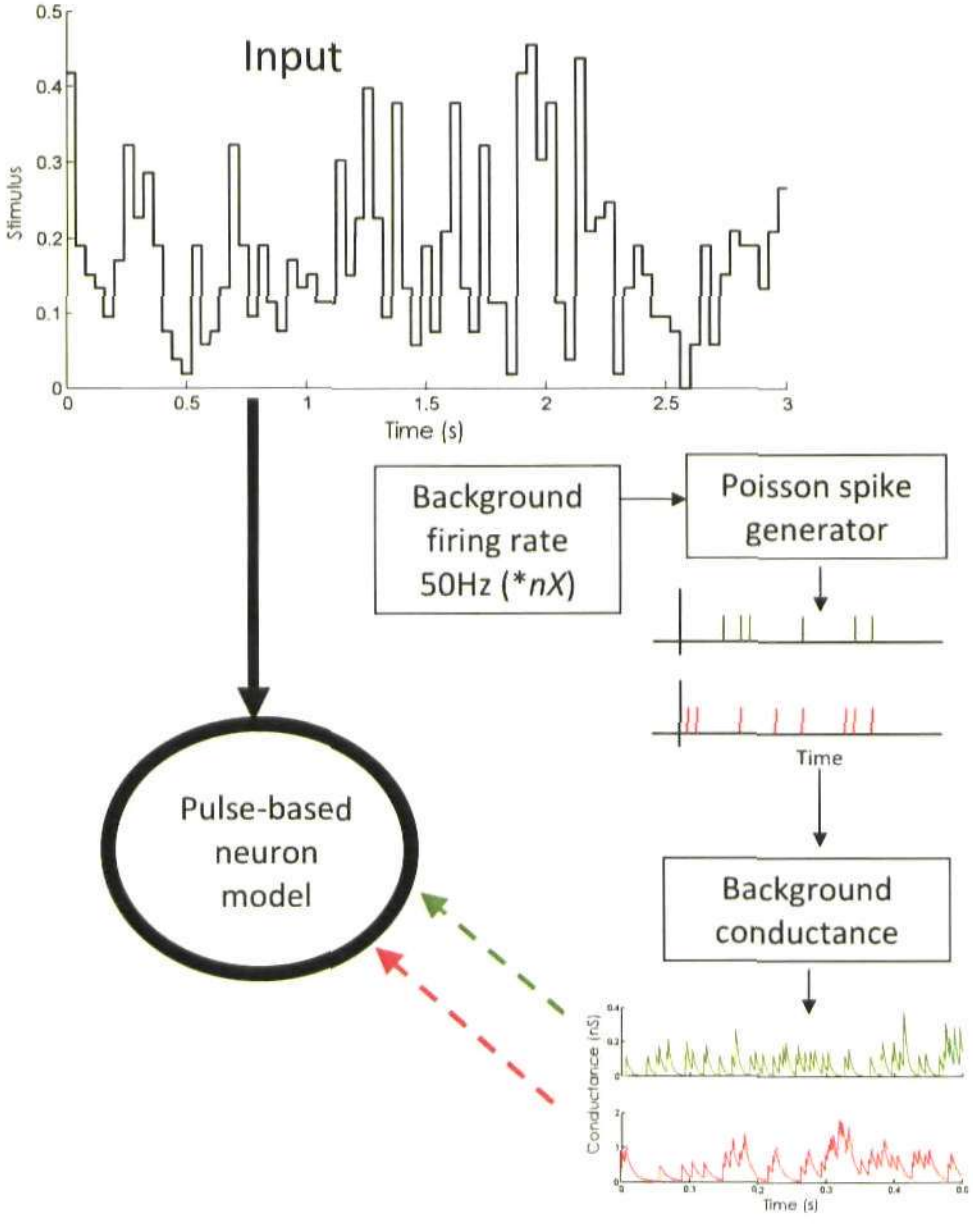

Figure 5.34: Schematic of background noise model
Driving input, represented by the thick, solid, black arrow, was a time-varying current injection (3s sample of the stimulus under the high-variance condition), where the amplitude was refreshed every 40ms. The base level of background firing rate, multiplied by the factor n (which increased with stimulus variance) was used to drive a Poisson spike-generating mechanism, from which the background (dashed arrows) conductance was derived. Green and red lines represent excitatory and inhibitory conductance, respectively.

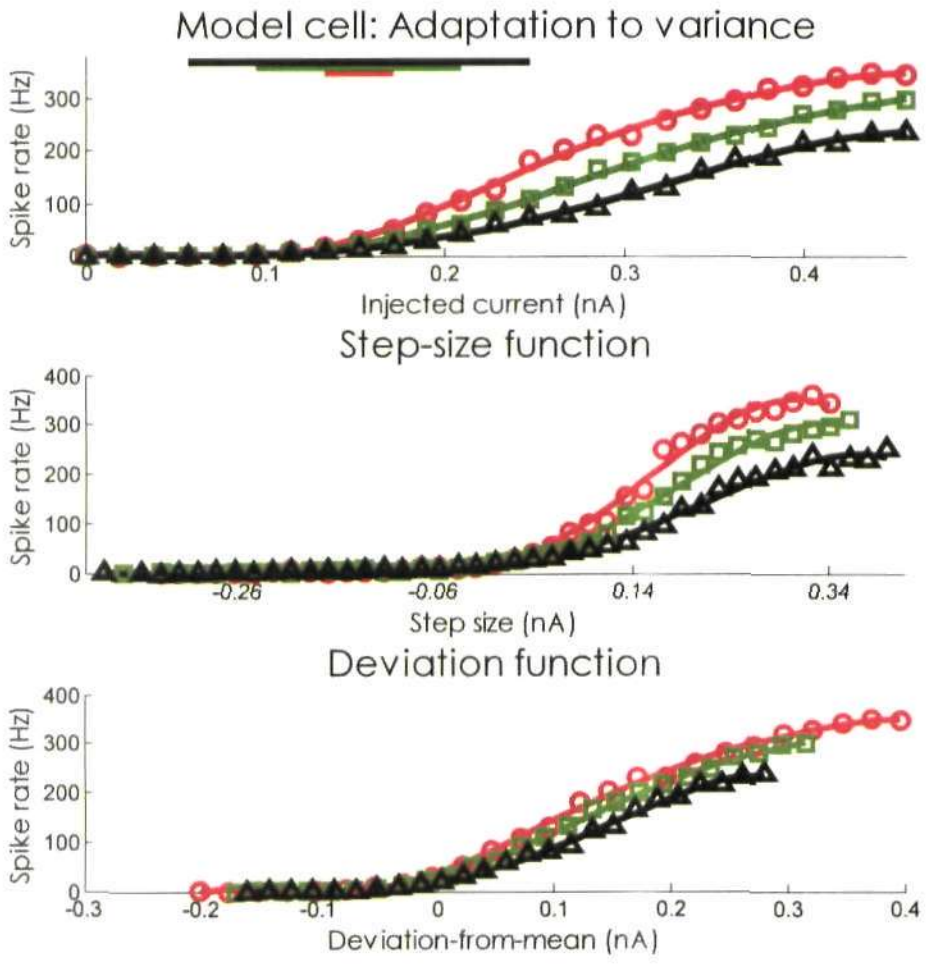

Figure 5.35: Adaptation to variance, in the model cell: background firing rates increase with stimulus variance
Left: The stimulus response function of the model cell; the widths of the high probability regions are plotted in vertical lines at the uppermost limit of the axis. **Right:** Step-size function of model cell. Solid lines represent the best-fit data (5th-order polynomial) and symbols represent the actual data; bars represent 1 standard error. Red line/circles represent the response to the low variance stimulus/1X firing rate condition, green line/squares gives the mid-variance/2X response and black line/triangle is the high-variance/3X condition.

The step-size and normalised, deviation-from-mean functions were not invariant, with respect to stimulus variance, and were in line with the experimental data (see Figure 5.35 and Chapter 4).

For the model to represent accurately the adaptive mechanisms present *in vivo*, two important assumptions were made (see Section 5.1.3): firstly, that background firing rates are dependent, to some degree, on the stimulus itself and, secondly, that background input rates increase in direct proportion to stimulus variance. Experimentally, variance levels, under the adaptation to variance paradigm, were 14,640μm^2, 19,600μm^2 and 26,900μm^2 under the low-, mid- and high variance conditions, respectively. Under the terms of the background noise model laid out above (and assuming that the low-variance condition was the base level), this would be represented by an increase in background noise 1.34 and 1.84 times the base firing rate of 50Hz, for the mid- to high-variance adapting regime, respectively. In terms of the model cell, variance levels were, on average (across 100 trials), 0.6071nA^2 (±0.0055SE), 0.8068nA^2 (±0.0064SE) & 1.0984nA^2 (±0.0063SE), respectively. Therefore further simulations were performed in which background rates were increased in proportion to levels of variance within the stimulus, at each trial, relative to the low variance condition; on average firing rates increased by 1.33 and 1.81 times the base rate of 50Hz (low-variance) for the mid- to high-variance conditions. These rates were sufficient to generate decreases in the gain of the model cell that were comparable to that observed experimentally (see Figure 5.36).

The average value (across 100 trials) of the stimulus-response function slope at the half-maximal current amplitude was 0.1422 (±0.010 standard error [SE]), 0.1398 (±0.005SE) and 0.119Hz/nA (±0.006SE) for the low (1X), mid (1.33X) and high (1.81X) variance conditions respectively (see Figure 5.36); the average value of the half-maximal amplitude, or threshold, was 0.247nA (±0.004SE), 0.260nA (±0.007SE) and 0.278nA (±0.007SE). The step-size function was not invariant, with respect to stimulus variance, and

was in line with the experimental data (see Figure 5.33 and Chapters 3-4).

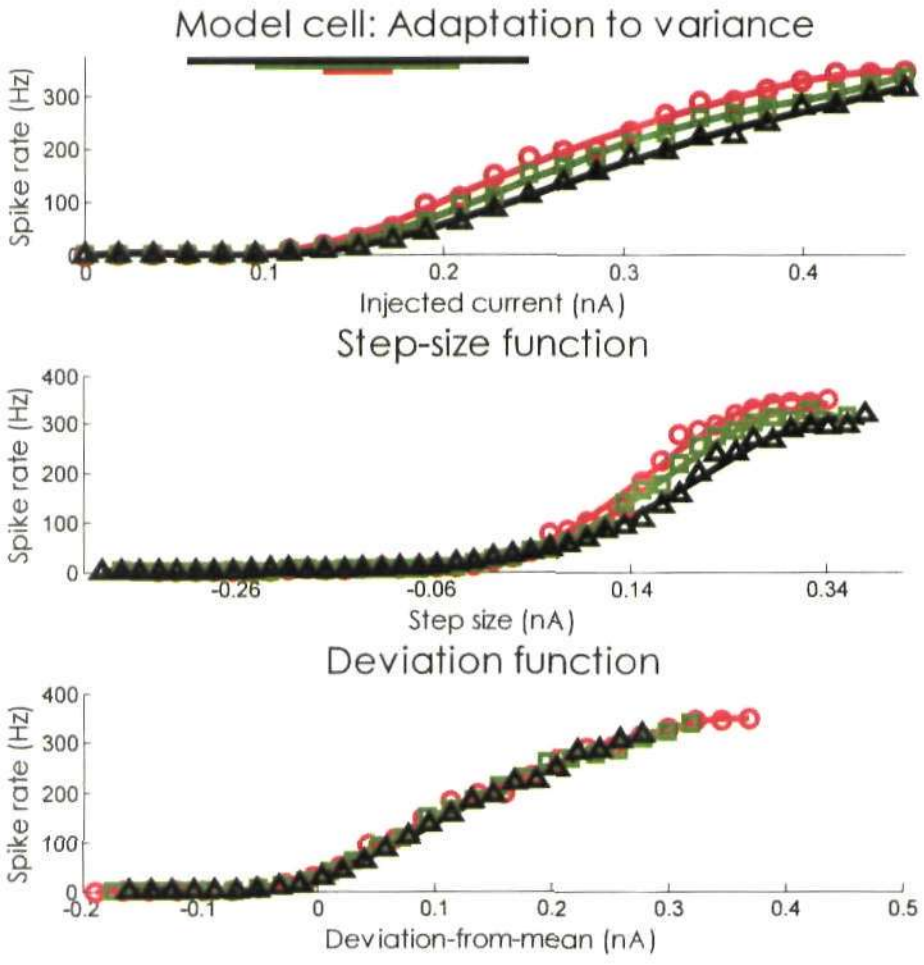

Figure 5.36: Adaptation to variance, in the model cell: background firing rates increase in exact proportion to stimulus variance
Solid lines represent the best-fit data (5th-order polynomial) and symbols represent the actual data. Red line/circles represent the response to the low variance stimulus/1X firing rate condition, green line/squares gives the mid-variance/1.33X response and black line/triangle is the high-variance/1.81X condition.

The next question to address was whether background firing rates would only increase with stimulus variance? The average output of the model cell, across all stimulus levels and trials for a given adapting condition, was calculated for all four mean (see Figure

5.37) and three variance (Figure 5.38) adapting levels. The model cell was either adapted or unadapted.

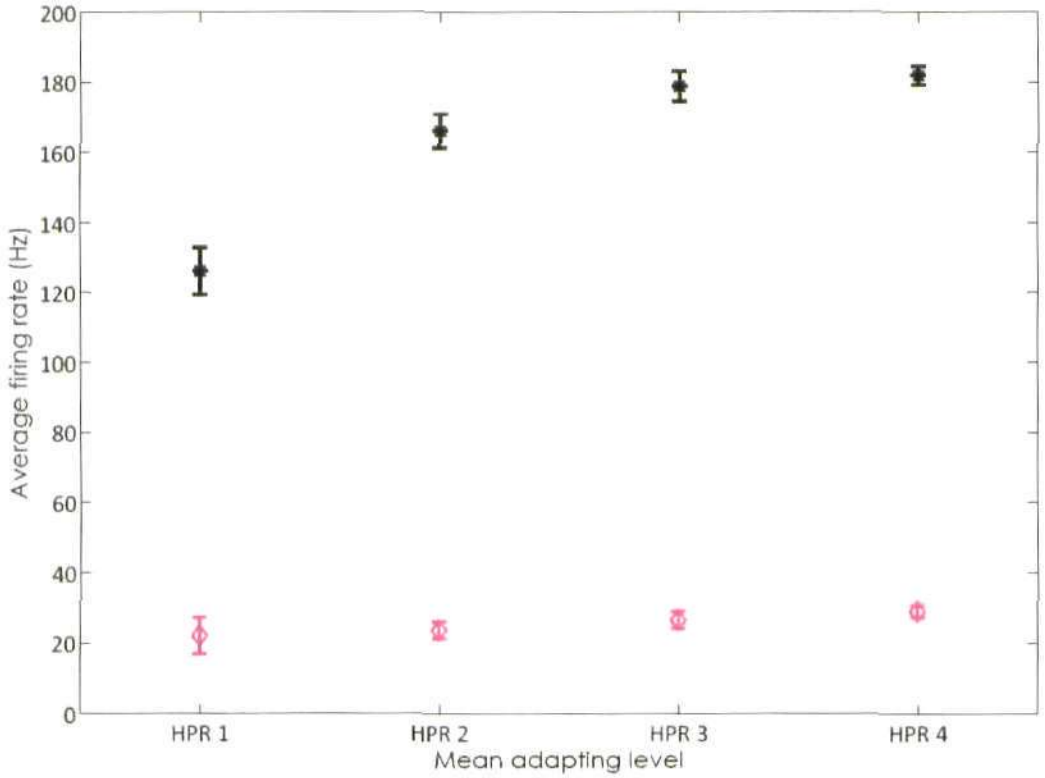

Figure 5.37: Average firing rates across all stimuli, as a function of the mean adapting level, for an adapted and unadapted model cell
Magenta diamonds plots the average firing rate for each trial (100 trials), of a given high-probability region (high-probability region), under adapting conditions (see text for details). Black asterisks plot the average firing rate of an unadapted model cell. Error bars plot 1 standard deviation from the mean.

In the case of adaptation to global mean, the average firing rate of the cell was fairly constant (Figure 5.37, magenta diamonds) across all conditions when the model cell was adapted (i.e. with a dynamic threshold and tonic inhibition). However, firing rates increased in the unadapted condition (Figure 5.37, black stars).

In the case of adaptation to global variance, the average firing rate of the cell increased when the model cell was unadapted (i.e. background firing rates were fixed at 50Hz, Figure 5.38, magenta circles). However, under adaptation, through doubling and tripling background firing rates with each successive increase in variance, overall firing rates decreased (Figure 5.38, blue squares).

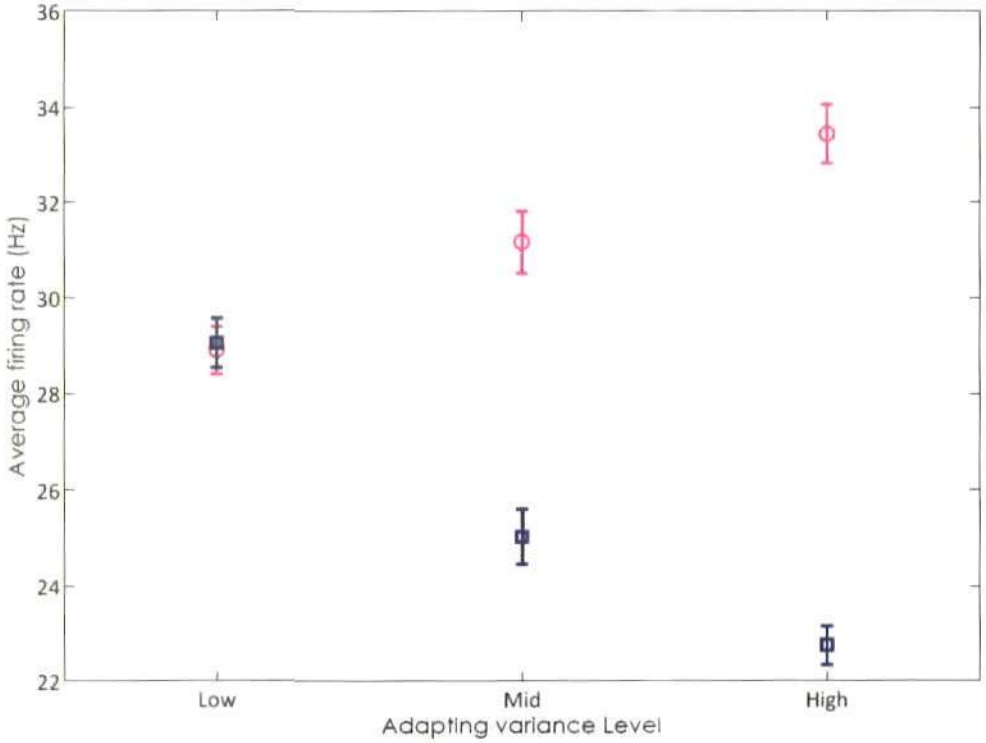

Figure 5.38: Average firing rates across all stimuli, as a function of the adapting levels of stimulus variance, for an adapted and unadapted model cell
Blue squares plot the average firing rate for each trial (100 trials), of a given variance level, under adapting conditions (see text for details). Magenta circles plot the average firing rate of an unadapted model cell. Error bars plot 1 standard deviation from the mean.

Assuming that the output of the model cell represents the typical levels of average firing rates for adapting and unadapting cells encountered *in vivo*, one can argue that to create the conditions necessary for adaptation to variance, background input to the target cell would have to arise from neurons that adapted only to global mean and not variance.

Extension to the background-noise model

It was decided to test this hypothesis by providing modulatory input to a pulse-based neuron model derived from the activity of such a population of cells. Background firing rates were approximated by running the driving stimulus, at each time step, through a population ($n = 100$) of parallel, non-linear filters. Each non-linear filter, f, was a perfect rectifier of the stimulus current, I, with an absolute threshold, θ and gain, A; threshold was set to be slightly higher than the average amplitude of the stimulus, to approximate adaptation to global-mean:

$$f(I) = [A(I - \theta)]^{+} \qquad \textbf{5.15}$$

Parameters values for threshold and gain ranged from 0.162 to 0.226nA (mean value: 0.181nA, ±0.022 standard deviation, or SD) and 1.39 to 2.88 (mean: 2.06±0.57 SD), respectively.

The relationship between the output of the filters, as a function of the stimulus, and the stimulus distribution is plotted in Figure 5.39 (top plot). All filter thresholds were slightly higher than the central amplitude of the high-probability region. The output of 3 example filters, to a 0.5s sample of the input stimulus is plotted in Figure 5.39 (bottom plot). Under the low variance regime (Figure 5.39, bottom plot, left column), the output of all three filters rarely rises above the base rate.

The output of each filter was converted to a firing rate (Hz) by multiplying it with a conversion constant, c, set to 40 (see Section 5.2.3, above); an extra term of 5 Hz was also added to the output, at each time step, to generate a source of stimulus-independent noise. This firing-rate was used to drive a Poisson spike-generator, from which the background conductance (both excitatory and inhibitory) was derived (for parameters of the background

conductance, please see Section 5.3.1 and 5.4.2).

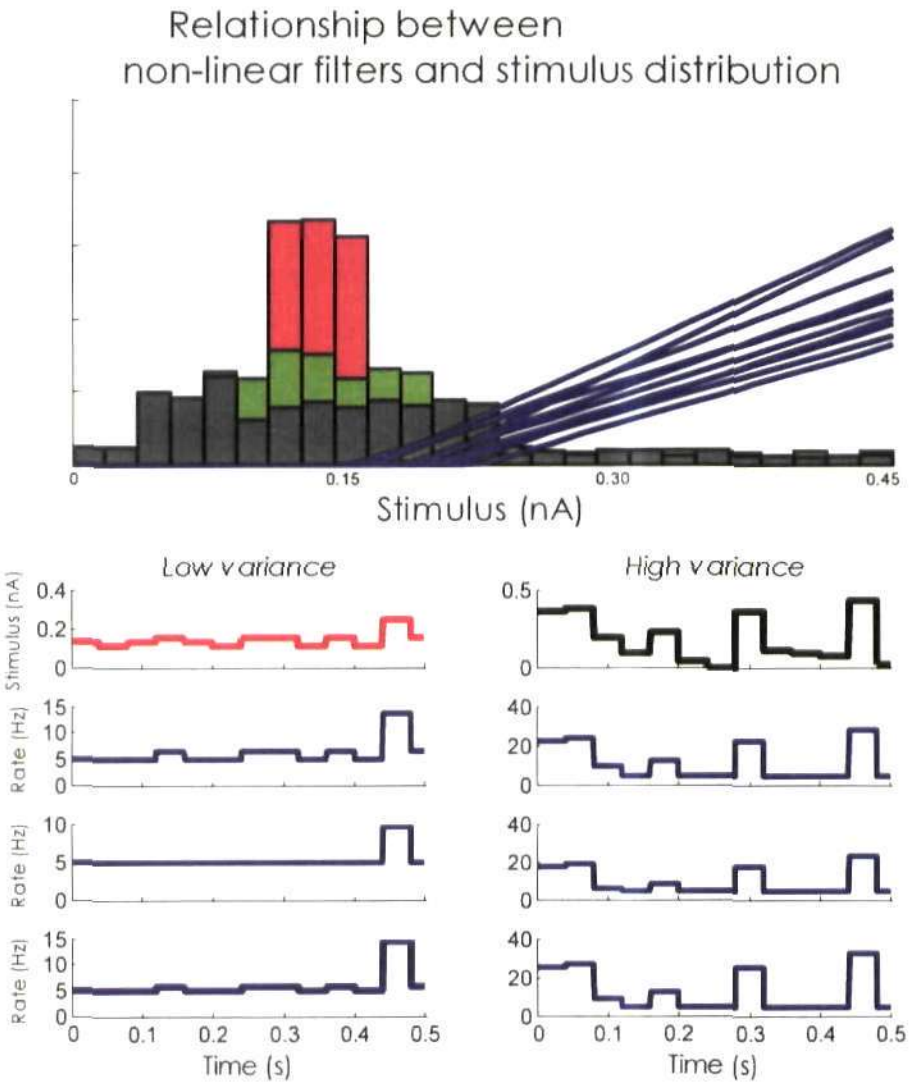

Figure 5.39: Relationship between non-linear filters of the input stimulus and the stimulus distribution

Top plot: Histogram is a frequency count of the inputs presented during a 30s stimulus sequence. Red, green and dark-grey bars show the distribution of the low-, mid- and high variance stimulus distribution, respectively. Blue lines plot the input-output relationship of 10 sample filters. All the filters have thresholds greater than the central amplitude of the high-probability region of the stimulus distribution. **Bottom plot, left column**: 0.5s sample of the time-dependent, low-variance, input (red line) and the corresponding outputs of three filters (blue lines). **Bottom plot, left column:** As for the left column (same three example filters), but for a high-variance stimulus.

Samples of excitatory conductances derived from the 3 filters in Figure 5.39 (bottom plot) are plotted in Figure 5.40.

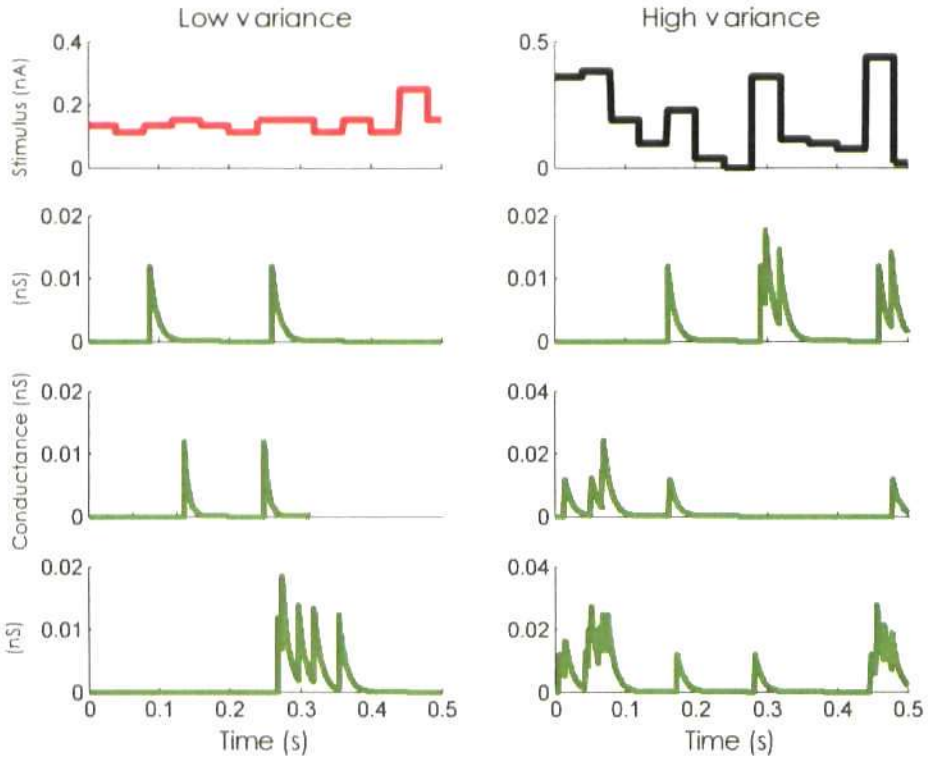

Figure 5.40: Sample excitatory conductances derived from the output of 3 non-linear, rectified stimulus filters
Left column: Uppermost plot gives a 0.5s sample of low-variance stimulus (red line) and the three lower plots show 3 sample conductances (derived from the filters in Figure 5.39) generated in response to the input. Right column: As for left column, but for a high-variance stimulus (black line)

The summed conductance, from all filters in the population, was used to provide background, modulatory input to a single, pulse-based model neuron, which was also driven by the stimulus current. A schematic description of the model is given in Figure 5.41.

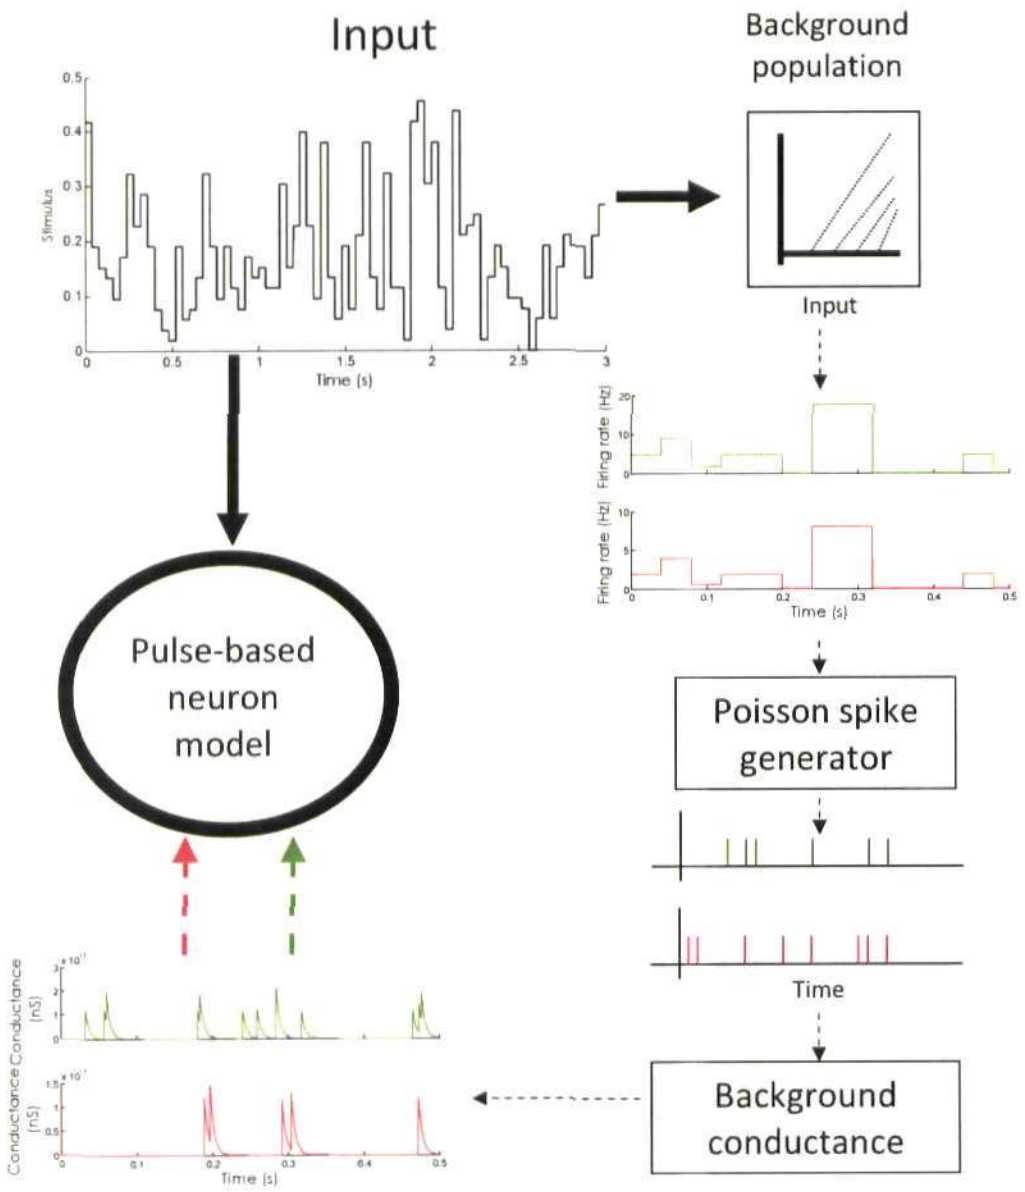

Figure 5.41: Schematic description of the background-conductance model, with modulatory noise derived from non-linear inputs
Driving input, represented by the thick, solid, black arrow, was a time-varying current injection (3s sample of the stimulus under the high-variance condition). The input was also fed into a bank of parallel non-linear filters, whose output was converted to a firing rate and used to drive a Poisson spike-generating mechanism, from which the background (dashed arrows) conductance was derived. Green and red lines represent excitatory and inhibitory conductance, respectively.

The response of the model neuron, averaged over 100 trials, is plotted in Figure

5.42. The slope of the stimulus-response function (Figure 5.42, top plot) displayed a significant (one-factor ANOVA test, with a significance level, $p<0.001$) decrease in gain as stimulus variance increased. The average value of the slope at the half maximal point (S_{50}) was 0.0353Hz/nA^2 (±0.00053 SE), 0.0314Hz/nA^2 (±0.00058 SE) and 0.0268Hz/nA^2 (±0.00037 SE) and the stimulus amplitude that elicited the S_{50} was 0.1812nA (±0.0012 SE), 0.1897nA (±0.0013 SE) and 0.1984nA (±0.0012 SE) for the low-, mid- and high-variance conditions, respectively.

The step size function (Figure 5.42, middle plot) was displaced slightly, as a function of stimulus-variance and the normalised deviation-from-mean function (Figure 5.42, bottom plot) did not scale-up, at the tails of the function (especially for positive deviations-from-the-mean); they were therefore also in line with the experimental data.

The average strength of the excitatory and inhibitory conductance, as a function of the stimulus, is plotted in Figure 5.43 and did not show any significant increase with stimulus-variance. However, the average value of the excitatory conductance, during each 30s simulation, increased steadily from 0.0077nS (±0.0199 SD) to 0.0088nS (±0.0204 SD) and 0.0115nS (±0.0223 SD) and the overall, mean inhibitory conductance increased from 0.0594nS (±0.1208 SD) to 0.0688nS (±0.1302 SD) and 0.0870nS (±0.1405 SD) for the low-, mid- and high-variance conditions. The overall average increase in background conductance was sufficient to generate a decrease in the gain of the model neuron, however this was observed only when combined with spike-rate adaptation (implemented as a dynamic membrane threshold) and the addition of a tonic inhibition that was proportional in strength to stimulus mean.

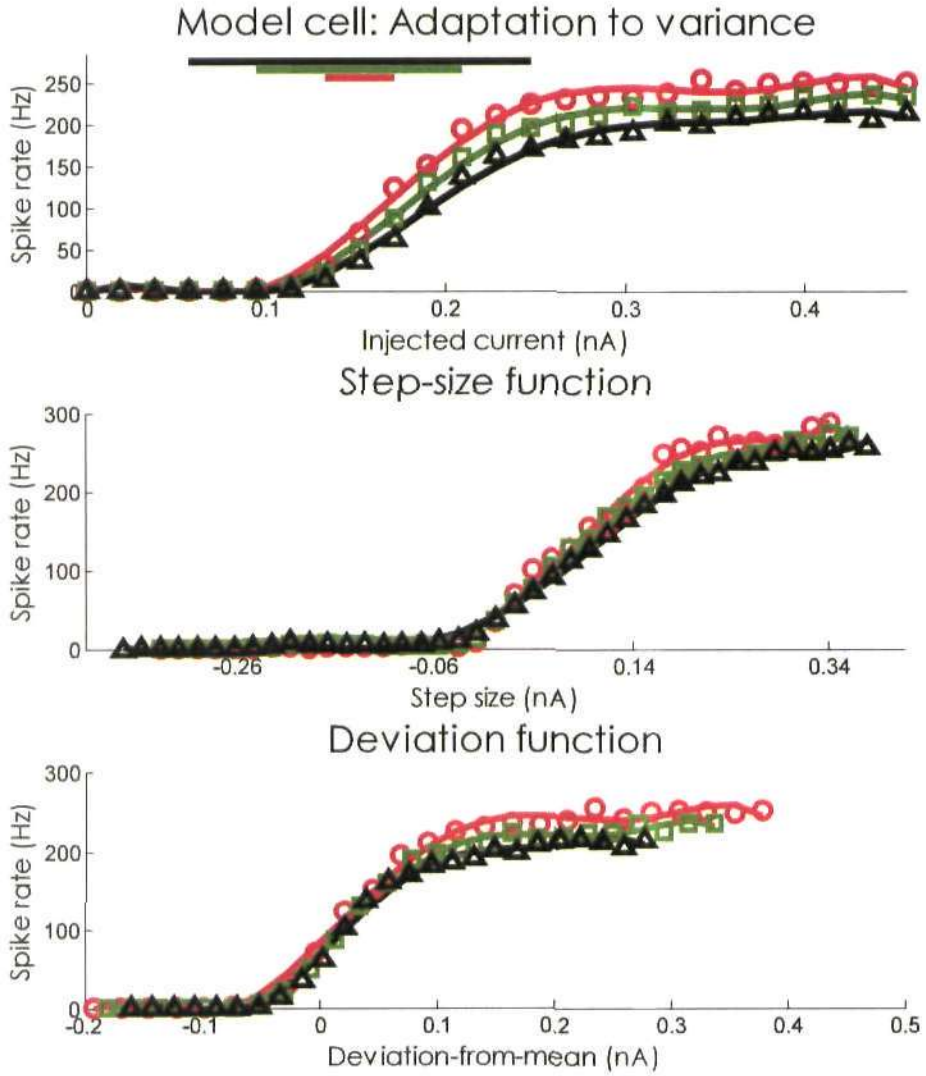

Figure 5.42: Adaptation to variance, in the model cell: background firing rates derived from a bank of non-linear filters of the input
Solid lines represent the best-fit data (5th-order polynomial) and symbols represent the actual data. Red line/circles represent the response to the low variance stimulus, green line/squares gives the mid-variance response and black line/triangle is the high-variance condition.

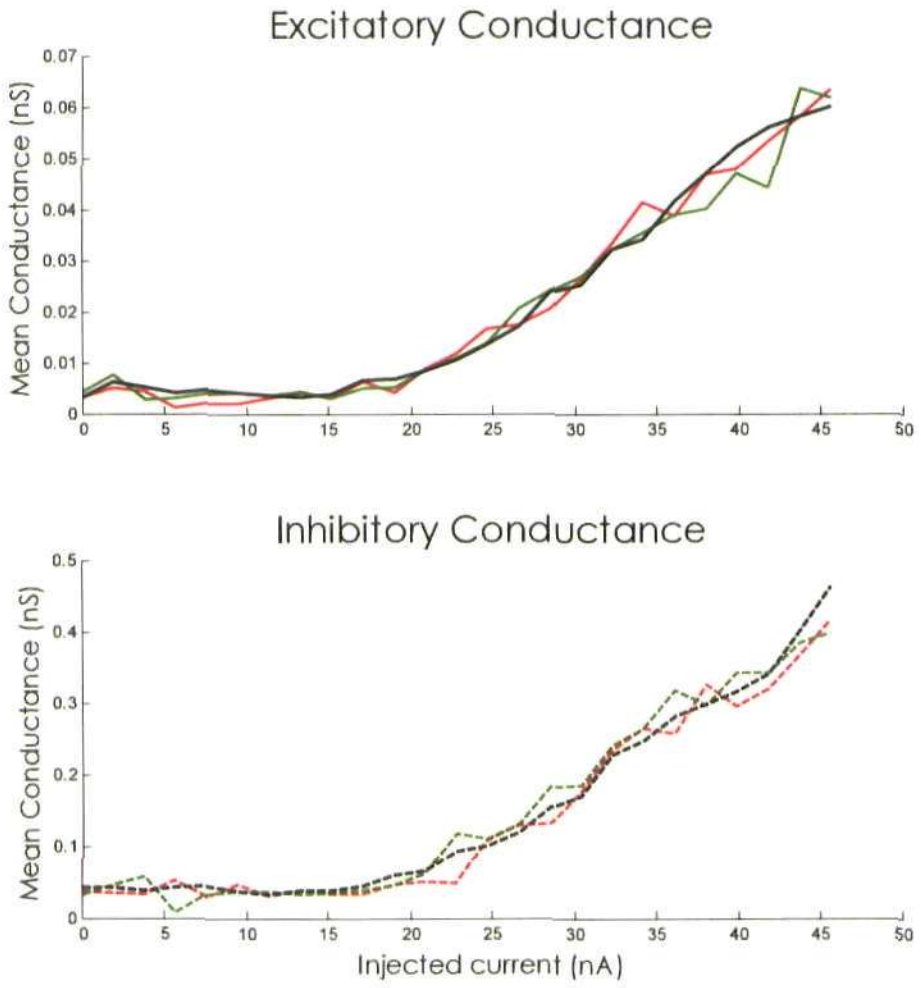

Figure 5.43: Average background synaptic conductance as function of the stimulus Red, green and black lines plot the average value of the background synaptic conductance (derived from the non-linear filters of the stimulus). Solid and dashed lines plot the excitatory and inhibitory conductances, respectively. The initial, un-filtered, input is plotted on the abscissa.

Does the introduction of balanced, noisy background conductances affect adaptation to stimulus mean, in the model cell?

As argued in the introduction to this Chapter (see Section 5.1.2, above), one can assume that under conditions of adaptation to mean the background firing rates remain

fairly constant. Thus in order to check that the lateral, adaptive, shift (observed in response to adaptation to stimulus mean) in the model neural response-profile was not abolished by the introduction of noisy, balanced, background excitatory and inhibitory conductances (50Hz, for all conditions), the model cell was tested again under conditions of adaptation to global mean. The levels of tonic inhibition, current noise and spike-rate adaptation were as for Section 5.4.1 with the only difference being the introduction of stochastic background conductances.

As can be seen in Figure 5.44, the lateral shift in the stimulus-response function was not abolished by the introduction of stochastic background excitation and inhibition. Thresholds (stimulus amplitudes that elicited the half-maximal response) increased on average (between the lowest and highest mean adapting levels) from 0.202nA to 0.302nA (compared to a threshold increase of 0.169nA-0.297nA in the absence of background noise, see Section 5.4.1 and Figure 5.29), thus the introduction of background stochastic conductances into the model increased the displacement of the neural rate-function due to tonic inhibition and spike-rate adaptation; the roughly invariant relationship between output and local differences in stimulus amplitude was also retained and replicated the experimental data (please see Figure 5.25).

5.4.3: Balance of inhibitory and excitatory synaptic depression dependent on levels of variance in adapting stimulus

Within the introduction to this Chapter (Section 5.1.3) it was stated that neural gain can be influenced by both the balance between afferent excitatory and inhibitory levels of steady-state depression and the ratio of the initial excitatory to inhibitory firing rates. Thus the question to be addressed here is whether the adaptive decrease in gain, as a function of

increased stimulus-variance, can be replicating through adjusting the balance between excitatory and inhibitory depression in a population of neurons, in which inhibition has a greater influence than excitation.

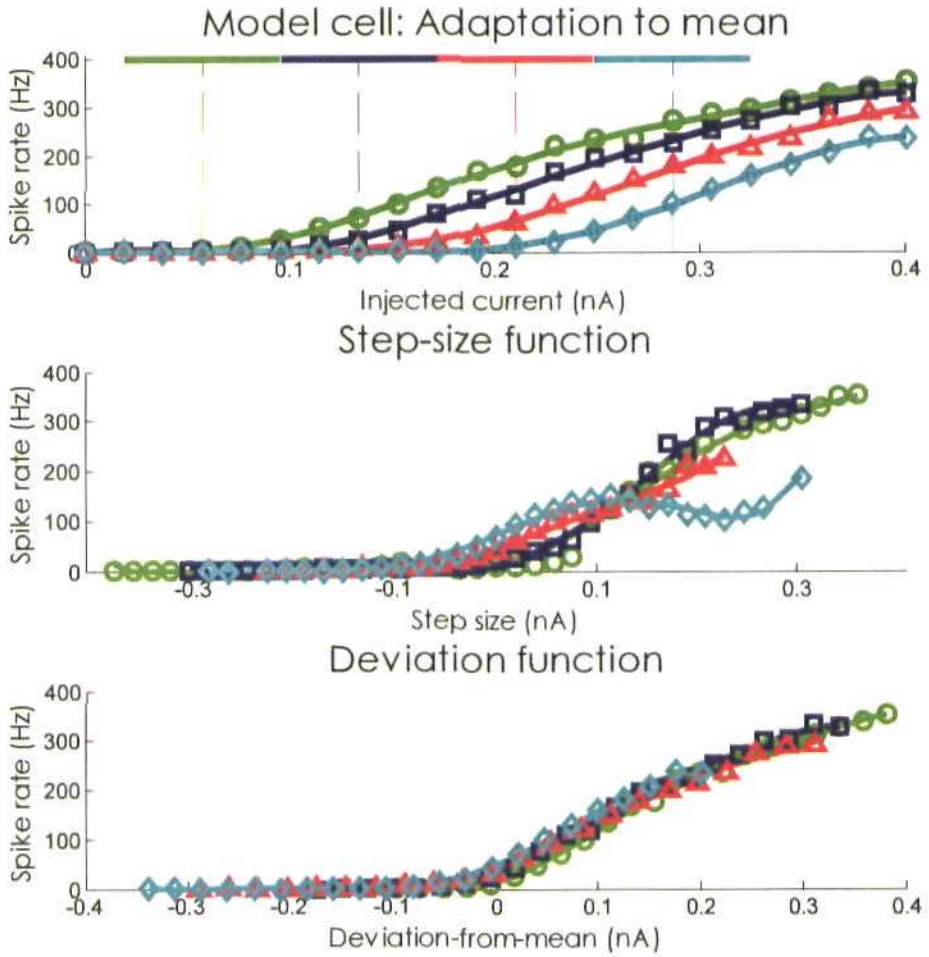

Figure 5.44: Adaptation to mean, in the model cell, with noisy background conductances, as in Chance et al. (2002)
Solid lines represent the best-fit data (5th-order polynomial) and symbols represent the actual data. **Top:** Stimulus-response function. **Middle**: Step-size function. **Bottom**: Deviation-from-mean function. Colours/symbols as for Figure 5.26, above

As for Section 5.3.3, above, the stimulating current was used as input to two sets

of population equations (see equations 5.12-5.14) which gave the excitatory and inhibitory responses to the stimulus; the weighted difference of which was then used to stimulate the model cell. Of course, it should be noted that in this case the stimulus was not a single current step but a time varying current injection with a change in amplitude every 40ms.

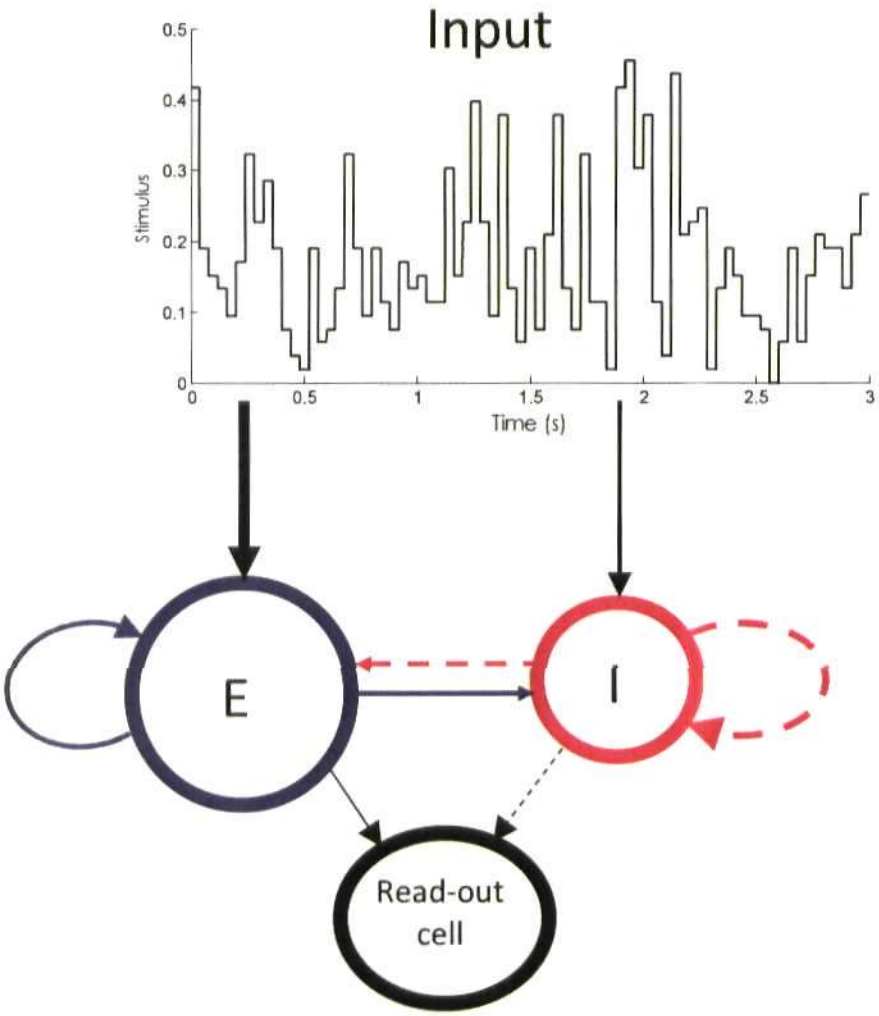

Figure 5.45: Schematic description of the synaptic depression model
The input is time varying current (sample is from high-variance condition); the inhibitory (I, represented by red circle) population receives less driving current that the excitatory (E, represented by blue circle) population (governed by constant k, which was set to 0.4 in the simulations below). Dashed lines indicate inhibition, whilst solid lines represent excitation; red arrows indicate recurrent inhibition, blue arrows: recurrent excitation and black arrows: feedforward connections.

The population equations only return the steady-state response to the current step, irrespective of stimulus history, thus they act as a type of linear filter on the stimulus, with the output dependent on the relative strength of synaptic depression.

The inhibitory depression factor p was fixed at 0.8 for all variance conditions. For excitation, p was set to 0.5, 0.7 and 0.8 for the low-, mid- and high-variance conditions, respectively; with the assumption that as variance in the stimulus increases, excitatory depression becomes weaker and approaches the level of inhibitory depression. The output of the model cell is plotted in Figure 5.46. The effect of reducing excitatory depression, as variance increased, was to reduce the gradient of the model stimulus-response function and overall firing rate. Slope at the S_{50} decreased, on average, from 0.0139Hz/nA^2 (±0.0016SE), under the low-variance condition to 0.0130Hz/nA^2 (±0.0010SE) and 0.0112Hz/nA^2 (±0.0008SE) for the mid- to high-variance conditions, respectively. There was also a slight increase in threshold from 0.2874nA (low-variance; ±0.0055SE) to 0.2902nA (±0.0064SE) and 0.3138nA (±0.0060SE), for the mid- to high-variance trials.

To see how changing the balance of depression affects the afferent input to the model cell, the individual population afferents were used to drive the cell, as opposed to the weighted difference between the two population used in the simulation above. The responses were plotted for the excitatory and inhibitory inputs in Figure 5.47. As depression levels approached symmetry, afferent inhibition increased dramatically more than the relevant increase in afferent excitation, thus responses were decreased, relative to the base response (low-variance/ excitatory $p = 0.5$, red lines in Figure 5.47). Overall, the steepness of the slope of the inhibitory rate-function also increased more than the relevant excitatory response; thereby accounting for the decrease in gain seen in the simulation

above (see Figure 5.46).

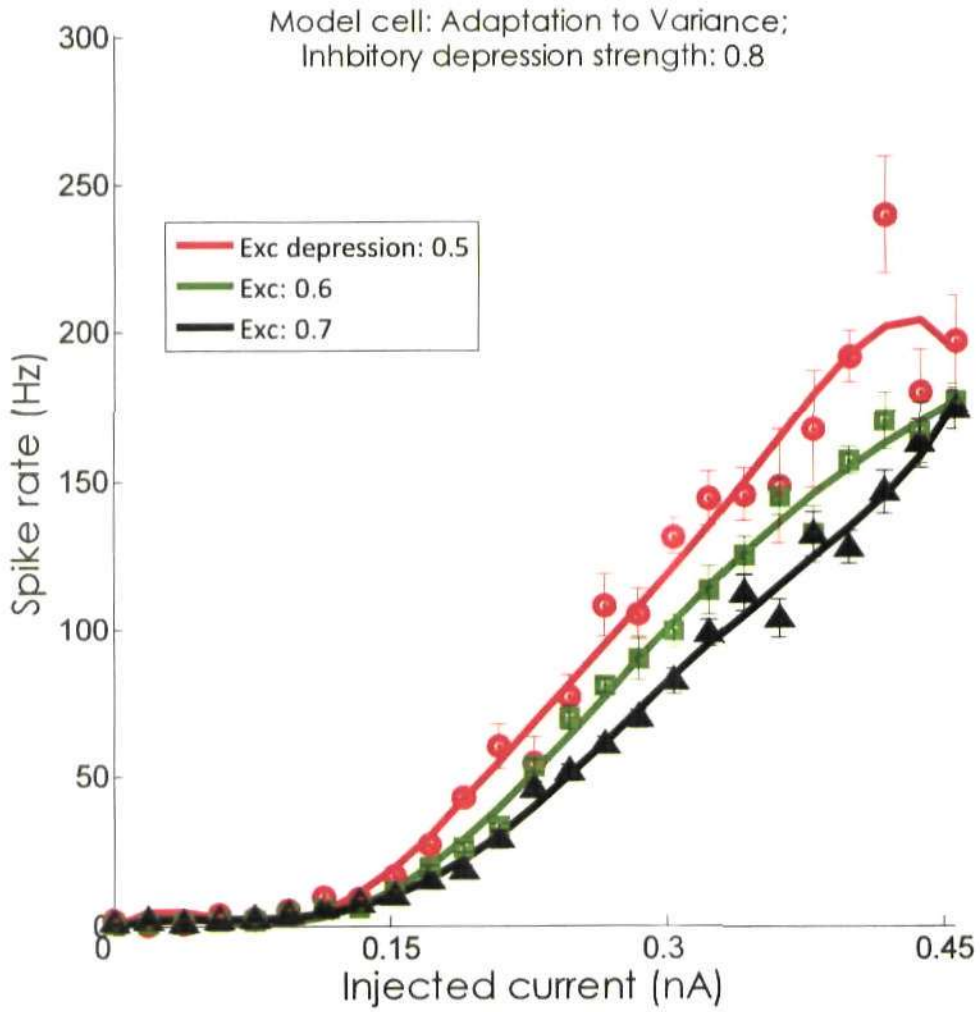

Figure 5.46: Gain modulation through asymmetric depression
Steady-state model of asymmetric depression, see text for details

The output of the model cell, following adaptation to mean and with a fixed ratio of depression is plotted in Figure 5.48. Excitatory depression levels were assumed to be fairly strong to reflect the saturation of depression due to the sustained nature of the stimulus;

excitatory and inhibitory depression levels were fixed at $p = 0.4/0.8$, respectively.

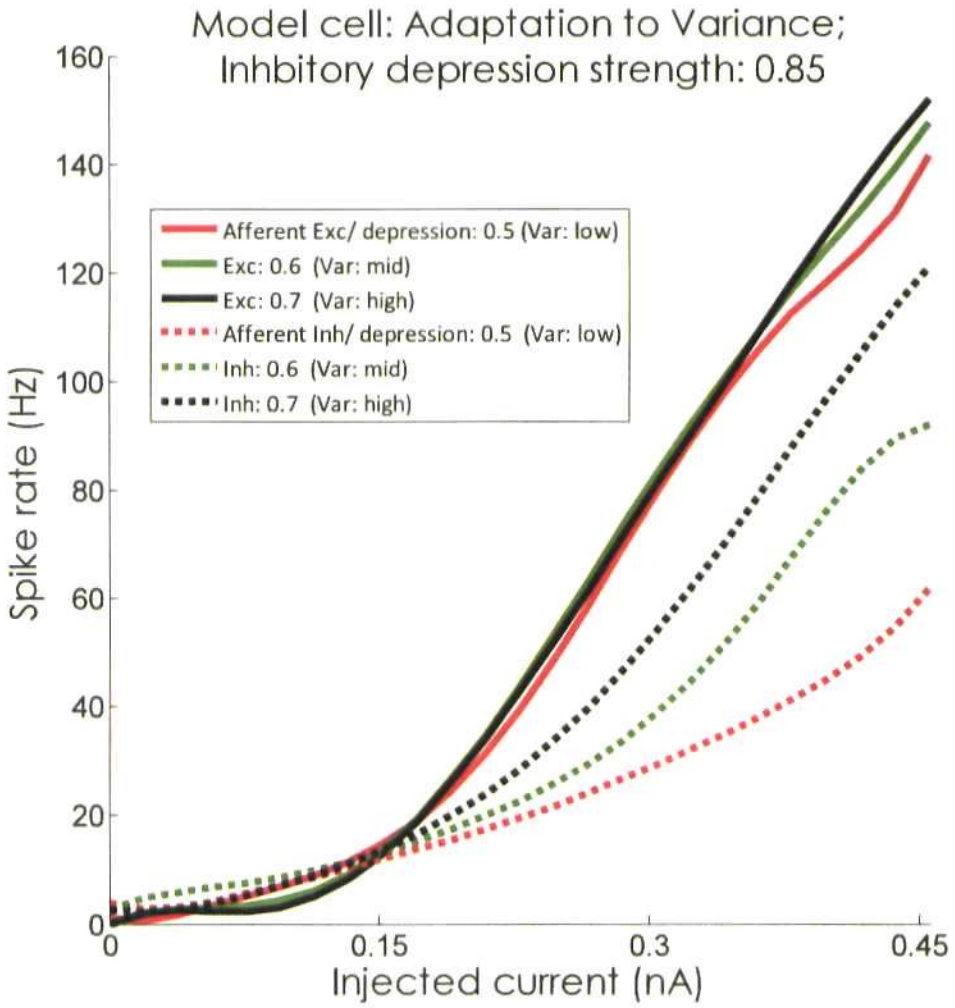

Figure 5.47: Relationship between excitation and inhibition, under conditions of weakening depression
Solid lines plot the excitatory input, dashed, the inhibitory input to produce the model output in Figure 5.46. Colours represent the different variance/excitatory depression levels and are given in the inset.

The model cell undergoes a lateral shift in the stimulus-response function, as expectated, and with the addition of a tonic inhibition the step-size function was relatively

invariant (bar some variance for the highest high-probability region condition). Therefore the model of gain modulation by asymmetric synaptic depression does not abolish the change in threshold associated with adaptation to stimulus mean.

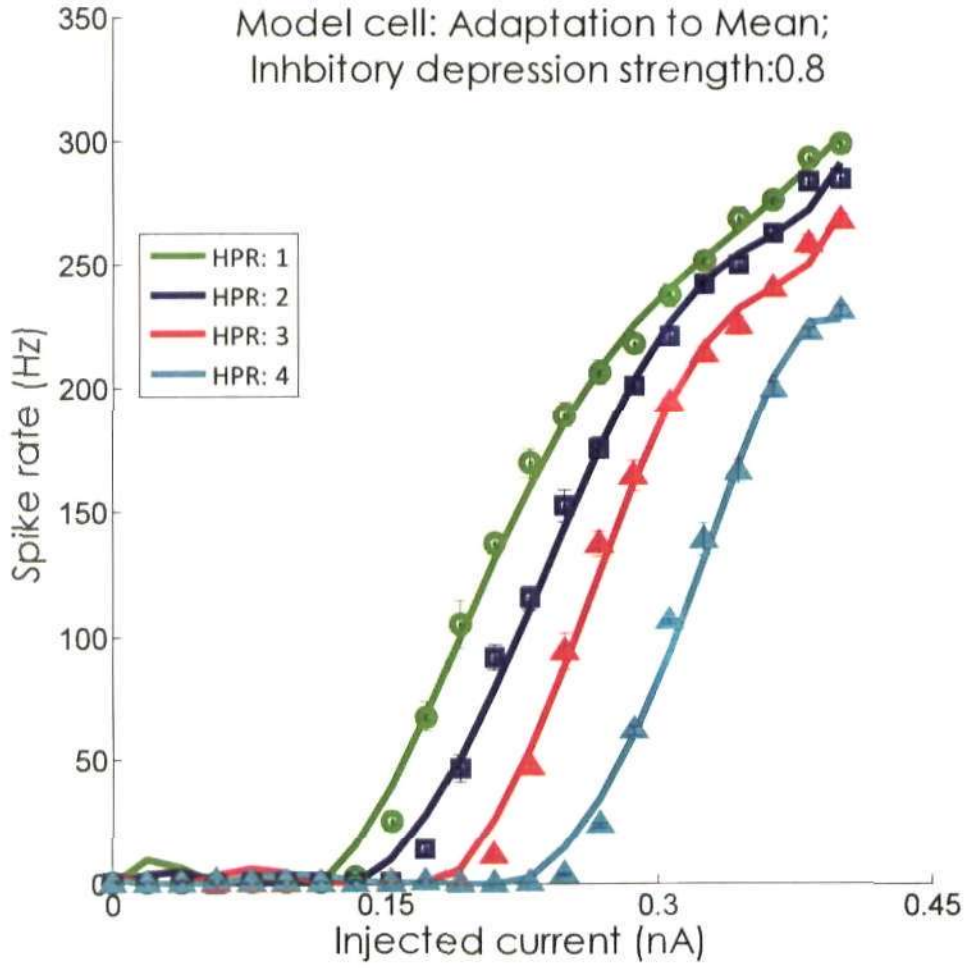

Figure 5.48: Adaptation to mean, Asymmetric depression
Excitatory depression was fixed at $p = 0.4$ for all conditions

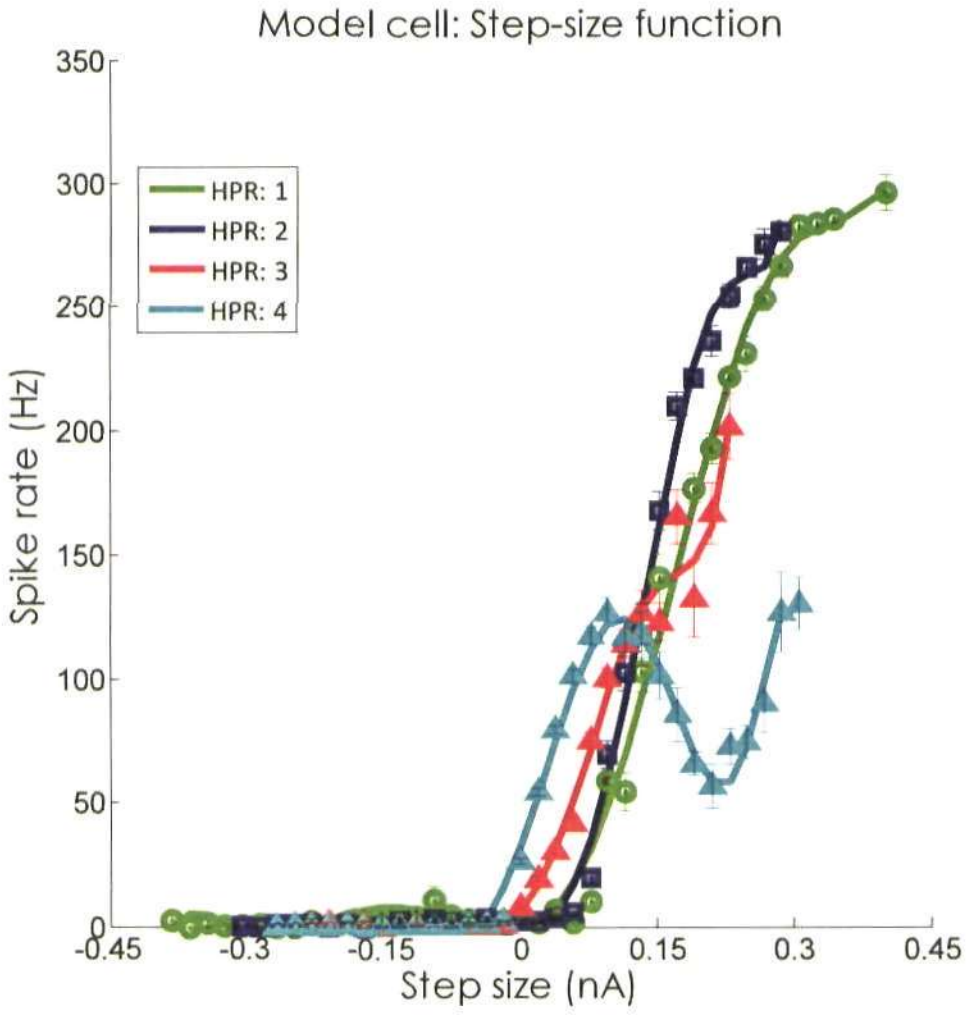

Figure 5.49: The step-size function following adaptation to mean, with asymmetric depression
Step-size function for the rate function plotted in Figure 5.48

As previously stated, these results are derived from a steady-state response to the stimulus. Of course, synaptic depression is highly dependent on stimulus history and, as stimuli amplitudes change too rapidly to allow for full recovery of synaptic dynamics during each, successive, 40ms presentation, it's important to test how the model cell responds to a time-varying input, under the different depression conditions. The adapting

stimuli were used to drive a recurrently connected circuit of four excitatory/one inhibitory I & F units, which provided afferent input to a read-out cell. For a description of the parameters see Section 5.2.6.

The output neuron was a pulse-based neural model, with a dynamic threshold (the spike-rate adaptation constant, b, was set to unity). Synaptic depression of the afferents targeting the output cell were the same for both inhibition and excitation (depression factor, $p = 0.8$, across all conditions), whilst recurrent inhibitory depression levels were fixed at $p = 0.85$ and excitatory depression was inversely proportional to stimulus variance (i.e. p increased from a base level of $p = 0.3$, under the low-variance condition, to $p = 0.45$ & 0.6 for the mid- and high-variance conditions, respectively); the results are plotted in Figure 5.50.

The value of the model cell's slope and S_{50} stimulus amplitude are plotted in Table 5.6. Overall the stimulus-response function was scaled downward as a function of stimulus variance, however, there was no successive decrease in the gradient of the curve, at the half-maximal point; thus the model generated a downward scaling of the response, but not a pure change in gain modulation.

What contribution do the excitatory and inhibitory cells make to the output of the model cell? The average firing rate response, as a function of stimulus intensity, of the four excitatory input neurons and the single inhibitory unit are plotted in Figure 5.51. For stimulus amplitudes of less than (roughly) 0.13nA, the firing rates of the excitatory units (Figure 5.51, solid lines) were similar under all adapting conditions; however, for stronger stimulus intensities, the responses diverged: excitatory firing rates decreased in line with the concurrent increase in stimulus variance and decrease in excitatory depression strength.

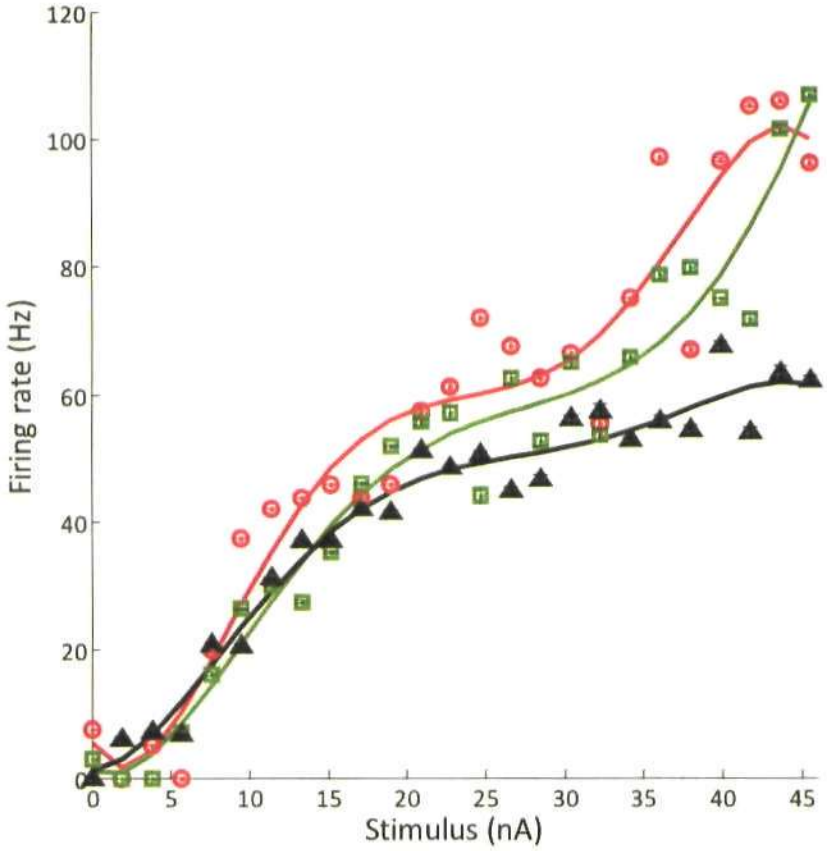

Figure 5.50: Adaptation to variance, in the model cell: Asymmetric synaptic depression
Inhibitory depression levels were fixed at 0.85 whilst excitatory depression was set to $p = 0.3$ for the low-variance stimulus (Red line/circles), $p = 0.45$ for the mid-variance condition (green line/squares) and finally $p = 0.6$ under high-variance stimulation (black line/triangles). Solid lines represent the best-fit to the data (5th-order polynomial) and symbols represent the actual data.

Conversely, inhibitory firing rates (Figure 5.51, dashed lines) increased with stimulus variance, thus the mechanism of gain modulation was subtly different from the steady-state simulations, described above, especially with reference to afferent excitation.

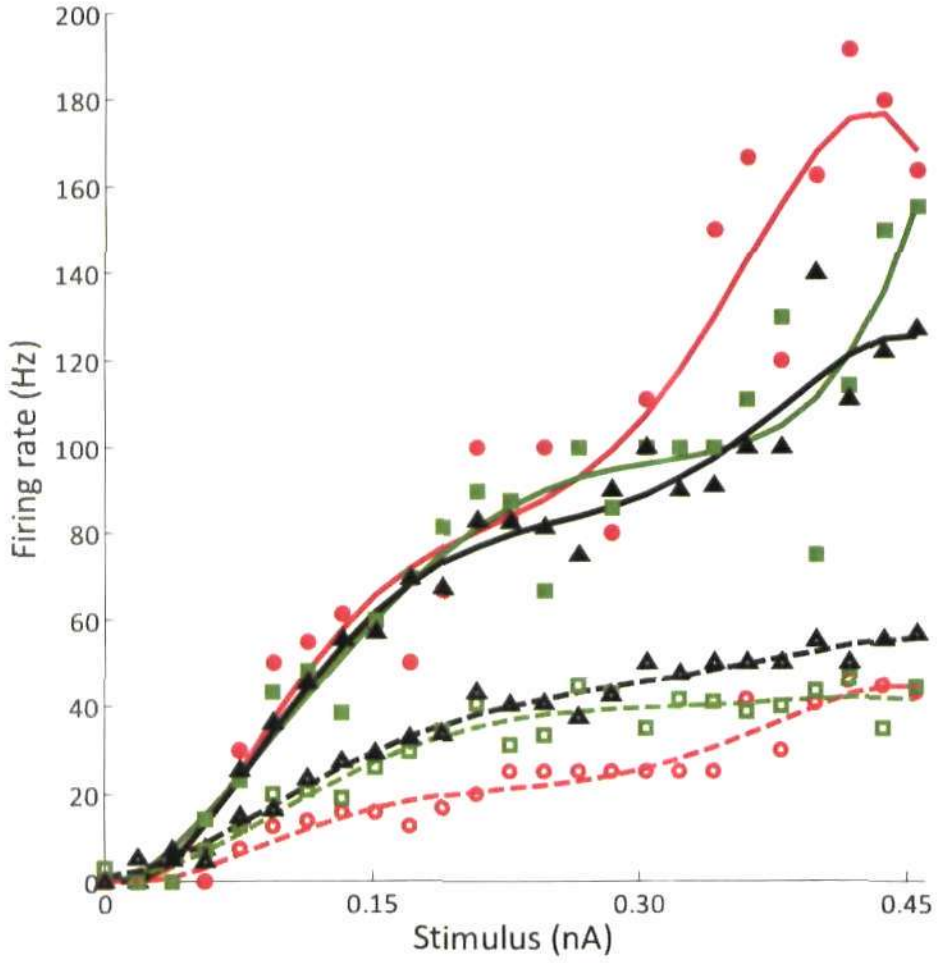

Figure 5.51: Average excitatory and inhibitory firing rates in the model circuit, for different ratios of excitatory and inhibitory depression levels
Inhibitory depression levels were fixed at 0.85 whilst excitatory depression was set to $p = 0.3$ for the low-variance stimulus (Red line/circles),$p = 0.45$ for the mid-variance condition (green line/squares) and finally $p = 0.6$ under high-variance stimulation (black line/triangles). Inhibitory inputs were represented by dashed lines and empty symbols, whilst the average firing rate, across all four excitatory units, was plotted with sold lines and symbols. Lines represented the best-fit to the data (5th-order polynomial) and symbols plotted the actual data.

Table 5.6: Amplitude and slope of model stimulus-response function, under different ratios of excitatory and inhibitory depression (inhibition: $p = 0.8$)

Variance Level	**Low**	**Mid**	**High**
Excitatory depression (p)	*0.3*	*0.45*	*0.6*
S_{50} *(nA)*	0.1662	0.2228	0.1240
(±SE)	(0.0016)	(0.0023)	(0.0009)
Slope (Hz/nA)	0.0429	0.0243	0.0502
(±SE)	(0.0014)	(0.0015)	(0.0004)

Extension to the model of asymmetric synaptic depression

Within the asymmetric depression model, the balance between levels of excitatory and inhibitory depression was manipulated in order to generate gain modulation. The depression-factor, p, for both excitation and inhibition was fixed, depending on the level of stimulus variance. However, in general, one would not expect the depression factor to vary as function of stimulus input; thus it could be argued that the model above does not represent a truly biologically valid mechanism of cortical adaptation.

The primary effect of reducing levels of excitatory depression, as stimulus variance increased, was to strengthen the excitatory drive acting on inhibition (please see Figure 5.47 and Figure 5.51) within the model; inhibition was dominant within the network and the gain of the read-out cell was thus reduced. One possible mechanism by which a shift in the overall balance between excitation and inhibition would emerge, without the need for explicit hand-tuning of parameters, could be through relative changes in overall levels of release from excitatory depression, as a function of stimulus-variance. The hypothesis

being that, as the high-probability region of the stimulus distribution widens (thereby increasing global variance), there is a relative increase in the presentation of stimulus amplitudes that only elicit sub-threshold responses; this should, theoretically at least, allow for relatively more recovery from depression as stimulus-variance increases.

Gain change through release from depression

In order to test whether increased diversity in the stimulus was sufficient to generate the necessary release from excitatory depression, simulations were run, using an I & F model (similar to that described in Section: 5.4.3), in which depression levels were fixed for all three stimulus-variance conditions. In brief, the model contained 1 inhibitory and 4 excitatory, self-connected, I & F cells, that were also reciprocally connected to each other. The population of I & F cells provided the sole, feed-forward, depressing input to an excitatory I & F model neuron that acted as the read-out cell. The read-out cell was implemented with spike-rate adaptation (i.e. a dynamic membrane threshold) and a tonic inhibitory conductance that was proportional, in strength, to the global-mean amplitude of the adapting stimulus. All parameters were as for the I & F model in Section 5.2.6 and 5.4.3, above.

For the 1st simulation the ratio of excitatory to inhibitory recurrent depression was set to $p = 0.3$ for excitation and $p = 0.8$ for inhibitory depression (AS_{EXC} configuration). In the 2nd simulation depression levels were balanced, i.e. $p = 0.8$ for both excitatory and inhibitory depression and for the 3rd simulation depression was asymmetric with inhibitory depression set to be stronger than excitation (AS_{INH}, $p = 0.8$ for excitatory synapses and $p = 0.3$ for inhibition). Feedforward depression was set to $p = 0.8$, for both feedforward inhibition and excitation. A simulation was also performed in which there was no recurrent

depression acting within the network; the results for all 4 simulations (across 100 trials) are plotted in Figure 5.52.

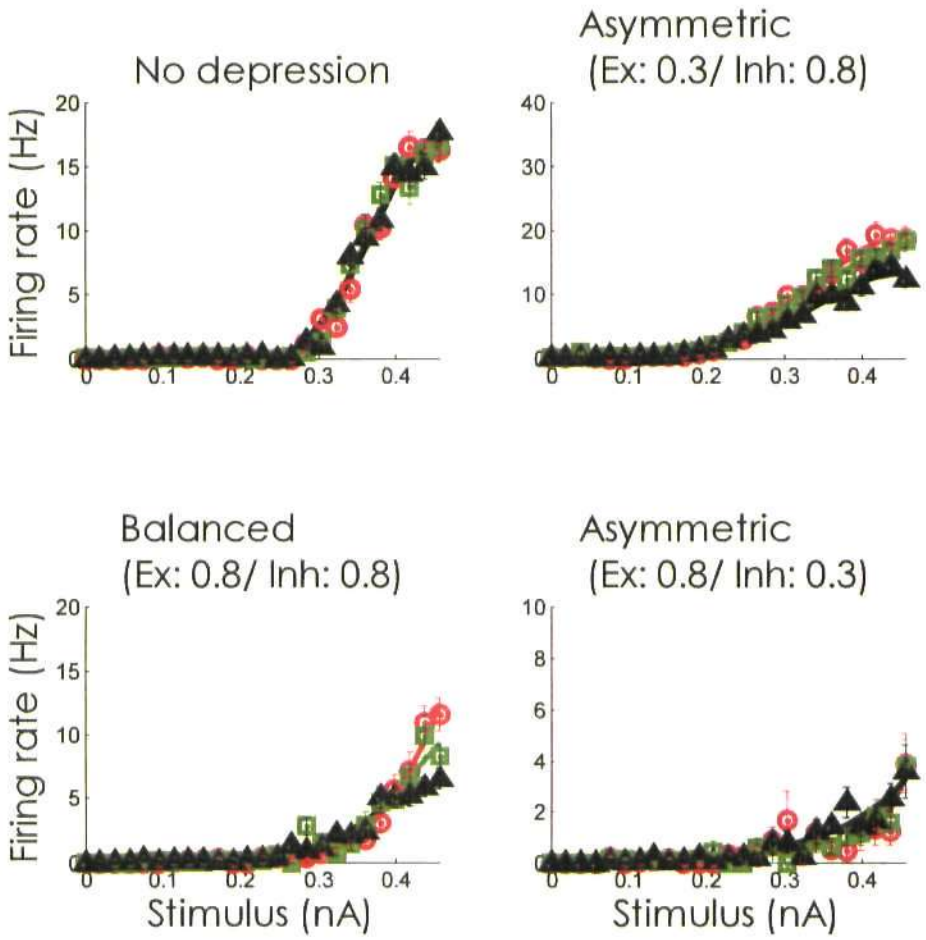

Figure 5.52: Stimulus-response functions, under all stimulus configurations
Red, green and black lines/symbols plot the low-, mid- and high-variance conditions respectively. **Top left:** No recurrent depression. **Top Right:** Asymmetric depression: 0.3:0.8 (Excitatory : Inhibitory). **Bottom Left**: Balanced: 0.8:0.8. **Bottom Right:** Asymmetric depression: 0.8:0.3.

Increasing adapting variance only affected the output of the model cell, in the presence of recurrent depression. There was a slight decrease in neural gain for the AS_{EXC}

(top right) configuration but it was not significant ($p > 0.1$, non-parametric Kruskal-Wallis test). Overall, no significant decrease in gain or shift in stimulus amplitude at the half-maximal point (S_{50}) was observed during any of the simulations.

As described in the introduction, the combination of tonic inhibition and synaptic depression has been shown, experimentally, to elicit multiplicative changes in gain (Rothman et al. 2009). Within the model, tonic inhibition is proportional to the global-mean of the input. However, as global stimulus-mean increases slightly with stimulus variance (due to the skewed nature of the stimulus distribution) there is a concomitant increase (albeit insignificant) in tonic inhibition acting on the read-out cell. Thus the combination of feedforward depression and tonic inhibition within the model could result in the slight multiplicative downward scaling of the stimulus-response function.

In order to test whether the changes in the model output (as a function of stimulus variance) were dependent on the membrane dynamics of the read-out cell, the simulations were re-implemented without spike-rate adaptation and tonic inhibition (see Figure 5.53).

In the absence of spike-rate adaptation and tonic inhibition, none of the simulations generated a decrease in the gain of the model output cell, as a function of stimulus-variance (Figure 5.53). Thus it would appear that the increased diversity in the stimulus was not *sufficient to shift the overall balance of synaptic activity to favour inhibition.* Within the model, the temporal dynamics of synaptic recovery from depression were on the order of 100s of milliseconds and therefore did not capture the rapid changes in stimulus amplitude, on the 40ms scale. Manipulating either (or a combination of) the presentation time for each stimulus (e.g. increasing the time-window of presentation from 40ms to 100ms) or the time constant of recovery from depression (e.g. decreasing τ_{REC} from 300ms to 100ms) did not

elicit a decrease in gain (data not shown).

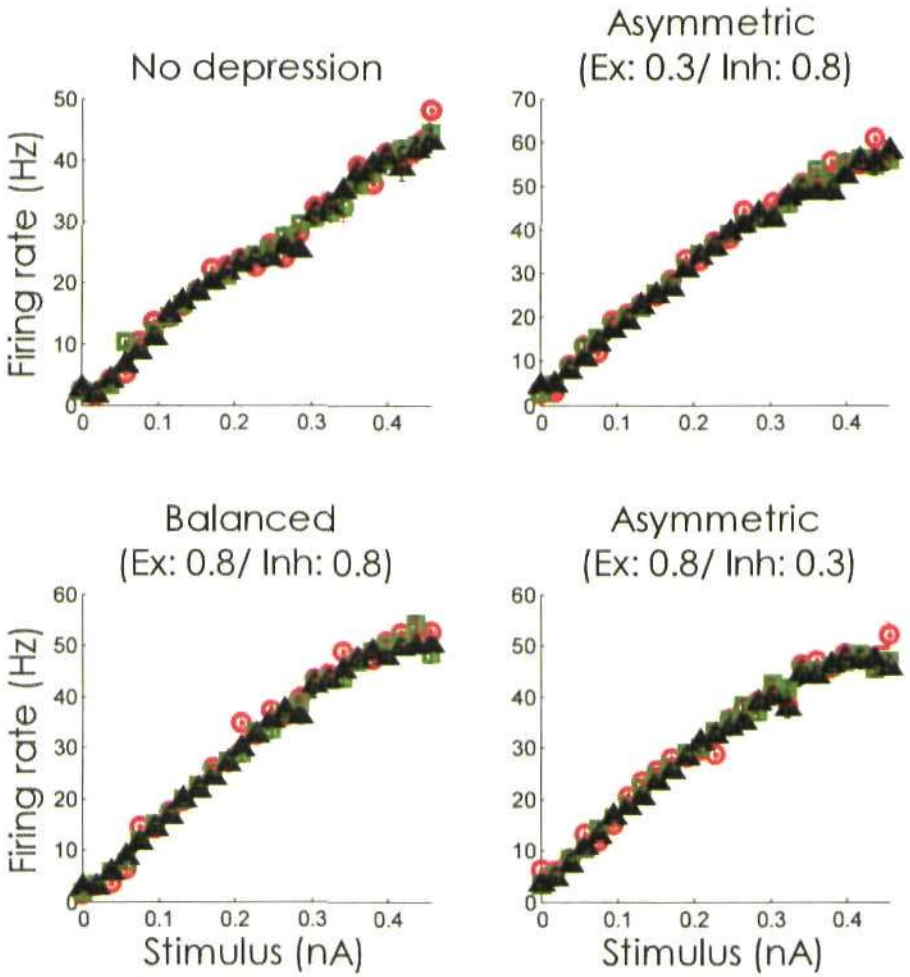

Figure 5.53: Stimulus-response functions, under all stimulus configurations, in the absence of spike-rate adaptation
Red, green and black lines/symbols plot the low-, mid- and high-variance conditions respectively. **Top left:** No recurrent depression. **Top Right:** Asymmetric depression: 0.3:0.8 (Excitatory : Inhibitory). **Bottom Left**: Balanced: 0.8:0.8. **Bottom Right:** Asymmetric depression: 0.8:0.3.

Synaptic Facilitation

Within the hand-tuned asymmetric depression model outlined above, the main effect of reducing excitatory depression was to increase the excitatory drive acting on

recurrent inhibition. In order to implement this, it was decided to incorporate synaptic facilitation into the I & F model. There is evidence that many cortical synapses display both facilitation and depression (e.g. Thomson 1997; Varela et al. 1997; Markram, Wang, Tsodyks 1998). Under conditions of facilitation, each time a pre-synaptic spike is fired, subsequent responses (within a given time-window) are enhanced, relative to the initial response. Within V1, facilitation has been observed to have faster temporal dynamics than depression (Varela et al. 1997) and thus only influences the initial synaptic response to supra-threshold stimuli.

Under conditions of increasing variance, the frequency of high-intensity stimuli also increases, thus it could be argued that facilitation would act to boost the otherwise suppressed synaptic response. Incorporating facilitation into the recurrent excitatory-to-inhibitory synapses, within the model network, could generate the relative increase in inhibitory drive necessary to elicit a decrease in gain. Facilitation has been observed in excitatory synapses impinging onto GABAergic inter-neurons (e.g. Thompson 1997).

Facilitation was implemented as for Varela et al. (1997). For each spike the probability of vesicle release was increased (in an additive fashion) by a given facilitatory factor, which was set to 0.6 (Varela et al. 1997); the use-dependent strength of the synapse decayed exponentially back to the base level of unity, with a time constant of 0.1s. Facilitation was restricted to recurrent excitatory-to-inhibitory synapses only, thus enhancing excitatory drive acting on inhibition within the network. In the presence of spike-rate adaptation and tonic inhibition, facilitation resulted in a significant decrease in response gain, but only under the AS_{EXC} configuration (results plotted in Figure 5.54).

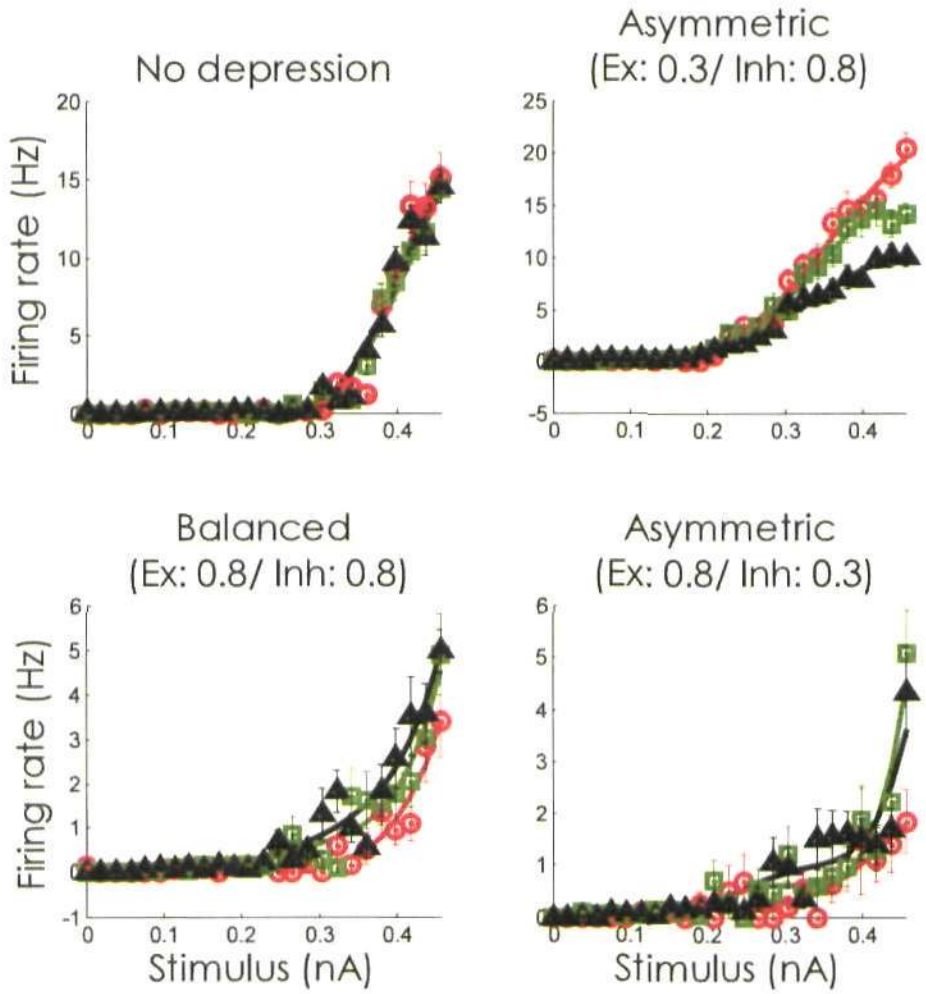

Figure 5.54: Stimulus-response functions in presence of inhibitory-to-excitatory facilitation, with spike-rate adaptation
Red, green and black lines/symbols plot the low-, mid- and high-variance conditions respectively. **Top left:** No recurrent depression. **Top Right:** Asymmetric depression: 0.3:0.8 (Excitatory : Inhibitory). **Bottom Left**: Balanced: 0.8:0.8. **Bottom Right:** Asymmetric depression: 0.8:0.3.

For the AS_{EXC} configuration (Figure 5.54, top right), the average value of the slope at the S_{50} decreased steadily (low to high-variance condition) from 0.0024 (±0.00022 SE) to 0.0013 (±0.00011 SE) with a significance level of $p = 0.0012$ (using a non-parametric Kruskal-Wallis test).

In order to examine the effect of facilitation on the network the reciprocal excitatory and inhibitory recurrent connections are plotted in for the both the AS_{EXC} and balanced configurations

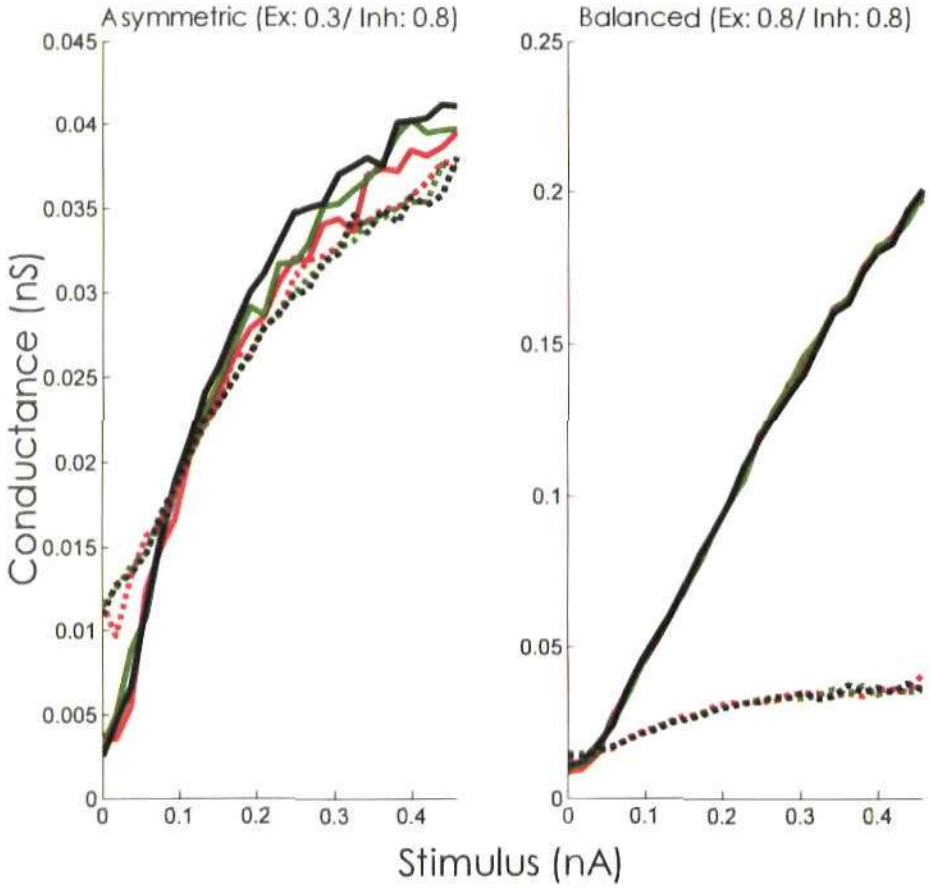

Figure 5.55: Recurrent, reciprocal conductances in the presence of both facilitation and depression. Solid lines plot the excitatory-to-inhibitory conductance and dashed lines plot the inhibitory-to-excitatory conductance. Red, green and black lines/symbols plot the low-, mid- and high-variance conditions respectively

Under the AS_{EXC} configuration, recurrent excitation (solid lines, left plot, Figure 5.61) is generally suppressed (relative to balanced depression), due to the high levels of depression. However, the influence of facilitation has a much stronger effect than that

observed in the balanced configuration (solid lines, right plot, Figure 5.61). Excitatory drive acting on inhibition was increased for the higher variance stimulus conditions, as predicted, thus generating a relative increase in the amount of inhibition in the network and enhancing the small decrease in gain observed in the absence of facilitation (see Figure 5.52).

The average response, in the absence of spike-rate adaptation and tonic inhibition, and across all depression configurations, is plotted in Figure 5.56 below.

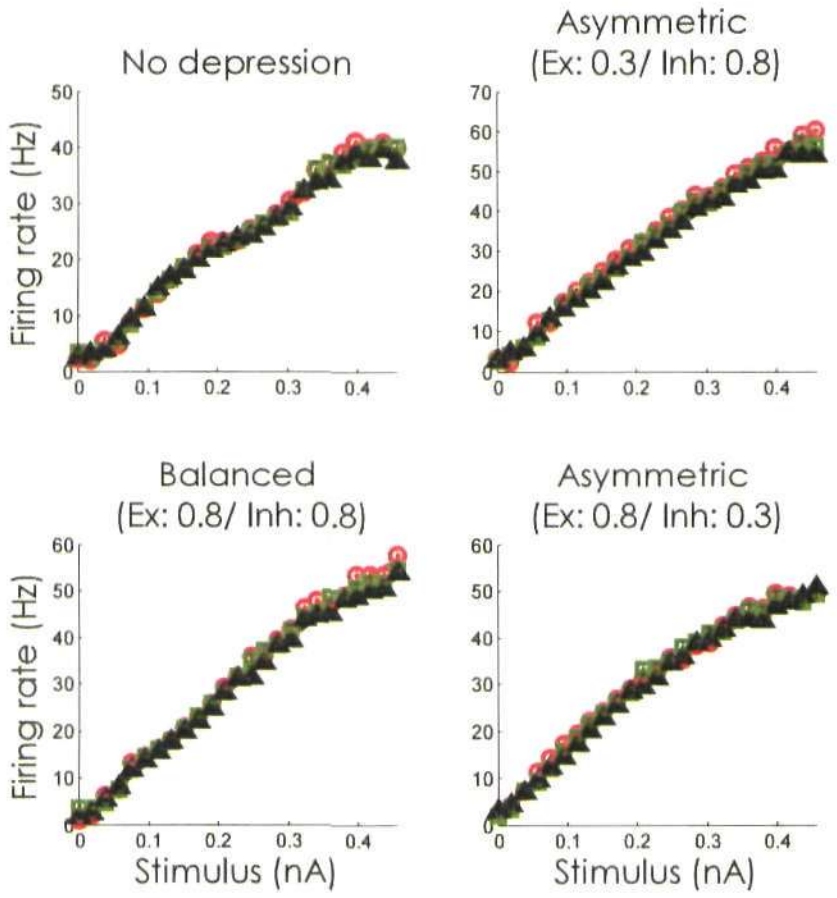

Figure 5.56: Stimulus-response functions in presence of inhibitory-to-excitatory facilitation
Red, green and black lines/symbols plot the low-, mid- and high-variance conditions respectively. **Top left:** No recurrent depression. **Top Right:** Asymmetric depression: 0.3:0.8 (Excitatory : Inhibitory). **Bottom Left**: Balanced: 0.8:0.8. **Bottom Right:** Asymmetric depression: 0.8:0.3.

Facilitation alone did not elicit any significant change in output. There was however a slight decrease in firing rates for the AS_{EXC} configuration, as variance increased. In combination with tonic inhibition, this was sufficient to generate a significant multiplicative decrease in gain, as presented in Figure 5.54 (top right).

5.4.4: Non-linear relationship between stimulus and response

Whilst the addition of facilitation into the model did elicit small, significant, changes in gain, the model could not be considered robust. The decrease in gain was not only reliant on the ratio of depression between the excitatory and inhibitory units but also on the presence of spike-rate adaptation and tonic inhibition. Adaptation to variance appears to be relatively common within the barrel cortex (Garcia-Lazaro et al. 2007), therefore a more robust model would better account for the experimental data.

It was considered earlier, within this chapter, that background noise might increase if derived predominantly from cells that are adapted to stimulus mean but not to stimulus variance. As demonstrated, the non-linear transformation of the input signal elicited a decrease in gain when used to generate noisy, balanced, background conductances (see the extended model in Section 5.4.2, above).

It was decided to use this non-linear input to drive the I & F model directly. The stimulus current was filtered through a bank of parallel non-linear rectified filters, prior to injecting it into the model (for detailed description of filters and parameters see Section 5.4.2, above). It was hoped that the increase in overall levels of firing rates, for the higher variance adapting conditions, would be strong enough to shift the bias of the network towards inhibition.

The stimulus-filters were as for the background noise model above, except that their outputs were used as a current injection into the model (i.e. without conversion into a synaptic conductance). A graphical description of the relationship between the average, summed driving current and stimulus distribution is plotted in Figure 5.57 and a schematic outline of the model and inputs is plotted in Figure 5.58

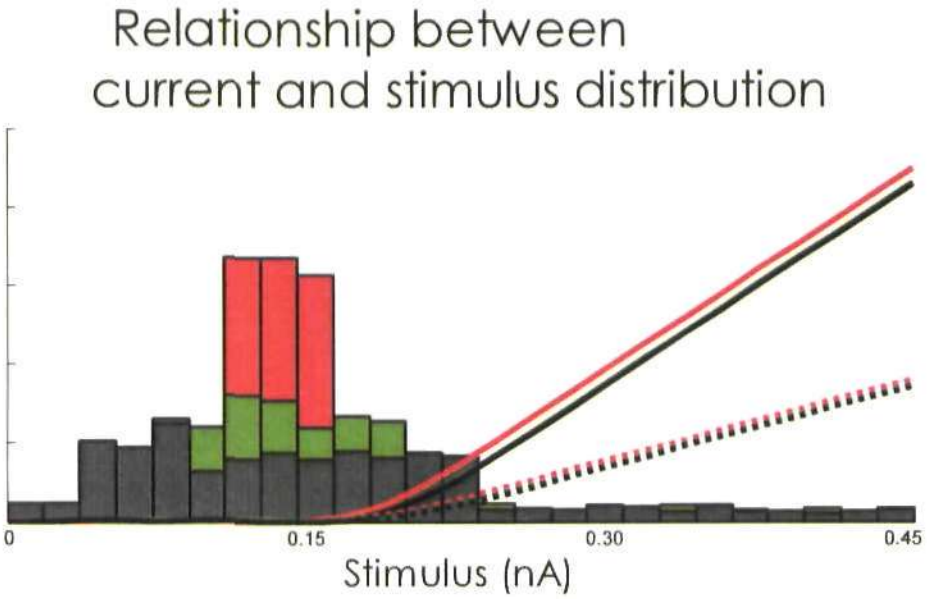

Figure 5.57: Relationship between stimulus distribution and summed filtered input current , plotted as a function of stimulus.
Histogram of number of times an input amplitude is presented during 30s stimulus sequence. Solid (excitatory input) and dashed (inhibitory) lines plot the relationship between the rectified non-linear filter of the input, as a function of stimulus variance. Red, green and black (dark-grey on histogram) plot the low-, mid- and high variance conditions respectively.

The results of the simulation, across all depression configurations, are plotted in Figure 5.59. There was a significant decrease in slope (see Table 5.7), under all depression configurations (significance level $p = 0.088$ for the no depression model and $p<0.0001$ for the remaining configurations).

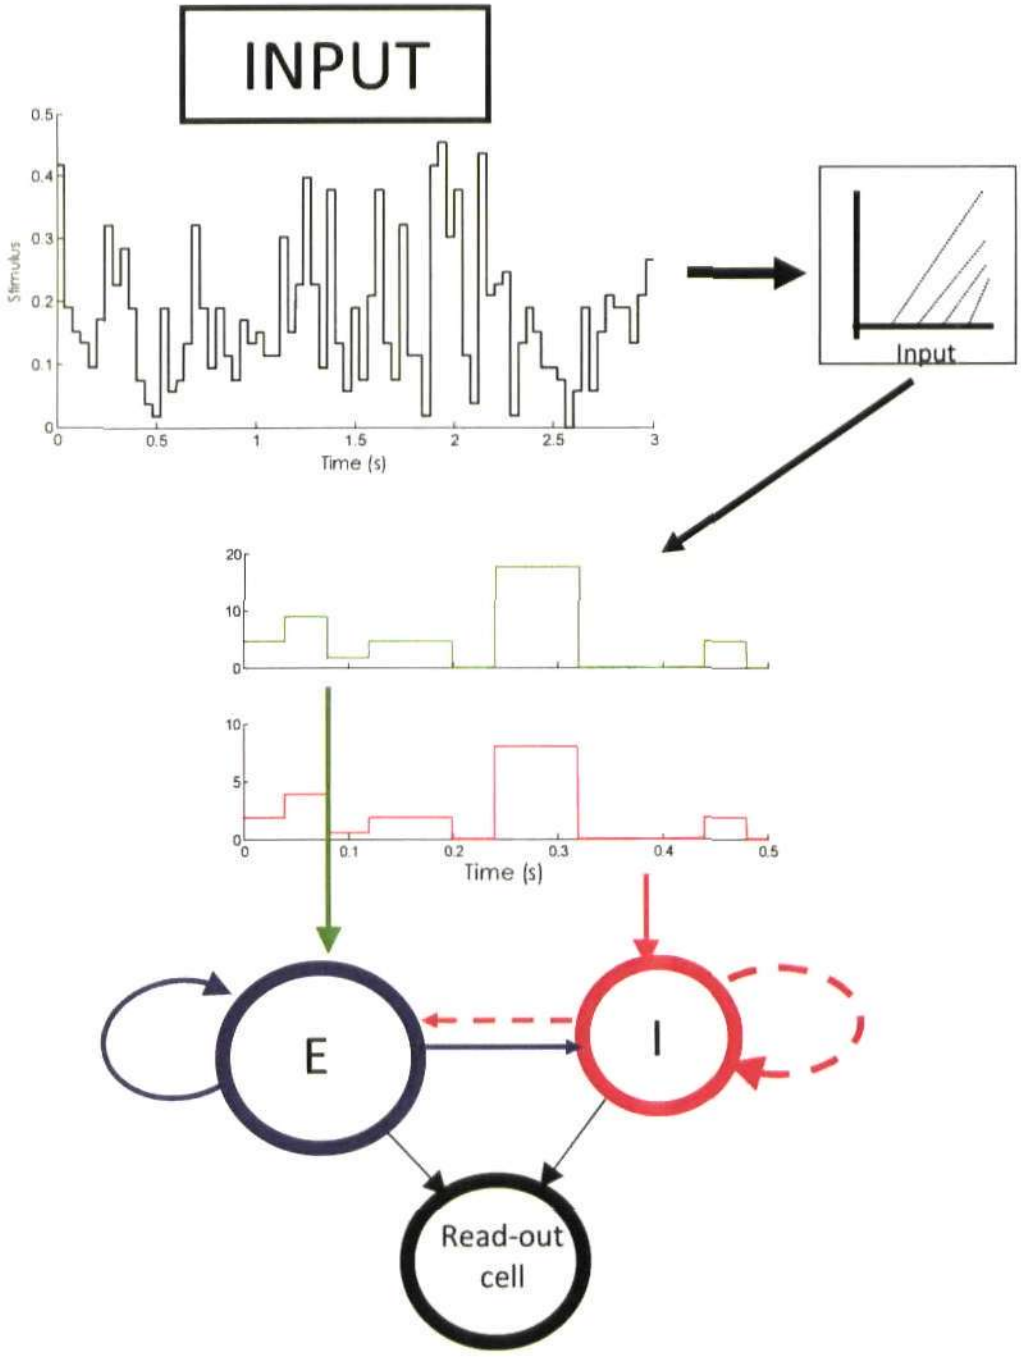

Figure 5.58: Schematic description of model
The input is fed to parallel non linear, rectifiers to generate a filtered stimulus, I, which drives recurrently connected I & F neurons (E and I). These provide feedforward input (black arrows) to the read-out cell; inhibitory stimulus is scaled by k (k=0.4). Green and red solid arrows show excitatory and inhibitory stimulus, solid blue and dashed red arrows show recurrent excitatory and inhibitory connections, respectively.

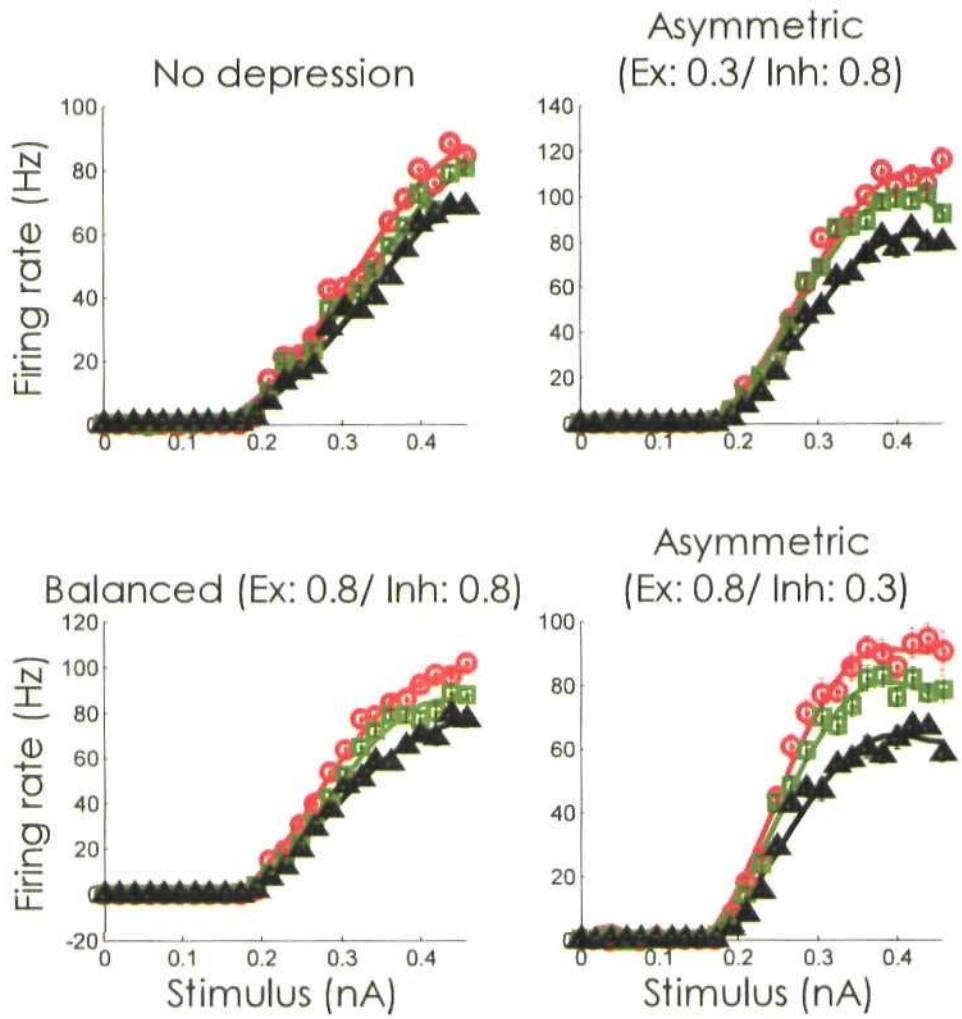

Figure 5.59: Stimulus-response functions for non-linear stimulus input, under all synaptic configurations
Red, green and black lines/symbols plot the low-, mid- and high-variance conditions respectively. **Top left:** No recurrent depression. **Top Right:** Asymmetric depression: 0.3:0.8 (Excitatory : Inhibitory). **Bottom Left**: Balanced: 0.8:0.8. **Bottom Right:** Asymmetric depression: 0.8:0.3.

There was no significant change in the stimulus amplitude that elicited the S_{50} under any condition. The effect of the stimulus on recurrent excitatory and inhibitory synaptic conductances are plotted in Figure 5.60

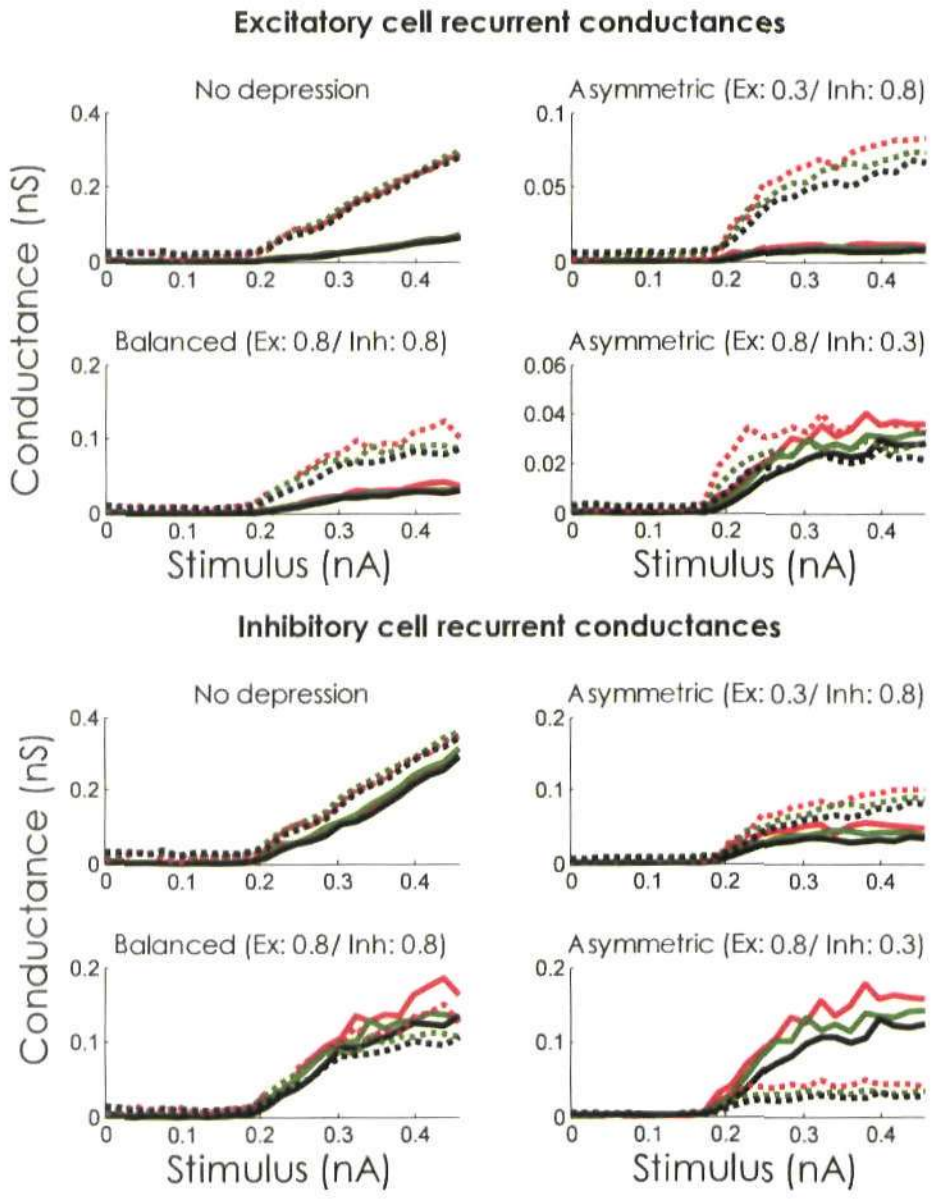

Figure 5.60: Recurrent excitation and inhibition
Each axis, plots the average, recurrent conductances impinging upon either an excitatory (uppermost plots) or inhibitory model cell (lowermost plots), as a function of stimulus input. Red, green and black lines plot the low-, mid- and high-variance conditions respectively. **Top left:** No recurrent depression. **Top Right:** Asymmetric depression: 0.3:0.8 (Excitatory : Inhibitory). **Bottom Left**: Balanced: 0.8:0.8. **Bottom Right:** Asymmetric depression: 0.8:0.3. Solid and dashed lines plot the excitatory and inhibitory conductances, respectively

In the absence of recurrent depression (Figure 5.60, top left), the synaptic conductances were, essentially, linear transformations of the driving current, however in the presence of depression the synaptic conductances tended to saturate. For high intensity stimuli (above the average, non-linear input threshold), the conductance, plotted as a function of stimulus, scaled downwards as adapting variance increased; this is due to greater levels of depression suppressing the synaptic response. As stimulus variance increased the presentation frequency of high-intensity stimuli (above the input filter threshold) also increased. Thus, the synapses became more depressed, relative to lower variance conditions. For the filtered inputs, the low variance stimuli are extremely 'sparse' in nature, thus the synapses had ample time to recover from depression between spikes.

Table 5.7: The slope and stimulus amplitude at the half-maximal point, under all adapting variance conditions (columns) and synaptic configuration (rows)

Excitatory / Inhibitory Depression (E : I)	**Low Variance**		**Mid Variance**		**High Variance**	
	Slope (Hz/nA^2) *SE* $1*10^{-3}$	Stimulus (nA)	*Slope* (Hz/nA^2) *SE* $1*10^{-3}$	Stimulus (nA)	*Slope* (Hz/nA^2) *SE* $1*10^{-3}$	Stimulus (nA)
1 : 1	*0.0073 (±0.298 SE)*	0.3125 (±0.003 SE)	*0.0066 (±0.251 SE)*	0.3146 (±0.003 SE)	*0.0062 (±0.199 SE)*	0.318 (±0.002 SE)
0.3 : 0.8	*0.0137 (±0.488 SE)*	0.2866 (±0.005 SE)	*0.0124 (±0.374 SE)*	0.2782 (±0.004 SE)	*0.0102 (±0.364 SE)*	0.287 (±0.004 SE)
0.8 : 0.8	*0.0112 (±0.368 SE)*	0.2858 (±0.004 SE)	*0.0104 (±0.259 SE)*	0.2893 (±0.003 SE)	*0.0076 (±0.340 SE)*	0.297 (±0.003 SE)
0.8 : 0.3	*0.0117 (±0.478 SE)*	0.2628 (±0.005 SE)	*0.0105 (±0.498 SE)*	0.2630 (±0.005 SE)	*0.0076 (±0.330 SE)*	0.266 (±0.003 SE)

The adaptive response was not abolished by the removal of spike-rate adaptation and tonic inhibition (see Figure 5.61 below); overall, the decrease in slope was reduced but still significant at the 5% level ($p<0.02$) for all depression configurations.

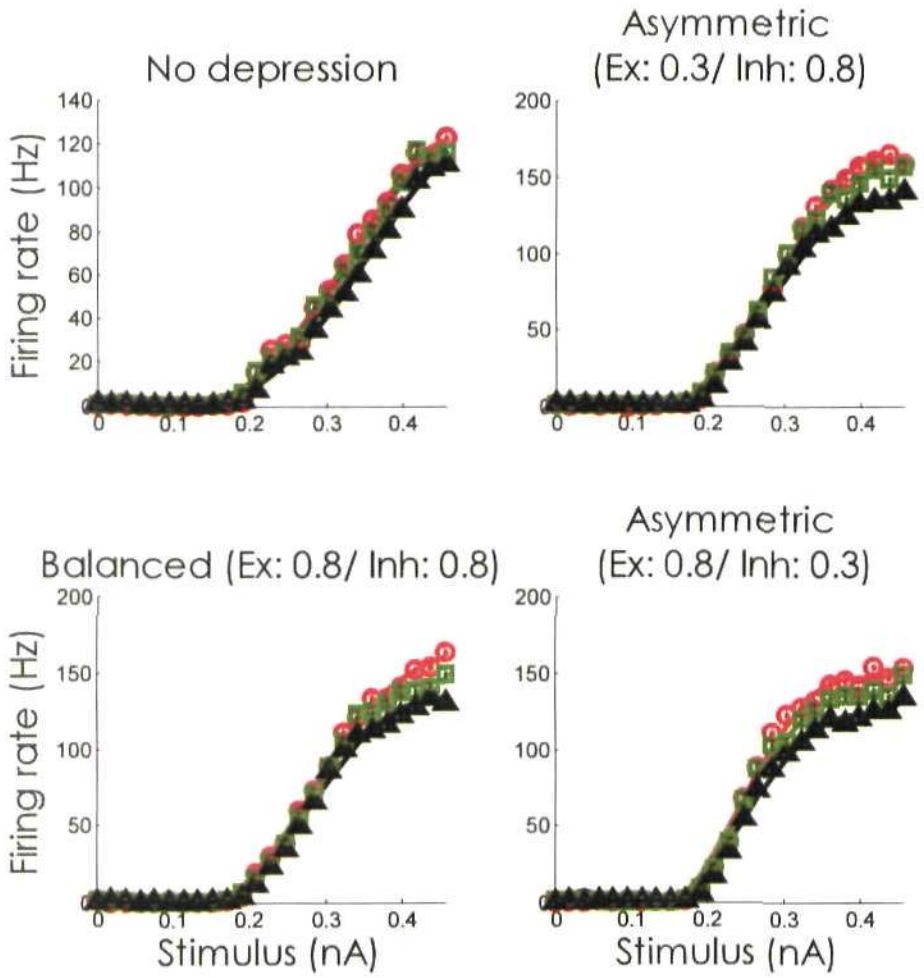

Figure 5.61: Stimulus-response functions for non-linear stimulus input, under all synaptic configurations; no spike-rate adaptation or tonic inhibition
Red, green and black lines/symbols plot the low-, mid- and high-variance conditions respectively. **Top left:** No recurrent depression. **Top Right:** Asymmetric depression: 0.3:0.8 (Excitatory : Inhibitory). **Bottom Left**: Balanced: 0.8:0.8. **Bottom Right:** Asymmetric depression: 0.8:0.3.

In the absence of spike-rate adaptation (Figure 5.61), the read-out cells' output reflects the feedforward input from the recurrent I & F cells within the network. Essentially, the feedforward drive, acting on the read-out cell, is the weighted difference of the excitatory and inhibitory feedforward conductance and is plotted in Figure 5.62. The decrease in inhibition, for each variance condition was proportional to the decrease in

excitation, thus the relative reduction in responsiveness was maintained.

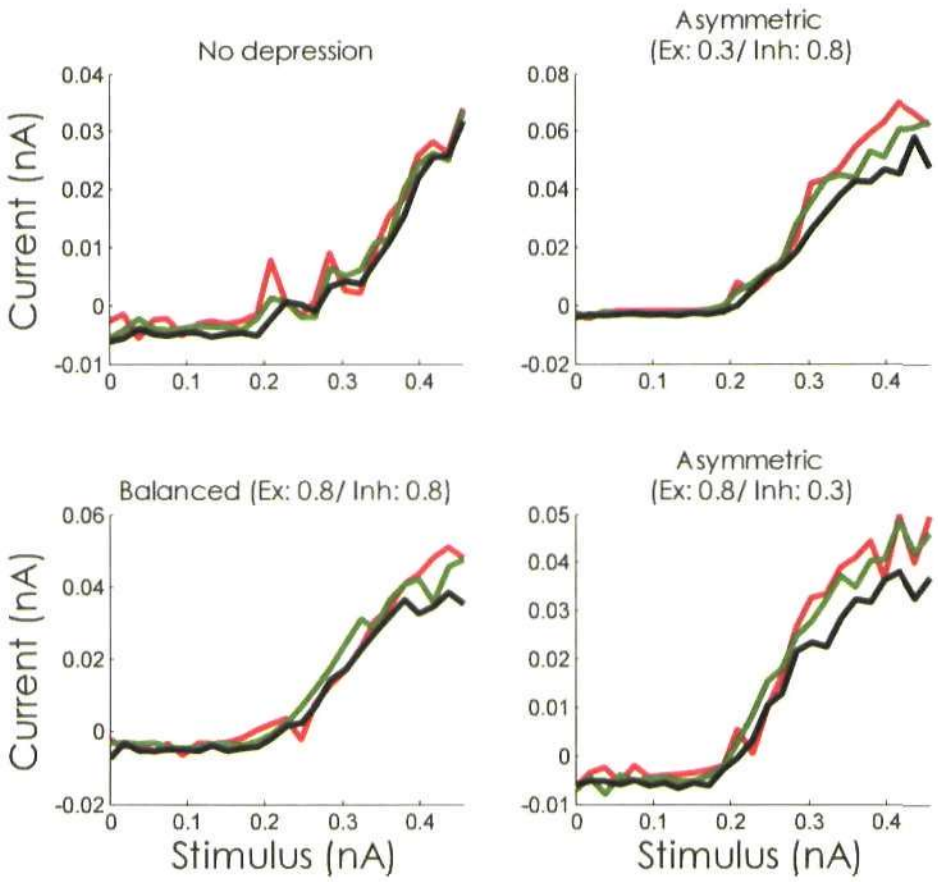

Figure 5.62: Net feedforward synaptic conductance for all depression configurations Red, green and black lines/symbols plot the low-, mid- and high-variance conditions respectively. **Top left:** No recurrent depression. **Top Right:** Asymmetric depression: 0.3:0.8 (Excitatory : Inhibitory). **Bottom Left**: Balanced: 0.8:0.8. **Bottom Right:** Asymmetric depression: 0.8:0.3.

Adaptation to mean

The model was tested under conditions of adaptation to stimulus mean, in order to confirm whether driving the model cell with the filtered input would elicit an adaptive

lateral displacement of the stimulus-response function. The relationship between the stimulus distributions and the summed, filtered input is plotted in Figure 5.63

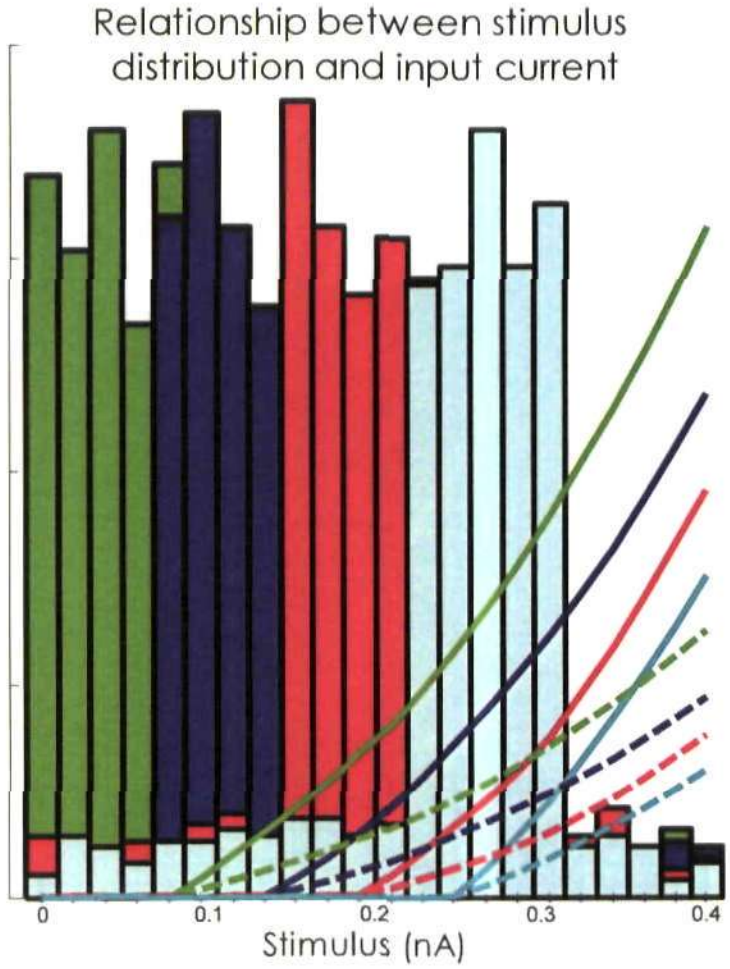

Figure 5.63: Relationship between non-linear input and stimulus distribution under conditions of adaptation to global-mean
The histogram plots frequency of stimulus presentation under the model equivalent of adaptation to global mean. The lines represent the stimulus-current functions, i.e. plot the relationship between driving input and its non-linear transformation. Green bars/lines: high-probability region centred on 0.0573nA; blue bars/lines: 0.1337nA; red bars/lines: 0.2101nA; cyan bars/lines: 0.265nA. Solid and dashed lines plot the excitatory and inhibitory current, respectively.

The response of the model cell, across all depression configurations is plotted in Figure 5.64. The model cell's stimulus- response function was displaced laterally, as a

function of stimulus-mean and was in line with the experimental data. The functions were displaced to centre neural threshold (i.e. the half-maximal point) outside the high-probability regions of the stimulus space.

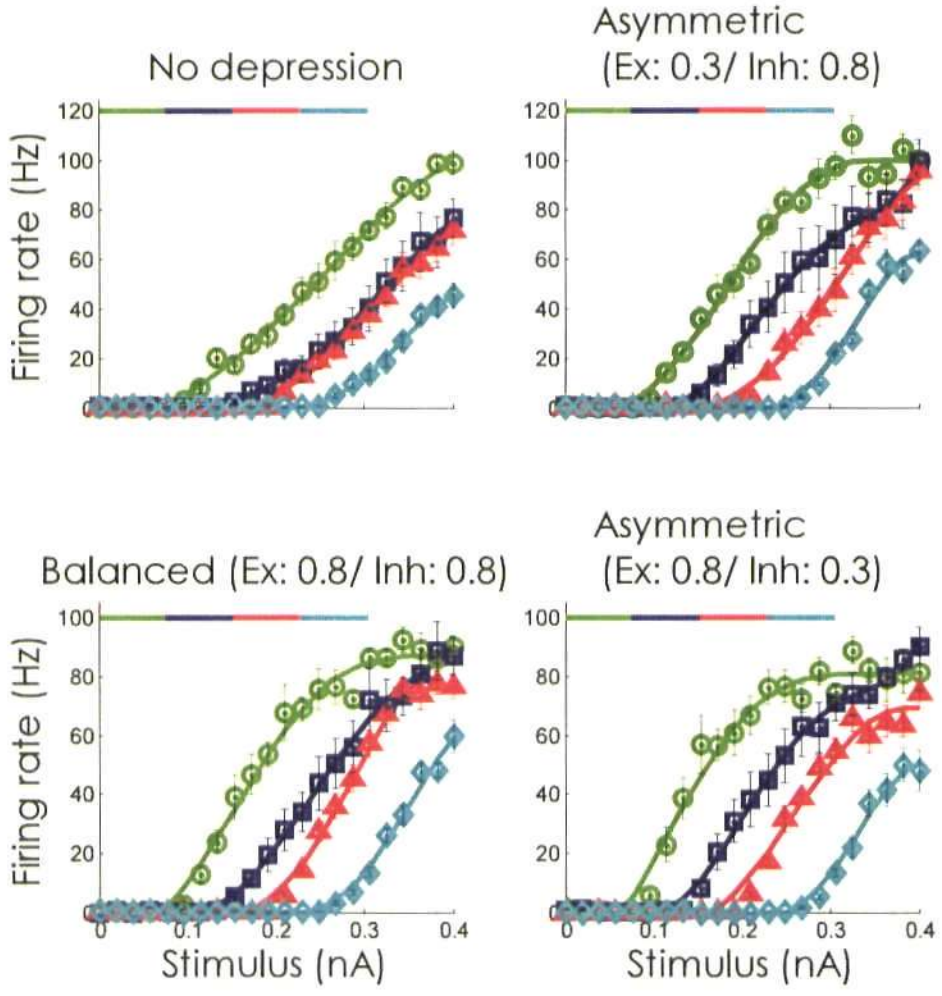

Figure 5.64: Stimulus-response functions, for non-linear input and for all synaptic configurations, under conditions adaptation to global mean.
Green lines/symbol: high-probability region centred on 0.0573nA; blue lines/symbol: 0.1337nA; red lines/symbol: 0.2101nA; cyan lines/symbol: 0.265nA, respectively. Lines plot the best fit to the data (symbols). **Top left:** No recurrent depression. **Top Right:** Asymmetric depression: 0.3:0.8 (Excitatory : Inhibitory). **Bottom Left**: Balanced: 0.8:0.8. **Bottom Right:** Asymmetric depression: 0.8:0.3.

The model step-size and deviation-from-the-mean functions were also calculated. Both the model step-size functions (see Figure 5.65) and normalised deviation functions (see Figure 5.66), across all depression configurations, displayed levels of invariance that were consistent with experimental data.

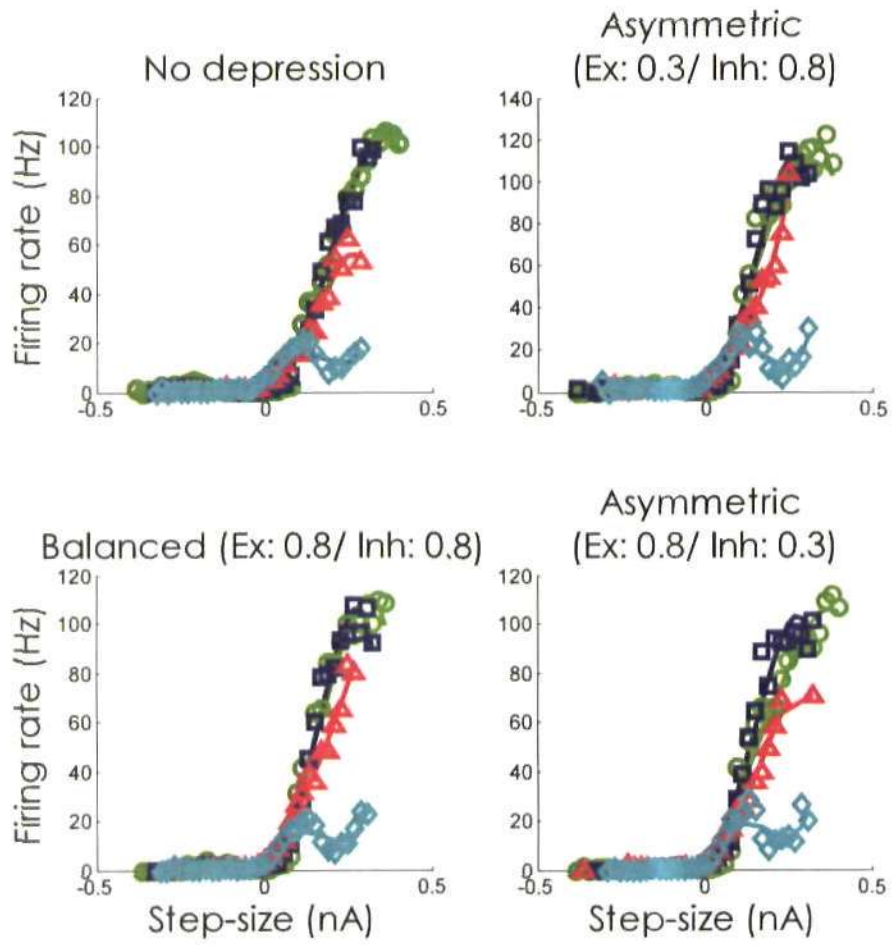

Figure 5.65: Model step-size functions, across all depression configurations; adaptation to mean
Green lines/symbol: high-probability region centred on 0.0573nA; blue lines/symbol: 0.1337nA; red lines/symbol: 0.2101nA; cyan lines/symbol: 0.265nA, respectively. Lines plot the best fit to the data (symbols). **Top left:** No recurrent depression. **Top Right:** Asymmetric depression: 0.3:0.8 (Excitatory : Inhibitory). **Bottom Left**: Balanced: 0.8:0.8. **Bottom Right:** Asymmetric depression: 0.8:0.3.

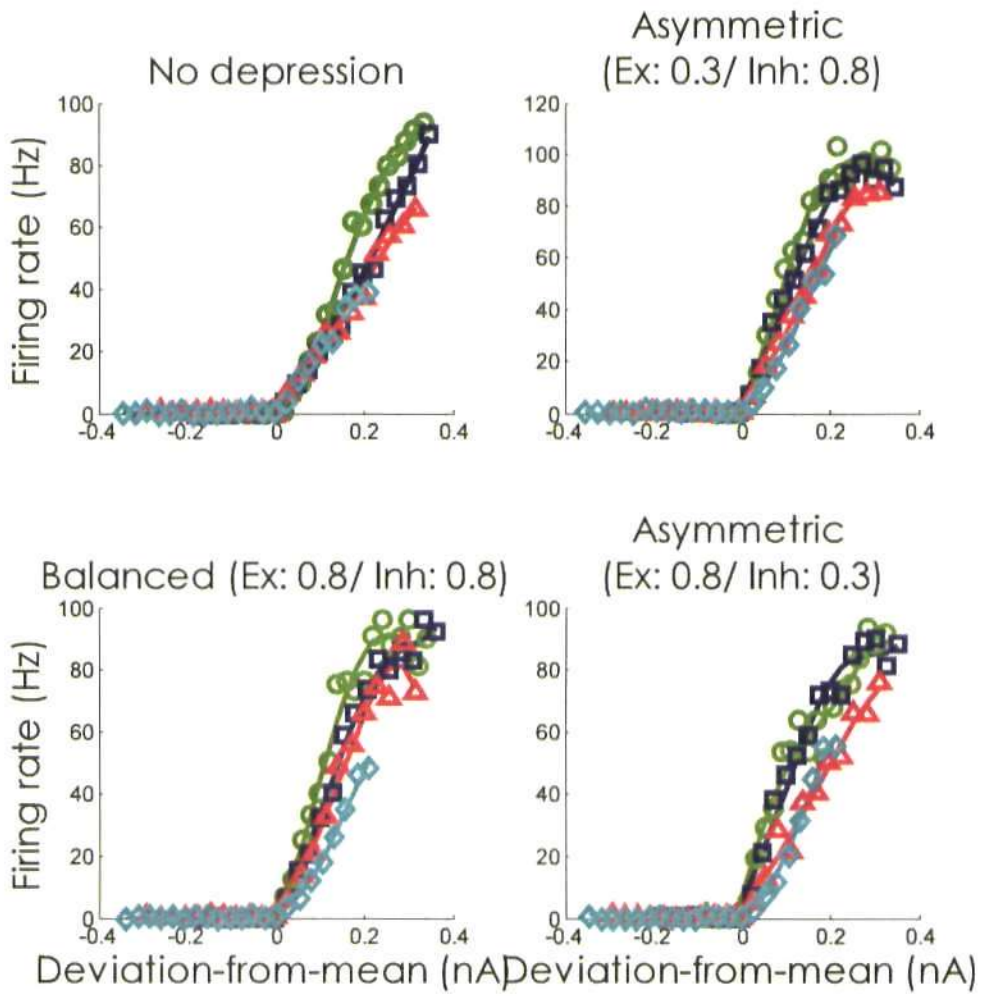

Figure 5.66: Model deviation-from-mean, across all depression configurations; adaptation to stimulus-mean
Green lines/symbol: high-probability region centred on 0.0573nA; blue lines/symbol: 0.1337nA; red lines/symbol: 0.2101nA; cyan lines/symbol: 0.265nA, respectively. Lines plot the best fit to the data (symbols). **Top left:** No recurrent depression. **Top Right:** Asymmetric depression: 0.3:0.8 (Excitatory : Inhibitory). **Bottom Left**: Balanced: 0.8:0.8. **Bottom Right:** Asymmetric depression: 0.8:0.3.

The adaptive response was not abolished by running the simulation without spike-rate adaptation and tonic inhibition (see Figure 5.67). However, the successive increases in degree of stimulus-response function displacement were less pronounced for all depression configurations; especially with reference to the no depression and AS_{EXC} models.

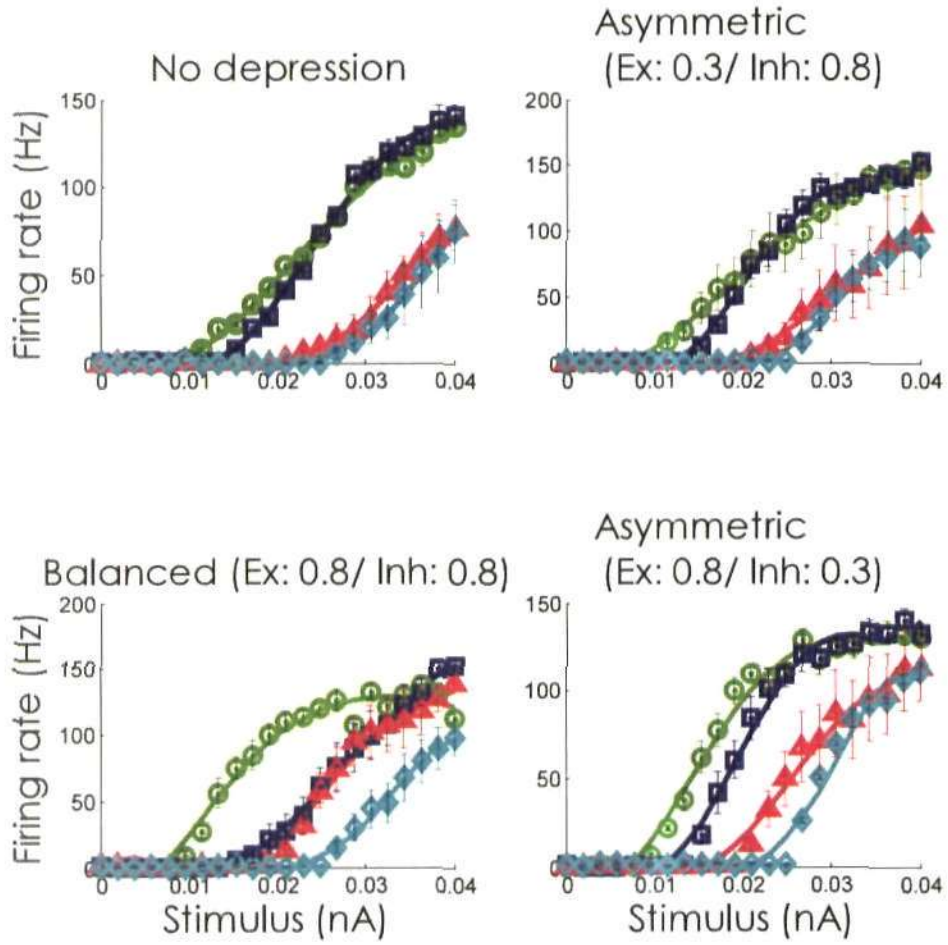

Figure 5.67: Stimulus-response functions, for non-linear input and for all synaptic configurations, under conditions of adaptation to global mean, without SRA and tonic inhibition.
Green lines/symbol: high-probability region centred on 0.0573nA; blue lines/symbol: 0.1337nA; red lines/symbol: 0.2101nA; cyan lines/symbol: 0.265nA, respectively. Lines plot the best fit to the data (symbols). **Top left:** No recurrent depression. **Top Right:** Asymmetric depression: 0.3:0.8 (Excitatory : Inhibitory). **Bottom Left**: Balanced: 0.8:0.8. **Bottom Right:** Asymmetric depression: 0.8:0.3.

5.5: Discussion

The main results derived from the computer simulations of various models of gain control are outlined below:

1. The lateral displacement of the neural rate function is best explained by a mechanism that incorporates both tonic inhibition and spike-rate adaptation

2. Gain modulation can be generated by several mechanisms:

- Increasing the firing rates of balanced, noisy, background conductances acting on the cell (Chance et al. 2002)

- Manipulating the relationship between afferent inhibition and excitation - this was modelled through adjusting the balance of synaptic depression (Chelaru & Dragoi 2008).

-A combination of recurrent and feedforward synaptic depression acting on a non-linear transformation of the driving current

In terms of adaptation to global mean, the computer simulations replicated the experimental data by incorporating tonic inhibition and spike rate adaptation into the model. Both tonic inhibition and spike-rate adaptation act as high pass filters on the incoming stimulus. Efferent excitation must be strong enough to counteract inhibition, (which increases with global mean stimulus levels) and drive the cell over threshold to elicit firing; as the average intensity of the adapting stimulus strengthens, membrane threshold is also increasingly elevated. However, there are differences between the two actions these processes have on the cell; tonic inhibition is essentially a hard-threshold as it is determined by the global mean amplitude of the stimulus, at each trial, and doesn't change during the adaptation process. Conversely, levels of spike-rate adaptation are dependent on the stimulus in a time-dependent manner, as well as the overall global mean of the stimulus. Spike-rate adaptation, through the process of increasing the dynamic threshold, is a self-limiting process. For high intensity stimuli, the dynamic membrane

threshold will increase to such a level that relatively fewer spikes will be elicited to a given stimulus amplitude; thereby allowing for threshold to decay back to its base level. Thus the relative suppressive effects of the dynamic membrane threshold are reduced; spike-rate adaptation alone will not elicit the requisite displacement of the stimulus-response function necessary to generate the invariance of the both the step-size and deviance functions. The introduction of a tonic inhibition increases the stimulus-dependent threshold at which spikes can be elicited; thus the model cell only responds to stimuli of a relatively high intensity, which then initiates the process of spike-rate adaptation. As to the biological validity of the model, there is certainly a wealth of evidence that tonic inhibition is present in the barrel cortex (e.g. Yamada et al. 2007; Kyriazi et al. 1998; Salin & Prince 1996; Kyriazi et al. 1996; Swadlow 2003; Krook-Magnuson, Li, Paluszkiewicz & Huntsman 2008; Krook-Magnuson & Huntsman 2005) and that, within layer 4, it acts to prevent firing to all but the most potent inputs (Swadlow 2002).

In terms of modelling the adaptive response to increases in stimulus variance, three different models were implemented. The first was that of gain modulation through increased background firing rates, as proposed by Frances Chance (Chance et al. 2002). The experimental results were replicated in this model, i.e. gain was decreased, by increasing the firing rates of excitatory and inhibitory background conductances in proportion to stimulus variance. Incorporating tonic inhibition into the model did not affect gain modulation and introducing background conductances into the tonic inhibition model did not adversely effect the lateral displacement of the neural rate-function, as a function of stimulus mean.

The model was not truly adapting in the sense that background firing rates were

explicitly hand-tuned to increase with stimulus variance. The argument was presented that, due to adaptation in the local population of cells providing modulatory inputs, background firing rates would be constant with respect to global mean (due to the relative suppression of output for the high-probability stimuli), but increase with global variance. In order to test this hypothesis, background firing rates were determined by a population of stimulus filters, with parameters set to mimic a population adapted to global mean but not global variance. The overall increase in background noise (averaged over the 30s stimulus presentation time) was sufficient to generate a decrease in gain as a function of variance. Therefore, it can be argued that the model is robust and may represent the mechanism underpinning adaptation to global variance *in vivo*, as well as gain modulation *in vitro*.

A cogent argument for the case that background firing rates could increase with stimulus variance was presented within the introduction to this Chapter (i.e. that background firing rates increase with variance due to the increased number of high-amplitude stimuli presented more frequently). The hypothesis relied on the premise that all cells, providing modulatory input to the target neuron, would be adapting to and correlated with the stimulus. It was found that the average firing rate of the model cell only increased in line with variance when the model cell was unadapted to the stimulus (see Figure 5.38). Numerous cell types can be found within the barrel cortex and not all are prone to gain modulation, e.g. fast-spiking inhibitory interneurons (McCormick, Connors, Lighthall & Prince 1985; Swadlow 2003), thus it is feasible that modulatory inputs could be derived from unadapting sources. The other point of the argument was that background firing rates would remain constant in the face of increasing global mean amplitudes. This implies that modulatory sources would be adapted to global mean levels, thus firing rates to high-

probability stimuli would be suppressed. If one recalls from Chapter 3, adaptation to variance was practically non-existent in the Inferior Colliculus, whilst adaptation to global mean was a common feature of neuronal output. The Inferior Colliculus provides much of the sub-cortical input into the primary auditory cortex, thus if levels of thalamic and cortical adaptation are similar in the somatosensory pathway, one could argue that a possible source of modulatory input, that only adapts to stimulus mean, is derived from cells that inherit their adaptive properties directly from thalamic input.

Another point is that random input by its very nature is not supposed to be correlated with any stimulus, thus it is arguable whether one can consider the background conductances as purely stochastic! Instead one can argue that the firing rates of spike trains, with Poisson statistics, must increase in a balanced manner for both modulatory excitation and inhibition, and that firing rates must be dependent on the overall variance in the adapting stimulus and **not** global mean.

Whilst the model described above relied on increasing the levels of modulatory excitation and inhibition in a balanced manner, the second model relied on the relative balance between inhibition and excitation to elicit gain modulation (Chelaru & Dragoi 2008). The evidence for asymmetry in both synaptic dynamics (i.e. Varela et al. 1999, *Galarreta & Hestrin, 1998) and overall influence of inhibition (e.g. McCormick1985)* was highlighted in the introduction. Changes in the relationship between afferent inputs was achieved by manipulating levels of synaptic depression, such that excitatory levels of depression became weaker as stimulus variance strengthened, thereby becoming closer to inhibitory depression. Whilst it seems counterintuitive to elicit a reduction in neural responsiveness by increasing the amount of excitation in the network, the increased

excitatory drive acted on recurrent inhibition, resulting in increased levels of inhibition within the network, relative to the low-variance case, and thereby decreasing gain. Adaptation to variance was successfully simulated in the steady-state case (as was adaptation to stimulus mean), and less successfully in the time-dependent case, in that the model cell displayed reduced responsiveness as a function of increased variance, but there was no steady decrease in slope values, at the S_{50}. However this could be due to the limitations in the implementation of the model; as it did not strictly elicit gain change in the same manner as the steady-state model. Inhibition levels increased in the same manner for both implementations of the model but within the time-dependent version excitation was highest under the low-variance condition, this could account for lack of clear decrease in neural slope.

Gain modulation was achieved through explicit hand-tuning of the depression factor, *p*. It was decided to see whether increased variance in the stimulus would be sufficient to elicit a release from depression and thus a change in gain. Experimentally, the amount of synaptic resource available is highly dependent on presynaptic firing rates (e.g. thalamocortical depression is often saturated due to tonic activity in the thalamus, Castro-Alamancos & Oldford 2002), thus a weakening of depression could be linked directly to the stimulus configuration; as the high-probability region widens, the number of low intensity current amplitudes presented increases, thereby allowing for more recovery from depression.

The model cell was tested under several distinct synaptic arrangements (and, for one simulation, the addition of facilitation on recurrent excitatory-to-inhibitory synapses) but no significant change in gain was observed. Synaptic dynamics were too slow to capture the

local changes in stimulus amplitude.

Decrease in gain was only observed following the introduction of tonic inhibition and spike-rate adaptation into the model read-out cell and only for the asymmetric depression configuration (where excitatory depression was stronger than inhibition) with facilitation. Tonic inhibition has been shown to elicit multiplicative changes in neural gain in the presence of synaptic depression (Murphy & Miller 2003, Rothman et al. 2009). Under the AS_{EXC} configuration, facilitation increased excitatory-to-inhibitory conductances, in proportion with stimulus variance; thus eliciting a significant decrease in neural gain. Therefore, only under very strict conditions, could the inherent diversity in the adapting stimulus elicit the requisite change in synaptic dynamics to influence neural gain.

For the final model, it was decided to drive the depressing I & F model (as described above) using a non-linear transformation of the stimulus, with the aim of approximating feedforward input from cells that were already displaying adaptation to global mean. The nature of the non-linear input meant that under the low-variance adapting condition, only a few stimulus inputs would be strong enough to elicit a response in the model cells. As adapting variance increased, the number of synaptic events also increased thereby generating greater levels of depression and reducing the overall synaptic response to a given stimulus. The effect was observed in both recurrent inhibitory and excitatory model cells. However the relative decreases in excitation and inhibition were proportional, thus the overall net effect on the model read-out cell was a reduction in driving input as stimulus variance increased.

The decrease in gain was more pronounced in the presence of spike-rate-adaptation and tonic inhibition. Spike-rate adaptation acts to displace the stimulus-response function

rightwards in the presence of higher amplitude stimuli and, as discussed, tonic inhibition has a multiplicative effect on the stimulus-response function. These mechanisms enhance the gain changes induced by depression in the recurrent circuit and initiate gain changes for lower-intensity stimuli than that observed under depression alone.

The model above assumes that, *in vivo*, the majority of the driving input will be both tuned to the stimulus and adapted to global mean. As already stated, within the auditory system, this closely matches the response profile of inferior colliculus cells, thus, if levels of adaptation to intensity are the same in the subcortical afferent pathway of the somatosensory system, this is a valid assumption. The model also requires synaptic depression, however again this is common feature of synaptic dynamics within the somatosensory system and other primary cortical areas.

Unlike the extended models of background noise and asymmetric depression, gain changes were initiated in the absence of spike-rate adaptation and tonic inhibition; thus under this mechanism neuron types which don't display firing-rate adaptation would still be sensitive to changes in stimulus-variance. Multiplicative gain modulation emerges from the model architecture and dynamics without the need for any other complex processes. Stimulus intensity or contrast is a relatively low-level feature of environmental input, hence the argument for it being encoded as efficiently as possible (Barlow 1961). The evidence presented in Chapter 4 indicates that neurons are adapting to encode local changes in the stimulus input, thus a model of gain change that emerges simply from synaptic dynamics represents a robust and economical mechanism to account for gain change in the cortex (Borst, Flanagin & Sompolinsky, 2005)

The model is of course similar in many features to mechanisms of gain modulation

proposed by Murphy & Miller (2003) and Rothman and colleagues (2009), where the combination of increased tonic inhibition, a non-linear input and synaptic depression leads to a decrease in neural gain. In the model outlined above, there is no increase in tonic inhibition instead the driving inputs become saturated for higher intensity stimuli and become relatively more suppressed as variance increases.

Finally, the model above suggests that adaptation to stimulus statistics in the barrel cortex is inherited primarily from the thalamic inputs arriving in layer 4, thus one would expect that the adaptive response would emerge rapidly, within the cortex. In Chapter 3, the analysis of the adaptation time course was inconclusive. However, it may be that adaptive responses to stimulus statistics are developing on too rapid a timescale to be observed within the experimental data presented within this thesis.

Chapter 6: Discussion

6.1: Summary of findings

In order to summarise the main conclusion of this thesis, the main points from each relevant chapter are presented below.

6.1.1: Chapter 2: Information theory and contrast adaptation

Within this chapter the main ideas underpinning Information Theory were introduced followed by a description of Barlow's Efficient Coding Hypothesis (Barlow 1959; Barlow 1962; Barlow 1972). Barlow saw that nervous systems operate under similar constraints to communication systems, in that environmental information is not only transformed into a binary signalling code (i.e. the all-or-nothing action potential) but is transmitted throughout the nervous system along finite channels and must therefore undergo an enormous amount of compression. Barlow devised the theory of redundancy reduction as a tool by which sensory information can be compressed as efficiently as possible, with minimal loss of information; sensory input can therefore be received further up in the neural hierarchy.

The Efficient Coding Hypothesis predicts that cells should have a dynamic response to the environment and adjust their output depending on the prevailing stimulus levels. Luminance contrast is a measure of the standard deviation of light intensity from mean luminance. There is strong evidence to suggest that neurons within the visual systems, across many species, adapt to centre their range of responsiveness onto the adapting contrast. As such, the contrast-response within the visual system was used to highlight some of the key features of gain control, e.g. the stimulus-response function and

the parameters used to describe it, as well as some putative mechanisms.

Extensions to the Efficient Coding Hypothesis were also introduced (albeit briefly), e.g. Infomax principle by which individual neural output is adjusted, in a stimulus-dependent manner (and/or neural systems as a whole are organised) to maximise the mutual information between input and output (where mutual information, $I(x \mid y)$, refers to the amount of information an output y can provide about an input x). Some alternatives to the Efficient Coding Hypothesis, as a stimulus encoding strategy, were also introduced, namely novelty and deviation detection.

Finally, Fisher Information was described. Formally, Fisher Information is *defined as the variance of the score (the negative, log-likelihood of observing x, given the* parameter set, θ, on which x is dependent). The inverse of the Fisher Information is the Cramer-Rao inequality, which defines the lower-bound of the variance of an unbiased maximum-likelihood estimator of θ; the inverse of the square-root of the Fisher Information defines the lower bound of the standard error of an unbiased maximum-likelihood estimator of θ. The Fisher Information of the neural, stimulus-response function quantifies the accuracy with which a cell is able to encode the stimulus. Fisher Information is highest on the slopes, rather than peaks, of the stimulus-response function and the inverse of the Fisher Information defines the discriminability threshold, i.e. where the cell is able to signal the just-noticeable-difference between θ and $\theta+\Delta\theta$.

6.1.2: Chapter 3: Adaptation to stimulus statistics in barrel cortex

Within this chapter the experimental adaptation data, from the barrel cortex, was introduced.

Recent evidence was presented that adaptation to intensity (i.e. adaptation to stimuli analogous to contrast in the visual system) is present in sensory modalities other than vision. Two studies were highlighted, in which sensory neurons were adapted using stimuli that allowed for dissociation between global mean and variance levels, thus adaptation to distinct measures of stimulus statistics could be recorded. Adaptation was observed in the rodent inferior colliculus (IC, Dean et al. 2005) and the barrel cortex (Garcia-Lazaro et al. 2007)

Whilst adaptation to global mean elicited a shift in stimulus threshold, as predicted by the Efficient Coding Hypothesis, adaptation to variance, elicited little adaptation in the IC. For those cells that did adapt, neural gain tended to increase with stimulus variance, in direct contradiction of the Efficient Coding Hypothesis. Conversely, increased adaptive gain, in response to relatively noisier adapting stimuli, was highly pronounced in the barrel cortex.

Jan Schnupp generously provided the experimental data, from the barrel cortex, which formed the basis of the analyses in this and the following chapter. The reanalysis of the experimental data revealed that independently increasing stimulus variance led to a decrease in gain of the population response, in direct contradiction to the published reporting of the experimental data. These results have been accepted by the authors of published data (Garcia-Lazaro et al. 2007). A further analysis of the data did not reveal any information about the time course of adaptation, suggesting that either the data was not suitable for this type of analysis or that adaptation to intensity, within the barrel cortex, developed very rapidly (with a latency of <400ms).

6.1.3: Chapter 4: Encoding strategy underpinning adaptation

Under conditions of adaptation to the global-mean amplitude of whisker deflection, sensitivity was greatest for stimuli outside the high-probability region (Garcia-Lazaro et al. 2007, Dean et al. 2005). This led to the suggestion that the adapted cells were primed for novelty detection (Dean et al. 2005).

Within somatosensory cortex (Garcia-Lazaro et al. 2007), stimulus-response function displacements were more extreme than those observed in the inferior colliculus (IC); it was therefore decided to test the hypothesis of novelty detection and other possible coding strategies, with the purpose of discovering which, if any, best described the adaptive response to global mean. Certainly, the adaptive response, within the barrel cortex, did not fall entirely within the expected bounds of the Information Maximisation principle.

The analysis based on stimulus-specific adaptation was inconclusive. Instead, it was found that there was a clear dependency of spiking activity on stimulus history. In general, neurons tended to fire in response to large positive differences between the stimulus amplitude at the time of firing and the one directly preceding it. However, the dependence on stimulus history could be found from approximately 280ms prior to spike firing. The general stimulus pattern preceding a spike was best described as a continuous decrease in stimulus intensity to below global mean amplitude until ≈120-80ms prior to spike-firing; followed by a slight increase in intensity, before the final, large step-size difference in amplitude that elicited the spike.

The average firing rate response, as a function of the relative difference in amplitudes between the stimuli presented within a given 40ms epoch of time and the one

preceding it (referred to as the step-size function), were plotted for all cells in the population. The population analysis revealed that adaptation to global mean acted to scale the neural-response function such that they were always invariant, with respect to step size. However, for those cells adapted only to stimulus variance, the step-size function was not invariant. The overall conclusion was that for those cells that were only adapted to stimulus variance, the step-size functions showed some dependence on the width of high-probability region in the stimulus space.

Amongst both populations (adaptation to mean or variance) some cells showed a preference for step-size irrespective of the direction of amplitude change and were classed as novelty-preferring cells. The majority of novelty-preferring cells had a response bias for positive step-changes; however a small proportion of these neurons displayed symmetrical functions. Novelty-preferring cells that displayed symmetry had longer time-to-peak and response-decay latencies than the remainder of the population.

The data was also analysed with respect to local, rather than global, variance. The average response was calculated, as either a function of the difference (deviation-from-the-mean) or the squared difference (variance-from-the-mean) between the input, at each time step and global-mean of the stimulus (averaged over the 30s presentation time). The population response showed a clear preference for positive deviations-from-the-mean. Under conditions of adaptation to global-mean, the deviation-from-the-mean functions for all 4 adapting conditions increased monotonically from local standard deviations of >0μm, and diverged for positive deviations-from-the-mean of > 200μm. The divergence was due to the skewed nature of the stimulus distribution; global standard-deviation decreased as global mean increased and local deviations ranged from -126.08 to

593.92μm for the 90μm condition, -217.68 to 502.32μm for the 210μm condition, -310.24 to 409.76μm for the 330μm condition and from -401.60 to 318.40μm for the 450μm condition. Normalising the functions, revealed that they were scaled according to levels of global standard deviation

Under conditions of adaptation to global-variance, neurons also displayed a preference for positive deviations-from-the-mean. The response functions tended to decrease in both slope-gain and maximum firing rates inversely with global-variance levels. Normalising the deviation-from-mean functions resulted in a considerable degree of overlap for the mid- to high-variance conditions, except for the positive tails of the functions, but not for the low-variance condition.

In conclusion then, adaptation to stimulus mean acts to rescale responses so as to maintain information transmission, about the local changes in stimulus amplitude, irrespective of global mean values. Adaptation to variance may act to allow for the encoding of the width of the high-probability region of the stimulus distribution; as suggested by the adaptive decrease in neural gain (as a function of either local stimulus intensity or deviation from global-mean) and the slight widening of the step-size functions.

6.1.4: Chapter 5: Mechanisms underpinning adaptation

The aim of this chapter was to explore some of the mechanisms underpinning adaptation, through the use of computer simulations. The majority of simulations were implemented using a conductance-based, Integrate & Fire, neuron developed by Alain Destexhe (1997).

Adaptation to global mean and variance were found to be distinct in terms of

their effects on neural output, thus, it was argued, they may be underpinned by distinct mechanisms; with the caveat that both mechanisms must be able to coexist, as stimulus modality was unchanged.

The adaptive shift in the neural-stimulus response function and the invariance of the step size function was best approximated by a model consisting of tonic inhibition and spike-rate adaptation; spike-rate adaptation was executed as a dynamic firing threshold mechanism.

Gain modulation was implemented in several models: increasing the firing rates of background conductances in proportion to stimulus variance (Chance et al. 2002), adjusting levels of synaptic depression in recurrent excitatory and inhibitory units, such that the difference between afferent inhibition and excitation was reduced (Chelaru & Dragoi, 2008) and by driving depressing, recurrent excitatory and inhibitory units with a non-linear transformation of the stimulus.

By increasing background firing rates in proportion to stimulus variance, it was possible to simulate the adaptive decrease in gain in the model cell. The simulations performed here indicate that this model of gain modulation can be applied to the *in vivo* case, with the following caveats: noisy, background firing rates can only increase with stimulus variance, therefore the presynaptic inputs, which generate background Poisson activity, must show dependence on the stimulus.

Within the introduction to Chapter 5, several models were reviewed in which the gain modulation was dependent on the presence of synaptic plasticity and the relative relationship between excitation and inhibition. A model of asymmetric synaptic depression was implemented, in order to generate these relative changes in inhibition and

excitation.

The model was tested in the both a steady-state and time-varying case. Excitatory depression was weakened as stimulus variance increased, thereby increasing the excitatory drive acting on afferent inhibition and reducing gain in the model cell. The output of the asymmetric depression model of gain modulation was not as robust as for the Frances Chance model, highlighted above.

It was hypothesised that the stimulus configuration itself would allow for release from depression, under the higher-variance conditions; thus depression levels would become increasingly balanced. However, the inherent variance in the stimulus was not sufficient to generate the relative relief from depression necessary to initiate a multiplicative decrease in neural-gain. It was argued that the temporal dynamics of depression were too slow to fully capture the changes in input signal on the 40ms scale. The addition of fast-facilitation (the time constant of recovery was 100ms, compared to 300ms for depression) onto the excitatory-to-inhibitory recurrent synapses did elicit a slight increase in recurrent excitatory drive as stimulus-variance increased, thus shifting the balance of the network towards inhibition, but only when recurrent excitatory connections were relatively more depressed than the inhibitory synapses . In the presence of spike-rate adaptation and tonic inhibition, acting on the read-out cell of the model, facilitation was sufficient to generate a significant decrease in neural gain, inversely to stimulus-variance levels.

The combination of tonic inhibition and synaptic depression has been shown to elicit multiplicative gain changes in the cerebellum (Rothman et al. 2009). The effects of facilitation (an overall reduction in firing-rates as a function of input signal) were

therefore enhanced by the non-linear interaction of tonic inhibition and feedforward synaptic depression.

It was decided to use a non-linear input to drive the I & F model directly. The stimulus current was filtered through a bank of parallel non-linear rectified filters, prior to injecting it into the model (the filters were identical to those used to drive the background modulatory conductances in the earlier model of gain modulation; the crucial feature was that the absolute threshold of each filter was set to be slightly higher than the global-mean value of the stimulus). It was hoped that the increase in overall levels of afferent input, for the higher variance adapting conditions, would be strong enough to shift the bias of the recurrent network towards inhibition.

There was a significant decrease in the slope of the read-out cell under all synaptic configurations (no recurrent depression, asymmetric depression with either excitatory depression stronger or weaker than inhibition and balanced depression configuration). There was no significant change in the stimulus amplitude that elicited the S_{50} under any condition. The presence of spike-rate adaptation and tonic inhibition, acting on the read-out cell, strengthened gain modulation.

6.2: General Discussion

6.2.1: Information theory and Efficient Coding Hypothesis

Information Theory arose as a method for quantifying information transfer along artificial communication systems (Shannon 1948). Horace Barlow, along with several other scientists (e.g. Attneave 1954; Craik 1948), made the important link between Information Theory and biological nervous systems, which are essentially faced with

many of the same issues as communication systems. Nervous systems must transform and compress vast amounts of information along highly restricted channels for receiving and processing of sensory input, further-up the neural hierarchy.

According to the principles of Efficient Coding Hypothesis (Barlow 1961), cells will adjust their responsiveness, in a stimulus-distribution dependent manner, such that information about the most common stimulus feature can be encoded with the minimum number of spikes; the main principle of Information maximisation holds that the underlying encoding strategy is to maximise the amount of mutual information a spike carries about the input that elicited it, thereby minimising the number of action potentials needed to encode sensory input accurately. Spikes can carry information about the stimulus in their average firing rate, over time; thus, the optimal solution (whether the underlying principle is maximising either mutual information or redundancy reduction) is for neural rate-functions to be centred on the global mean and broad enough to encode the full stimulus distribution. An example of a stimulus feature that must be encoded as simply and efficiently as possible is visual contrast, as it is a universal component of the visual scene.

As highlighted in Chapter 2, there is wealth of evidence to suggest that the encoding of contrast is based on Information Theoretical principles. For example contrast adaptation, shifts the neural contrast-response function such that cellular threshold (as determined by the half-maximal point on the tuning function) is set to the adapting contrast (Ohzawa et al. 1985). The contrast-response function itself is sigmoidal in nature, and reflects the Gaussian distribution of stimulus contrast within the environment (Laughlin 1981).

Of course, as been elucidated within this thesis, encoding of whisker deflection amplitude (a stimulus feature analogous to visual contrast, as it is also a measure of fluctuation about a global mean) is not strictly within the bounds of information theory. Certainly, adaptation to variance was in line with the principle of Efficient Coding Hypothesis in that the slope (or neural gain) of the neural rate function, for the majority of neurons within the population, decreased (as measured at the half-maximal point) as stimulus variance increased. The decrease in gain resulted in a much broader slope thus allowing for adapted neurons to respond, in a dynamic manner, to amplitudes within the high-probability distribution.

In terms of adaptation to stimulus mean, adaptation centred the population stimulus-response function to a location in the stimulus-space that was always just outside the high-probability region. The Fisher function also peaked just outside the high-probability region.

The Fisher function is essentially a measure of how well one can estimate the stimulus parameters from the neuronal response; thus, it was clear that, over the population, adaptation served to improve the transmission of sensory input that was of a higher intensity than the prevailing stimulus levels.

6.2.2: Encoding strategies

The location of adaptive neural threshold and peak Fisher Information tends to suggest that as well as encoding general stimulus levels, neurons may be primed for novelty detection (Dean et al. 2005). However, it was found that novelty detection was not necessarily the underpinning encoding strategy. Instead neuronal output was scaled,

across the different mean adapting levels, to maintain the relationship between firing rate and changes in the stimulus amplitude between one 40ms stimulus step and the next. For the majority of cells within the population, this relationship was biased towards positive step-sizes only, i.e. adapted neurons did not respond to differences in amplitude steps that were less than zero and the firing rate increased almost linearly as the gap between concurrent stimulus-step amplitudes widened.

Due to the stimulus configuration, the majority of step sizes encountered were in the range of ±150μm. The strongest firing rate responses were found for step-size differences outside this range. This was also reflected in the average spike-triggered stimulus which revealed that a firing event was most likely to occur when the preceding stimulus intensity was much weaker than the stimulus that elicited the spike (and that spike-triggered input was generally stronger than the global mean of the adapting stimulus).

Barrel cortex neurons were also sensitive to how far the stimulus, at each time step, deviated from global-mean levels. The response functions scaled up to each other when normalised by global levels of standard deviation, thus it would appear that neurons were not interested in the amplitude of sinusoidal whisker deflection, *per se*, but in the changes in deflection intensity every 40ms.

This has parallels with a recent paper by Ringach and Malone (2007) in which it was found, in monkey V1 at least, that sensory cells appear to act as deviation detectors, which amplify large signals at the expense of background noise. The discovery that the majority of neurons, within the barrel cortex population, had adaptive step-size and deviation-from-the-mean functions that were invariant, with respect to stimulus mean,

builds on a slowly growing body of evidence (e.g. Ringach & Malone 2007; Wimmer et al. 2008) that information maximisation may not be the only strategy underpinning the encoding of simple stimulus features, such as contrast or whisker deflection amplitude.

There was some evidence that a small number of cells were acting as novelty detectors, in that the step-size functions increased away from a local minimum value (close to zero) in both directions. Of this subsample, only four cells had symmetrical functions, implying that these neurons responded to the absolute value of the step-size difference in amplitude; responses were still invariant with respect to stimulus mean. However it could not be ruled out that the responses were due to latency differences between these cells and the majority of the population. The decay response to large step-size differences in amplitudes tended to be longer in the novelty-type cells, thus it could be argued that the apparent equal preference for negative step-sizes could be due to the response being inherited from the decay of the response to the preceding amplitude (which of course would be of stronger intensity).

The other interesting facet of adaptive behaviour to be revealed from this analysis was that the step-size functions of neurons, adapted to stimulus variance, were not invariant. The normalised deviation-from-mean functions also tended to diverge for strongest, positive deviations.

This implies that two different strategies underpin the encoding of stimulus variance and deviations-from-the-mean. Neurons encode the stimulus distribution through adjustments in neural gain and local deviations-from-mean by neural threshold. One can argue that neurons are able to adjust their encoding strategy depending on the adapting conditions. Under adaptation to increasingly noisy stimuli, neurons shift from a

regime of encoding deviation to one of encoding the stimulus distribution.

6.2.3: Mechanisms of adaptation

Adaptation to stimulus mean was modelled by incorporating a spike-rate adaptation mechanism and increased levels of tonic inhibition into the model. Spike-rate adaptation has been implicated as a possible mechanism for contrast adaptation (Sanchez-Vives et al. 2000) where it is mediated via sodium- and/or calcium-dependent, hyperpolarising, potassium current. The current is strongest after spiking and prevents repolarisation (it is termed the after-hyperpolarising current, or I_{AHP}), thereby reducing the ability of the cell to fire, with the same rate, to a given stimulus.

Within the model, spike-rate adaptation was implemented as a dynamic firing threshold that increased with every spike and decayed exponentially back to the absolute firing threshold between spikes. Thus for every spike fired, the threshold for the next spike (within a given time window) would have increased, thereby suppressing firing. Spike-rate adaptation generated a lateral displacement in the model cell's stimulus-response function (as per Garcia-Lazaro et al. 2007), but not invariance of the step-size function.

Spike-rate adaptation, via a mechanism of dynamic threshold, is essentially self-limiting process and is dependent on the stimulus in a time-dependent manner. As membrane threshold increases, the ability of the cell to fire is reduced thus allowing for membrane threshold to begin recovering back to base level. As stimulus variance becomes greater, the number of sub-threshold, as well as supra-threshold, stimuli increases, also allowing for recovery of the membrane threshold.

It was thus decided to incorporate tonic inhibition into the model cell. This was inspired partly through replicating a model developed by Murphy & Miller (2003), in which lateral displacement of the model stimulus-response function was achieved by increasing tonic inhibition and having a linear relationship between stimulus and driving current. The amount of tonic inhibition was directly related to the global mean amplitude of the stimulus and was therefore independent of local changes in stimulus intensity.

In the presence of both tonic inhibition and spike-rate adaptation, the degree of stimulus-response function displacement in the model cell, was comparable to that observed experimentally, especially with reference to the first three adapting global-mean conditions. Both model step-size and normalised deviation functions displayed invariance and the peak Fisher Information functions were centred on amplitudes outside and greater than the high-probability region.

Correspondence with the experimental data was less pronounced for the final mean adapting condition. In the experimental data, the step-size function for the strongest global-mean level was significantly different from the middle two high-probability conditions; however the degree of difference in the simulated data was greater than that observed experimentally. The model cell, peak Fisher function was located within the high probability region of the stimulus space, thus suggests the degree of displacement in the stimulus-response functions was not comparable to that observed experimentally.

It could be that the range of stimuli used to test the model cell was not wide enough. The rate-functions of all barrel cortex cells, following adaptation, were displaced to lie outside the high-probability region and reached response maximum at ≈600μm; this possibly represents the biophysical response limits within the barrel cortex. The model

cell is not subject to the same constraints as biological neurons, therefore the stimulus range may not have been extensive enough to drive the model cell to reach its firing limits.

As to the biological validity of the model, there is strong evidence to suggest that spike-rate adaptation is involved in the suppression of responses observed in response to sustained stimulation, within V1 (Sanchez-Vives et al. 2000). Certainly there is evidence that a tonic hyperpolarisation can elicit contrast adaptation (Carandini & Ferster 1997), however there is also evidence that blocking GABA inhibition has no effect on adaptation (e.g. Vidyasagar 1990; McLean & Palmer 1996). The important factor is how the mechanism operates: it is essentially a high pass filter, thus only those stimuli strong enough to counteract inhibition will elicit a response in the model cell. As inhibition increases with global mean, only relatively higher intensity stimuli will pass the inhibitory filter, thereby initiating an increase in firing threshold via spike-rate adaptation. The combination of the two factors (tonic inhibition and spike-rate adaptation) scales the output of the cell to be invariant with respect to stimulus step-size.

Adaptation to variance was also simulated by several different models. All the experimental results were replicated by a model of increasing the firing rates of background excitatory and inhibitory, stochastic conductances in a balanced manner (Chance et al. 2002), in proportion to stimulus variance. The model replicated the experimental data even in the presence of tonic inhibition and incorporating random background conductances into the model of adaptation to stimulus mean did not affect output adversely.

Therefore this model represents a universal mechanism for generating gain

control following both adaptation to either mean or variance. The model relies on background firing rates increasing with stimulus variance only. If one considers the stimulus configuration, it is fair to assume that firing rates would increase with stimulus variance as the number of high intensity stimuli presented frequently increases, especially if the sources of background noise are already adapted to global mean amplitude. However, this caveat only holds if one assumes that background activity is derived from cells that are also un-adapted to variance. In order to test this hypothesis, background firing rates were derived from a bank of non-linear filters, which approximated the activity of background noise sources that were adapted to global mean. Non-linear filters were set to have thresholds that were greater than the global-mean of the stimulus and the outputs were used to create Poisson-generated background conductances. There was no difference in the average level of either excitatory or inhibitory background conductance, as a function of stimulus amplitude; however, averaged over the stimulus presentation as a whole, conductance levels did increase overall with stimulus variance. The model replicated gain modulation, but only with tonic inhibition and spike-rate adaptation.

Within the Frances Chance model (Chance et al. 2002), gain modulation arose through a combination of shunting inhibition eliciting a lateral shift in response function and increased noise enhancing firing rates for lower amplitude stimuli (see Figure 6.1). Within the model of gain modulation presented above, the small increase in global-mean levels, due to increased stimulus-variance, elicited a slight lateral displacement in the stimulus-response function; this combined with slightly elevated firing rates, as a function of adapting variance levels, was sufficient to generate the decrease in gain.

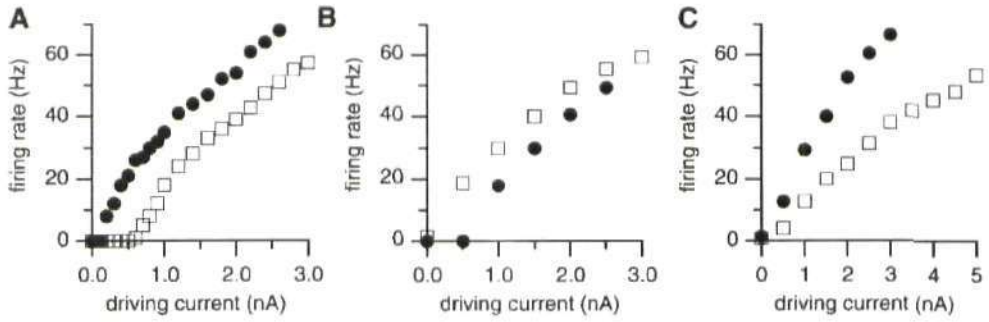

Figure 6.1: Effects of shunting inhibition and noise on neural response
A: Input-output functions for a neuron without (circles), and with (squares) 32 nS of additional constant conductance in the absence of any additional noise from background synaptic input. The input-output function is displaced laterally. **B:** Input-output functions for a different neuron in the 1X condition (circles) and with the same level of conductance but input noise equivalent to the 3X condition (squares). The effect is a relative increase in firing rates for the lowest inputs, resulting in a change in the slope of the firing-rate curve. **C:** The firing-rate function for another neuron under the 1X (circles) and 3X (squares) conditions. The combination of lateral displacement and enhanced firing rates for lower intensity inputs generates a decrease in gain. Taken from Chance et al. (2002)

Adaptation was also modelled by adjusting the relative strength of both afferent inhibition and excitation within a population of recurrently connected inhibitor and excitatory neurons, which provided feedforward input to a read-out cell. The mechanism of gain modulation was achieved by network inhibition becoming progressively stronger, as variance increased; thereby suppressing the overall responsiveness of the cell. Because afferent excitation was also increasing there was an overall decrease in neural gain rather than threshold in the cellular response profile. The change in output was brought about by weakening excitatory depression and thereby increasing the drive acting on inhibition. The model was tested with a steady-state model, in which the output of the cell for a given stimulus was dependent on the levels of depression between the two populations and was, therefore, fixed, as a function of the relative balance of depression between the

two populations of excitatory and inhibitory neurons. The results from this simulation replicated the experimental data well.

The model was also tested under a time-dependent condition, which could accurately model the temporal dynamics of the recurrent synapses. A small network of 4 excitatory and 1 inhibitory I & F cells were recurrently connected to each other and provided feedforward input to a read-out cell. Recurrent excitatory depression was fixed, such that it weakened with every increase in stimulus variance and approached levels of inhibitory depression. The model displayed overall decrease in firing rates; however this was not accompanied by a decrease in gain. Thus the model didn't replicate the experimental data as robustly as the steady-state model.

One key feature of the adapting stimulus is that the presentation frequency of sub-threshold intensity levels increases in line with stimulus-variance. In theory, this could result in an overall release from excitatory depression (as synapses would have more opportunity to recover) and the stimulus conditions necessary to generate a decrease in gain would arise. The model was tested under several different depressing configurations to test whether the inherent diversity within the stimulus would elicit the requisite relief from excitatory depression.

It was only achieved with both facilitation and depression acting on the recurrent excitatory-to-inhibitory synapse and under the asymmetric depression configuration with strongly depressing excitation. Facilitation was modelled to be faster (Varela et al. 1997) than depression (100ms recovery time) and was additive, and acted to boost the initial synaptic response to supra-threshold stimuli. Under the asymmetric synaptic depression configuration, excitatory synapses were relatively suppressed (as compared to the

balanced configuration) and the addition facilitation had an observable effect on the levels of recurrent excitatory drive acing on inhibition. Facilitation resulted in an overall, slight decrease in firing rates in the read-out cell. It has been demonstrated experimentally (Rothman et al. 2009) that the combination of tonic inhibition and synaptic depression elicits multiplicative changes in neural gain; thus the small decrease in overall firing rates due to facilitation of inhibition was enhanced by the multiplicative decrease in neural gain elicited by the slight increase in tonic inhibition (in combination with feedforward depression).

Facilitation is a feature of synaptic activity within the neocortex and can be highly synapse specific (e.g. Thompson 1997; Markram, Wang & Tsodyks 1998, Reyes, Lujan, Rozov, Burnashev, Somogyi & Sakmann 1998) thus it is entirely feasible that gain modulation could arise through such a mechanism. However, the effects were not robust, i.e. were only present for the asymmetric depression configuration and required the presence of tonic inhibition, feedforward depression and spike-rate adaption acting on the read-out cell.

It was decided to use aspects of all the models highlighted above to generate a final model of gain modulation. A network of 4 excitatory and 1 inhibitory I & F neurons were recurrently connected to each other and provided feedforward input to a read-out I & F cell, with tonic inhibition and spike-rate adaption acting on it. Feedforward depression was fixed to have a value of 0.8 and recurrent depression was non-existent, balanced or asymmetric. The non-linear input, as described in the background noise model above, was used to drive the I & F model directly, rather than via a background conductance. The stimulus current was filtered through a bank of parallel non-linear

rectified filters, with thresholds set to be higher than the global mean of the input.

The stimulus under the low variance condition rarely crossed the input-filter threshold thereby resulting in large intervals between the recurrent-network being driven to respond. As variance increased, the presentation of supra-threshold stimuli also increased, therefore the network responded more frequently and overall synaptic activity was increasingly suppressed. This led to a decrease in neural-gain as a function of stimulus-variance, in both excitatory and inhibitory feedforward signals. Gain modulation was enhanced by both tonic inhibition and spike-rate adaption, but did not dependent on it entirely.

Under conditions of adaptation to global-mean, the read-out cell displayed a lateral displacement that was dependent on global-mean levels; the step-size and normalised deviation-from –mean functions also displayed invariance and thus matched the experimental data.

The advantage of the background firing rate model of gain modulation (Chance et al. 2002) is that it is relatively simple, i.e. there is no direct reliance on synaptic dynamics to generate gain modulation. However it does require that background firing rates are correlated somehow with the stimulus and must also be derived from a non-adapting source. This leads to the suggestion that if this is the mechanism by which gain modulation is generated, the strength of adaptation, in a given cell, would be strongly dependent on levels of adaptation in surrounding neurons, assuming of course that much of the modulatory drive acting on a cell is local in origin. One possible suggestion is that, if adaptation to stimulus statistics in the inferior colliculus (Dean et al. 2005) is representative of the thalamic response both across stimulus modalities and rodent

species, modulatory sources are derived from neurons that inherit their response profiles from their thalamic afferents. An even simpler proposition, however, is that barrel cortex neurons receive their primary source of afferent input directly from adapted thalamic cells and gain modulation emerges through the interaction between the non-linear input and synaptic depression. This would allow for adaptation to simple stimulus features such as noise and intensity levels to emerge from the network without the need for extensive computations.

6.2.4: Functional significance

In terms of the functional significance of these results, it should be noted that intensity is not the most important feature of the stimulus, for example barrel neurons are highly sensitive to velocity of whisker deflection (e.g. Maravall et al. 2007). However, as with contrast-gain in the visual system, amplitude fluctuations about the mean are such a universal feature of sensory input, across all modalities, that encoding this information represents a fundamental feature of neural computation. By adjusting the relationship between firing rate and stimulus in response to global and local levels of intensity fluctuations, this allows for other, higher level features of the stimulus space to be encoded by other more complex mechanism, such as spike-timing for example. The argument here is that amplitude of whisker deflection is analogous to contrast in the visual system, therefore understanding its functioning and how it is encoded can only add to understanding about brain function in general. Visual contrast is known to have profound effects on responsiveness of simple cells within V1 (e.g. on spatial summation, Sceniak, Ringach, Hawken & Shapley 1999), therefore it is possible that whisker deflection amplitude may have similar effects on responsiveness to higher stimulus

features, within the barrel cortex.

Rats constantly use their whiskers, to convey information about texture within the environment. The analysis of the experimental results, reveal that during conditions of relatively low noise, barrel cortex neurons are primed to detect deviant signals. It is useful to consider what is most efficient in the low-noise case. There is clearly an advantage in being primed to signal rapid changes in input when the environment is rarely changing. Information theory was of course developed to account for communication systems, which are different from neural systems in one important facet - they have been designed! Nervous systems have evolved and thus, in engineering terms, are not perfect. Clearly there is evidence that encoding of contrast in V1 is in line with Information theory, therefore one can assume that it is advantageous to have contrast encoded in a manner that maximises information with the minimum number of spikes. Within barrel cortex the situation is somewhat different: whisker deflection amplitude can not cover the same order of magnitude that visual contrast does. Therefore, it may well be advantageous to have a trade off between efficiency and being ready to respond to large, novel stimuli. Clearly, in noisy environments the encoding strategy switches to one in which the distribution is encoded along information theoretical principles.

Finally, a recent study (von Heimendahl, Itskov, Arabzadeh & Diamond 2007) looking at the correlation between barrel cortex firing rates and texture identification, in rats trained in a texture discrimination task, found that the spike trains in the 75ms before a task was performed contained the most information about the stimulus. It was concluded that the animal's perceptual judgment of texture is highly dependent on firing rates within this time window. This suggests that the invariance of the step-size function

is a necessary component of texture-discrimination.

6.2.5: Future work

In terms of future work, there are a few questions that remain unanswered from this analysis of the data. For example, the time course of adaptation to the global statistics of the stimulus could not be assessed with any certainty from the data provided. An optimal experimental design to analyse the time course of adaptation would be a switching stimulus, such as that performed by (Dean et al. 2008) in the inferior colliculus (IC). It was found that adaptation to global mean in the IC was fairly rapid (<300ms); no experiments were performed exploring the time course of variance. It would be interesting to see whether a) there are any differences in the time course of adaptation between the thalamus and the barrel cortex and b) to see if the time course of adaptation to variance and mean are distinct from each other. If adaptation is predominantly 'inherited' from sub-cortical afferent input, one would expect that the development of the cortical adaptive response would occur within a time window only slightly longer (to allow for signal transmission) than that observed in the thalamus. This would lend weight to the model of adaptation that combined non-linear afferent input and synaptic dynamics; it would be expected that adaptation, within this model, would be reliant on transmission latency only, as it emerges without the need for complex processing or net changes in the balance between inhibition and excitation to develop.

It would also be useful to compare the adaptive response of the thalamic region that provides afferent input to the barrel cortex, with the experimental data presented here and in Garcia-Lazaro et al. (2007). The adaptive responses in the inferior colliculus (IC) were distinct from the barrel cortex data in 3 key areas:

1. that adaptation to variance was relatively uncommon in the IC,
2. that neural gain increased with stimulus variance and
3. that the step-size functions following adaptation to stimulus mean were not invariant with respect to adapting global mean levels.

These key features could be unique to the auditory system or, as the IC is subcortical, they could represent a transformation of adaptive strategies between thalamus and cortex.

In terms of the increase in gain observed within the IC, in response to increased stimulus variance, there is a possibility that the analysis of the data may have been subject to a similar error as was found for the experimental data presented within Chapter 3. In which case, adaptation to variance, within the IC, would result in a decrease in gain, as per the expectations of the Efficient Coding Hypothesis.

As Barlow himself noted, the Efficient Coding Hypothesis is focused on how information is packaged for processing of sensory information. If there is a difference in thalamus/cortical adapting strategies it could represent the fact that primary sensory areas are not simply relay stations but active processing areas, which can influence perception. When Barlow initially worked on his hypothesis the prevailing view of neural processing was hierarchical; for example simple visual features such as orientation, spatial frequency etc are detected in V1 then more complex processing is performed in V2 and further up the cortical hierarchy. This view is now falling out favour with the greater understanding of the role and function of feedback and feedforward connectivity in V1 (e.g. Angelucci & Bullier 2003; Angelucci et al. 2002 & Bullier 2001). Any differences between thalamic and cortical adaptation to stimulus mean could be an indication that thalamic

areas are more focused on transmitting sensory input, whilst the barrel cortex is performing more complex computations on the environmental data.

6.3: Conclusion

In conclusion, the main findings of this thesis are that:

1. Adaptation to variance results in a decrease in neural gain, in line with the principles of the Efficient Coding Hypothesis, as proposed by Horace Barlow.
2. The lateral displacement observed following adaptation to global mean serves to maintain the relationship between firing rate and local changes in the stimulus.
3. The lateral displacement and invariance of the step-size and normalised deviation-from-mean functions were replicated by introducing spike-rate adaptation and a tonic inhibition (whose strength was dependent on the global mean stimulating amplitude) into a model, conductance based, I & F, cell.
4. Gain modulation was modelled by increasing, in proportion to stimulus variance, the firing rates of Poisson-driven excitatory and inhibitory conductances in a balanced manner, thus confirming that the *in vitro* model of gain modulation though background noise (Chance et al. 2002) can be applied to *in vivo* data. This was also achieved by having background firing rates derived from a bank of non-linear filters of the stimulus (to approximate background noise arising from sources already adapted to global-mean); overall background firing rates increased, with stimulus variance, sufficiently to generate gain modulation, but only in the presence of spike-rate adaptation and tonic inhibition.
5. There is evidence to suggest that it is the changing relationship between afferent excitation and inhibition that generates changes in neural gain; this was

simulated in a model of asymmetric synaptic depression (Chelaru & Dragoi, 2008) but less successfully. The model best approximated the experimental data by incorporating both facilitation and depression onto recurrent excitatory-to-inhibitory synapses, when recurrent excitatory synapses were more strongly depressing than inhibitory synapses.

6. Gain modulation was elicited by a combination of the models above, whereby non-linear afferent inputs (derived from a bank of non-linear, rectified filters of the stimulus whose threshold was set to be slightly higher than the global mean amplitude of the stimulus) were used to drive the recurrent population of excitatory and inhibitory I & F cells. The combination of recurrent depression acting within the network and feedforward depression, spike-rate adaptation and tonic inhibition acting on the read-out I & F cell elicited a decrease in neural gain as global variance increased, under various depression configurations. The decrease in gain was diminished but not abolished in the absence of tonic inhibition and spike-rate adaptation.

Adaptation to amplitude of whisker deflection is clearly an important part of processing of sensory input. While it is not the most relevant feature of the stimulus, the results presented here demonstrate that neurons change their coding strategies depending upon the overall levels of mean amplitude and variance in the sensory input. The argument presented here, based on reanalysis of experimental data (Garcia-Lazaro et al. 2007), is that under low-noise conditions, neurons act as deviation detectors, primed to respond to large and rapid changes in the stimulus; however, under conditions of

increased noise their main function is to compute the full range of the stimulus distribution through an adjustment in neural gain.

List of references.

Abbott, L. F. & Dayan, P. (1999), 'The effect of correlated variability on the accuracy of a population code.', *Neural Computation*, 11, 91-101.

Abbott, L. F., Varela, J. A., Sen, K. & Nelson, S. B. (1997), 'Synaptic depression and cortical gain control.', *Science* 275(5297), 220–224.

Adorján, P., Piepenbrock, C. & Obermayer, K. (1999), 'Contrast adaptation and infomax in visual cortical neurons.', *Rev Neurosci* 10(3-4), 181–200.

Albrecht, D. G., Farrar, S. B. & Hamilton, D. B. (1984), 'Spatial contrast adaptation characteristics of neurones recorded in the cat's visual cortex.', *J Physiol* 347, 713–739.

Albrecht, D. G. & Geisler, W. S. (1991), 'Motion selectivity and the contrast-response function of simple cells in the visual cortex.', *Vis Neurosci* 7(6), 531–546.

Albrecht, D. & Geisler, W. (1994,), *Visual cortex neurons in monkey and cat: Contrast response nonlinearities and stimulus selectivity. Computational vision based on neurobiology,*, Bellingham, Wash: SPIE.

Albrecht, D. G., Geisler, W. S., Frazor, R. A. & Crane, A. M. (2002), 'Visual cortex neurons of monkeys and cats: temporal dynamics of the contrast response function.', *J Neurophysiol* 88(2), 888–913.

Albrecht, D. G. & Hamilton, D. B. (1982), 'Striate cortex of monkey and cat: contrast response function.', *J Neurophysiol* 48(1), 217–237.

Andersen, P., Sundberg, S. H., Sveen, O., Swann, J. W. & Wigström, H. (1980), 'Possible mechanisms for long-lasting potentiation of synaptic transmission in hippocampal slices from guinea-pigs.', *J Physiol* 302, 463–482.

Anderson, J. S., Carandini, M. & Ferster, D. (2000), 'Orientation tuning of input conductance, excitation, and inhibition in cat primary visual cortex.', *J Neurophysiol* 84(2), 909–926.

Angelucci, A. & Bullier, J. (2003), 'Reaching beyond the classical receptive field of v1 neurons: horizontal or feedback axons?', *J Physiol* Paris 97(2-3), 141–154.
URL: http://dx.doi.org/10.1016/j.jphysparis.2003.09.001

Angelucci, A., Levitt, J. B., Walton, E. J. S., Hupe, J.-M., Bullier, J. & Lund, J. S. (2002), 'Circuits for local and global signal integration in primary visual cortex.', *J Neurosci* 22(19), 8633–8646.

Armstrong-James, M. & Fox, K. (1987), 'Spatiotemporal convergence and divergence in the rat s1 "barrel" cortex.', *J Comp Neurol* 263(2), 265–281.

URL: http://dx.doi.org/10.1002/cne.902630209

Attneave, F. (1954), 'Some informational aspects of visual perception.', *Psychol Rev* 61(3), 183–193.

Baccus, S. A. & Meister, M. (2002), 'Fast and slow contrast adaptation in retinal circuitry.', *Neuron* 36(5), 909–919.

Baddeley, R. Abbott, L. F. Booth, M. C. A. Sengpiel, F. Freeman, T. Wakeman, E. A. & Rolls, E .T. (1997), 'Responses of neurons in primary and inferior temporal visual cortices to natural scenes', *Proc R Soc Lond B*, 264,1775–1783

Baden, T. & Hedwig, B. (2007), 'Neurite-specific ca2+ dynamics underlying sound processing in an auditory interneurone.', *Dev Neurobiol* 67(1), 68–80.
URL: http://dx.doi.org/10.1002/dneu.20323

Barlow, H. (1959), Sensory mechanisms, the reduction of redundancy, and intelligence., *in* 'NPL Symposium on the Mechanization of Thought Process', Vol. 10, H.M. Stationary Office, London, pp. 535–539.

Barlow, H. (1961), Possible principles underlying the transformations of sensory messages, *in* W.Rosenblith, ed., 'Sensory Communication', M.I.T. Press,, chapter 13, pp. 217–234.

Barlow, H. (2001), 'Redundancy reduction revisited.', *Network* 12(3), 241–253.

Barlow, H. B. (1972), 'Single units and sensation: a neuron doctrine for perceptual psychology?', *Perception* 1(4), 371–394.

Bi, G. Q. & Poo, M. M. (1998), 'Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type.', *J Neurosci* 18(24), 10464–10472.

Bliss, T. V. & Lomo, T. (1973), 'Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path.', *J Physiol* 232(2), 331–356.

Boudreau, C. E. & Ferster, D. (2005), 'Short-term depression in thalamocortical synapses of cat primary visual cortex.', *J Neurosci* 25(31), 7179–7190.
URL: http://dx.doi.org/10.1523/JNEUROSCI.1445-05.2005

Brady, N. and Field, D. J. (2000), 'Local contrast in natural images: normalisation and coding efficiency'. *Perception* , 29(9),1041–1055.

Brenner, N., Strong, S. P., Koberle, R., Bialek, W. & de Ruyter van Steveninck, R. R. (2000), 'Synergy in a neural code.', *Neural Comput* 12(7), 1531–1552.

Brunel, N. & Nadal, J.P. (1998), 'Mutual information, Fisher information, and population coding.', *Neural Comp.*,10,1731–1757.

Bullier, J. (2001), 'Integrated model of visual processing.', *Brain Res Brain Res Rev* 36(2-3), 96–107.

Butts, D.A. & Goldman, M. (2006), 'Tuning curves, neuronal variability and sensory coding.', *PLOS Biol.*, 4:639–646.

Cano, M., Bezdudnaya, T., Swadlow, H. A. & Alonso, J.-M. (2006), 'Brain state and contrast sensitivity in the awake visual thalamus.', *Nat Neurosci* 9(10), 1240–1242.
URL: http://dx.doi.org/10.1038/nn1760

Carandini, M. & Ferster, D. (1997), 'A tonic hyperpolarization underlying contrast adaptation in cat visual cortex.', *Science* 276(5314), 949–952.

Carandini, M. & Heeger, D. J. (1994), 'Summation and division by neurons in primate visual cortex.', *Science* 264(5163), 1333–1336.

Carandini, M., Heeger, D. J. & Movshon, J. A. (1997), 'Linearity and normalization in simple cells of the macaque primary visual cortex.', *J Neurosci* 17(21), 8621–8644.

Carandini, M., Heeger, D. J. & Senn, W. (2002), 'A synaptic explanation of suppression in visual cortex.', *J Neurosci* 22(22), 10053–10065.

Carandini, M. & Ringach, D. L. (1997), 'Predictions of a recurrent model of orientation selectivity.', *Vision Res* 37(21), 3061–3071.

Carvalho, T. P. & Buonomano, D. V. (2009), 'Differential effects of excitatory and inhibitory plasticity on synaptically driven neuronal input-output functions.', *Neuron* 61(5), 774–785.
URL: http://dx.doi.org/10.1016/j.neuron.2009.01.013

Castro-Alamancos, M. A. (2004a), 'Absence of rapid sensory adaptation in neocortex during information processing states.', *Neuron* 41(3), 455–464.

Castro-Alamancos, M. A. (2004b), 'Dynamics of sensory thalamocortical synaptic networks during information processing states.', *Prog Neurobiol* 74(4), 213–247.
URL: http://dx.doi.org/10.1016/j.pneurobio.2004.09.002

Castro-Alamancos, M. A. & Oldford, E. (2002), 'Cortical sensory suppression during arousal is due to the activity-dependent depression of thalamocortical synapses.', *J Physiol* 541(Pt 1), 319–331.

Challis, E., Yarrow, S. & Series, P. (2008), 'Fisher vs Shannon information in

populations of neurons', *NeuroComp 2008*

Chance, F. S., Abbott, L. F. & Reyes, A. D. (2002), 'Gain modulation from background synaptic input.', *Neuron* 35(4), 773–782.

Chance, F. S., Nelson, S. B. & Abbott, L. F. (1998), 'Synaptic depression and the temporal response characteristics of v1 cells.', *J Neurosci* 18(12), 4785–4799.

Chander, D. & Chichilnisky, E. J. (2001), 'Adaptation to temporal contrast in primate and salamander retina.', *J Neurosci* 21(24), 9904–9916.

Chelaru, M. I. & Dragoi, V. (2008), 'Asymmetric synaptic depression in cortical networks.', *Cereb Cortex* 18(4), 771–788.
URL: http://dx.doi.org/10.1093/cercor/bhm119

Chung, S., Li, X. & Nelson, S. B. (2002), 'Short-term depression at thalamocortical synapses contributes to rapid adaptation of cortical sensory responses in vivo.', *Neuron* 34(3), 437–446.

Clarke, B. S. & Barron, A. R. (1990), 'Information-theoretic asymptotics of Bayes methods.', *IEEE Trans.Information Theory*, 36, 453–471.

Condon, C. D. & Weinberger, N. M. (1991), 'Habituation produces frequency-specific plasticity of receptive fields in the auditory cortex.', *Behav Neurosci* 105(3), 416– 430.

Connors, B. W. & Gutnick, M. J. (1990), 'Intrinsic firing patterns of diverse neocortical neurons.', *Trends Neurosci* 13(3), 99–104.

Connors, B. W. & Kriegstein, A. R. (1986), 'Cellular physiology of the turtle visual cortex: distinctive properties of pyramidal and stellate neurons.', *J Neurosci* 6(1), 164–177.

Contreras, D. & Palmer, L. (2003), 'Response to contrast of electrophysiologically defined cell classes in primary visual cortex.', *J Neurosci* 23(17), 6936–6945.

Cover, T. M. & Thomas, J. A. (1991), *Elements of Information Theory.*, Wiley, New York.

Craik, K. J. W. (1948), 'Theory of the human operator in control systems; man as an element in a control system.', *Br J Psychol Gen Sect* 38(Pt 3), 142–148.

Crowder, N. A., Price, N. S. C., Hietanen, M. A., Dreher, B., Clifford, C. W. G. & Ibbotson, M. R. (2006), 'Relationship between contrast adaptation and orientation tuning in v1 and v2 of cat visual cortex.', *J Neurophysiol* 95(1), 271–283.
URL: http://dx.doi.org/10.1152/jn.00871.2005

de Boer, R. & Kuyper, P. (1968), 'Triggered correlation.', *IEEE Trans Biomed Eng*

15(3), 169–179.

de Kock, C. P. J. & Sakmann, B. (2009), 'Spiking in primary somatosensory cortex during natural whisking in awake head-restrained rats is cell-type specific.', *Proc Natl Acad Sci* U S A 106(38), 16446–16450.
URL: http://dx.doi.org/10.1073/pnas.0904143106

Dean, I., Harper, N. S. & McAlpine, D. (2005), 'Neural population coding of sound level adapts to stimulus statistics.', *Nat Neurosci* 8(12), 1684–1689.
URL: http://dx.doi.org/10.1038/nn1541

Dean, I., Robinson, B. L., Harper, N. S. & McAlpine, D. (2008), 'Rapid neural adaptation to sound level statistics.', *J Neurosci* 28(25), 6430–6438.
URL: http://dx.doi.org/10.1523/JNEUROSCI.0470-08.2008

Destexhe, A. (1997), 'Conductance-based integrate-and-fire models.', *Neural Comput* 9(3), 503–514.

Dragoi, V., Rivadulla, C. & Sur, M. (2001), 'Foci of orientation plasticity in visual cortex.', *Nature* 411(6833), 80–86.
URL: http://dx.doi.org/10.1038/35075070

Dragoi, V., Sharma, J. & Sur, M. (2000), 'Adaptation-induced plasticity of orientation tuning in adult visual cortex.', *Neuron* 28(1), 287–298.

Dragoi, V., Turcu, C. M. & Sur, M. (2001), 'Stability of cortical responses and the statistics of natural scenes.', *Neuron* 32(6), 1181–1192.

Duong, T. & Freeman, R. D. (2007), 'Spatial frequency-specific contrast adaptation originates in the primary visual cortex.', *J Neurophysiol* 98(1), 187–195.
URL: http://dx.doi.org/10.1152/jn.01364.2006

Dayan, P. & Abbott, L. F. (2002), '*Theoretical Neuroscience*.', MIT press, Cambridge, MA.

Ermentrout, B., Pascal, M. & Gutkin, B. (2001), 'The effects of spike frequency adaptation and negative feedback on the synchronization of neural oscillators.', *Neural Comput* 13(6), 1285–1310.

Fairhall, A. L., Lewen, G. D., Bialek, W. & de Ruyter Van Steveninck, R. R. (2001), 'Efficiency and ambiguity in an adaptive neural code.', *Nature* 412(6849), 787–792.
URL: http://dx.doi.org/10.1038/35090500

Feldman, D. E. (2000), 'Timing-based ltp and ltd at vertical inputs to layer ii/iii pyramidal cells in rat barrel cortex.', *Neuron* 27(1), 45–56.

Feldmeyer, D., Egger, V., Lubke, J. & Sakmann, B. (1999), 'Reliable synaptic connections between pairs of excitatory layer 4 neurones within a single 'barrel' of developing rat somatosensory cortex.', *J Physiol* 521 Pt 1, 169–190.

Feldmeyer, D., Lübke, J. & Sakmann, B. (2006), 'Efficacy and connectivity of intra-columnar pairs of layer 2/3 pyramidal cells in the barrel cortex of juvenile rats.', *J Physiol* 575(Pt 2), 583–602.
URL: http://dx.doi.org/10.1113/jphysiol.2006.105106

Field, D. (1994), 'What is the goal of sensory coding', Neural Computation, 6, 559–601.

Fisher, R. (1922), 'On the mathematical foundations of theoretical statistics', *Philosophical Transactions of the Royal Society*, A 222, 309–368.

Freeman, T. C. B., Durand, S., Kiper, D. C. & Carandini, M. (2002), 'Suppression without inhibition in visual cortex.', *Neuron* 35(4), 759–771.

Galarreta, M. & Hestrin, S. (1998), 'Frequency-dependent synaptic depression and the balance of excitation and inhibition in the neocortex.', *Nat Neurosci* 1(7), 587–594.
URL: http://dx.doi.org/10.1038/2882

Garcia-Lazaro, J. A., Ho, S. S. M., Nair, A. & Schnupp, J. W. H. (2007), 'Shifting and scaling adaptation to dynamic stimuli in somatosensory cortex.', *Eur J Neurosci* 26(8), 2359–2368.
URL: http://dx.doi.org/10.1111/j.1460-9568.2007.05847.x

Gershon, E. D., Wiener, M. C., Latham, P. E. & Richmond, B. J. (1998), 'Coding Strategies in Monkey V1 and Inferior Temporal Cortices' *J. Neurophysiol.* 79: 1135-1144.

Gil, Z., Connors, B. W. & Amitai, Y. (1999), 'Efficacy of thalamocortical and intracortical synaptic connections: quanta, innervation, and reliability.', *Neuron* 23(2), 385– 397.

Gilbert, C. D., Li, W. & Piech, V. (2009), 'Perceptual learning and adult cortical plasticity.', *J Physiol* 587(Pt 12), 2743–2751.
URL: http://dx.doi.org/10.1113/jphysiol.2009.171488

Gonzalez, R. & Wintz, P. (1977), *Digital image processing*, LONDON, UK, ADDISON-WESLEY, UK, ISBN 0-201-02596-5.

Hebb, D. (1949), *The organization of behavior*,, Wiley, New York.

Heeger, D. J. (1992), 'Normalization of cell responses in cat striate cortex.', *Vis Neurosci* 9(2), 181–197.

Heiss, J. E., Katz, Y., Ganmor, E. & Lampl, I. (2008), 'Shift in the balance between excitation and inhibition during sensory adaptation of s1 neurons.', *J Neurosci* 28(49), 13320–13330.
URL: http://dx.doi.org/10.1523/JNEUROSCI.2646-08.2008

Higley, M. J. & Contreras, D. (2006), 'Balanced excitation and inhibition determine spike timing during frequency adaptation.', *J Neurosci* 26(2), 448–457.
URL: http://dx.doi.org/10.1523/JNEUROSCI.3506-05.2006

Hines, M. L. & Carnevale, N. T. (1997), 'The neuron simulation environment.', *Neural Comput* 9(6), 1179–1209.

Hodgkin, A. L. & Huxley, A. F. (1990), 'A quantitative description of membrane current and its application to conduction and excitation in nerve. 1952.', *Bull Math Biol* 52(1-2), 25–71; discussion 5–23.

Holt, G. R. & Koch, C. (1997), 'Shunting inhibition does not have a divisive effect on firing rates.', *Neural Comput* 9(5), 1001–1013.

Jacob, V., Brasier, D. J., Erchova, I., Feldman, D. & Shulz, D. E. (2007), 'Spike timing-dependent synaptic depression in the in vivo barrel cortex of the rat.', *J Neurosci* 27(6), 1271–1284.
URL: http://dx.doi.org/10.1523/JNEUROSCI.4264-06.2007

Jin, D. Z., Dragoi, V., Sur, M. & Seung, H. S. (2005), 'Tilt aftereffect and adaptation-induced changes in orientation tuning in visual cortex.', *J Neurophysiol* 94(6), 4038–4050.
URL: http://dx.doi.org/10.1152/jn.00571.2004

Kang, K. & Sompolinsky, H. (2001), 'Mutual information of population codes and distance measures in probability space.', *Phys Rev Lett*, 86(21):4958–4961

Kayser, A., Priebe, N. J. & Miller, K. D. (2001), 'Contrast-dependent nonlinearities arise locally in a model of contrast-invariant orientation tuning.', *J Neurophysiol* 85(5), 2130–2149.

Kim, K. J. & Rieke, F. (2001), 'Temporal contrast adaptation in the input and output signals of salamander retinal ganglion cells.', *J Neurosci* 21(1), 287–299.

Krook-Magnuson, E. I. & Huntsman, M. M. (2005), 'Excitability of cortical neurons depends upon a powerful tonic conductance in inhibitory networks.', *Thalamus Relat Syst* 3(2), 115–120.
URL: http://dx.doi.org/10.1017/S1472928807000192

Krook-Magnuson, E. I., Li, P., Paluszkiewicz, S. M. & Huntsman, M. M. (2008), Tonically active inhibition selectively controls feedforward circuits in mouse barrel cortex.', *J Neurophysiol* 100(2), 932–944.
URL: http://dx.doi.org/10.1152/jn.01360.2007

Kyriazi, H., Carvell, G. E., Brumberg, J. C. & Simons, D. J. (1998), 'Laminar differences in bicuculline methiodide's effects on cortical neurons in the rat whisker/barrel system.', *Somatosens Mot Res* 15(2), 146–156.

Kyriazi, H. T., Carvell, G. E., Brumberg, J. C. & Simons, D. J. (1996), 'Quantitative effects of gaba and bicuculline methiodide on receptive field properties of neurons in real and simulated whisker barrels.', *J Neurophysiol* 75(2), 547–560.

Lamsa, K., Heeroma, J. H. & Kullmann, D. M. (2005), 'Hebbian ltp in feed-forward inhibitory interneurons and the temporal fidelity of input discrimination.', *Nat Neurosci* 8(7), 916–924.
URL: http://dx.doi.org/10.1038/nn1486

Laughlin, S. (1981), 'A simple coding procedure enhances a neuron's information capacity.', *Z Naturforsch* C 36(9-10), 910–912.

Laughlin, S. B., Howard, J. & Blakeslee, B. (1987), 'Synaptic limitations to contrast coding in the retina of the blowfly calliphora.', *Proc R Soc Lond B Biol Sci* 231(1265), 437–467.

Laughlin, S. & Hardie, R. (1978), 'Common strategies for light adaptation in the peripheral visual systems of fly and dragonfly', *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology* 128, 319–340.

Lewicki, M. S. (2002), 'Efficient coding of natural sounds.', *Nat Neurosci* 5(4), 356–363.
URL: http://dx.doi.org/10.1038/nn831

Li, B., Thompson, J. K., Duong, T., Peterson, M. R. & Freeman, R. D. (2006), 'Origins of cross-orientation suppression in the visual cortex.', *J Neurophysiol* 96(4), 1755–1764.
URL: http://dx.doi.org/10.1152/jn.00425.2006

Li, W., Piëch, V. & Gilbert, C. D. (2004), 'Perceptual learning and top-down influences in primary visual cortex.', *Nat Neurosci* 7(6), 651–657.
URL: http://dx.doi.org/10.1038/nn1255

Linsker, R. (1990), 'Perceptual neural organization: some approaches based on network models and information theory.', *Annu Rev Neurosci* 13, 257–281.
URL: http://dx.doi.org/10.1146/annurev.ne.13.030190.001353

Liu, Y. H. & Wang, X. J. (2001), 'Spike-frequency adaptation of a generalized leaky

integrate-and-fire model neuron.', *J Comput Neurosci* 10(1), 25–45.

Mackay, D. and McCullogh, W. S. (1952), 'The limiting information capacity of neuronal link.' *Bull. Math. Biophys.* 14,127–135, 1952.

Madison, D. V. & Nicoll, R. A. (1984), 'Control of the repetitive discharge of rat ca 1 pyramidal neurones in vitro.', *J Physiol* 354, 319–331.

Maffei, L. & Fiorentini, A. (1973), 'The visual cortex as a spatial frequency analyser.', *Vision Res* 13(7), 1255–1267.

Mainen, Z. F. & Sejnowski, T. J. (1996), 'Influence of dendritic structure on firing pattern in model neocortical neurons.', *Nature* 382(6589), 363–366.
URL: http://dx.doi.org/10.1038/382363a0

Malone, B. J. & Semple, M. N. (2001), 'Effects of auditory stimulus context on the representation of frequency in the gerbil inferior colliculus.', *J Neurophysiol* 86(3), 1113–1130.

Maravall, M., Petersen, R. S., Fairhall, A. L., Arabzadeh, E. & Diamond, M. E. (2007), 'Shifts in coding properties and maintenance of information transmission during adaptation in barrel cortex.', *PLoS Biol* 5(2), e19.
URL: http://dx.doi.org/10.1371/journal.pbio.0050019

Marder, C. P. & Buonomano, D. V. (2004), 'Timing and balance of inhibition enhance the effect of long-term potentiation on cell firing.', *J Neurosci* 24(40), 8873–8884. **URL**: http://dx.doi.org/10.1523/JNEUROSCI.2661-04.2004

Markram, H., Lübke, J., Frotscher, M. & Sakmann, B. (1997), 'Regulation of synaptic efficacy by coincidence of postsynaptic aps and epsps.', *Science* 275(5297), 213– 215.

Markram, H. & Tsodyks, M. (1996), 'Redistribution of synaptic efficacy between neocortical pyramidal neurons.', *Nature* 382(6594), 807–810.
URL: http://dx.doi.org/10.1038/382807a0

Markram, H., Wang, Y & Tsodyks, M. (1998), 'Differential signaling via the same axon of neocortical pyramidal neurons.', *Proc Natl Acad Sci USA*, 95:5323-5328.

Martinez, L. M., Wang, Q., Reid, R. C., Pillai, C., Alonso, J.-M., Sommer, F. T. & Hirsch, J. A. (2005), 'Receptive field structure varies with layer in the primary visual cortex.', *Nat Neurosci* 8(3), 372–379.
URL: http://dx.doi.org/10.1038/nn1404

McCormick, D. A., Connors, B. W., Lighthall, J. W. & Prince, D. A. (1985), 'Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex.', *J Neurophysiol* 54(4), 782–806.

McLean, J. & Palmer, L. A. (1996), 'Contrast adaptation and excitatory amino acid receptors in cat striate cortex.', *Vis Neurosci* 13(6), 1069–1087.

Monier, C., Chavane, F., Baudot, P., Graham, L. J. & Frégnac, Y. (2003), 'Orientation and direction selectivity of synaptic inputs in visual cortical neurons: a diversity of combinations produces spike tuning.', *Neuron* 37(4), 663–680.

Movshon, J. A. & Lennie, P. (1979), 'Pattern-selective adaptation in visual cortical neurones.', *Nature* 278(5707), 850–852.

Murphy, B. K. & Miller, K. D. (2003), 'Multiplicative gain changes are induced by excitation or inhibition alone.', *J Neurosci* 23(31), 10040–10051.

Ohzawa, I., Sclar, G. & Freeman, R. D. (1982), 'Contrast gain control in the cat visual cortex.', *Nature* 298(5871), 266–268.

Ohzawa, I., Sclar, G. & Freeman, R. D. (1985), 'Contrast gain control in the cat's visual system.', *J Neurophysiol* 54(3), 651–667.

Okajima, K. (2001), 'An infomax-based learning rule that generates cells similar to visual cortical neurons.', *Neural Netw* 14(9), 1173–1180.

Olshausen, B. A. & Field, D. J. (2004), 'Sparse coding of sensory inputs.', Curr Opin Neurobiol 14(4), 481–487.
URL: http://dx.doi.org/10.1016/j.conb.2004.07.007

Petersen, C. C. H. (2002), 'Short-term dynamics of synaptic transmission within the excitatory neuronal network of rat layer 4 barrel cortex.', *J Neurophysiol* 87(6), 2904–2914.

Priebe, N. J. & Ferster, D. (2006), 'Mechanisms underlying cross-orientation suppression in cat visual cortex.', *Nat Neurosci* 9(4), 552–561.
URL: http://dx.doi.org/10.1038/nn1660

Reig, R. Gallego, R. Nowak, L. G. & Sanchez-Vives, M. V. (2006), 'Impact of cortical network activity on short-term synaptic depression.', *Cereb Cortex* , 16(5):688–695.

Reyes, A., Lujan, R., Rozov, A., Burnashev, .N, Somogyi, P. & Sakmann, B. (1998), 'Target-cell-specific facilitation and depression in neocortical circuits.', *Nat Neurosci,* 1:279-285.

Rieke, F. (2001), 'Temporal Contrast Adaptation in Salamander Bipolar Cells', *J. Neurosci* 21, 9445 - 9454.

Rieke, F. Warland, D. de Ruyter van Steveninck, R. & Bialek, W. (1996)

'*Spikes: Exploring the neural code*' , MIT Press, Cambridge

Ringach, D. L. & Malone, B. J. (2007), 'The operating point of the cortex: neurons as large deviation detectors.', *J Neurosci* 27(29), 7673–7683.
URL: http://dx.doi.org/10.1523/JNEUROSCI.1048-07.2007

Rissanen, J. (1996) 'Fisher Information and Stochastic Complexity.', *IEEE Trans. Information Theory*, IT-42(1), 40–47

Rothman, J. S., Cathala, L., Steuber, V. & Silver, R. A. (2009), 'Synaptic depression enables neuronal gain control.', *Nature* 457(7232), 1015–1018.
URL: http://dx.doi.org/10.1038/nature07604

Sadakane, O., Ozeki, H., Naito, T., Akasaki, T., Kasamatsu, T. & Sato, H. (2006), 'Contrast-dependent, contextual response modulation in primary visual cortex and lateral geniculate nucleus of the cat.', *Eur J Neurosci* 23(6), 1633–1642.
URL: http://dx.doi.org/10.1111/j.1460-9568.2006.04681.x

Salin, P. A. & Prince, D. A. (1996), 'Spontaneous gabaa receptor-mediated inhibitory currents in adult rat somatosensory cortex.', *J Neurophysiol* 75(4), 1573–1588.

Sanchez-Vives, M. V., Nowak, L. G. & McCormick, D. A. (2000), 'Cellular mechanisms of long-lasting adaptation in visual cortical neurons in vitro.', *J Neurosci* 20(11), 4286–4299.

Sato, T. R., Gray, N. W., Mainen, Z. F. & Svoboda, K. (2007), 'The functional microarchitecture of the mouse barrel cortex.', *PLoS Biol* 5(7), e189.
URL: http://dx.doi.org/10.1371/journal.pbio.0050189

Sceniak, M. P., Ringach, D. L., Hawken, M. J. & Shapley, R. (1999), 'Contrast's effect on spatial summation by macaque v1 neurons.', *Nat Neurosci* 2(8), 733–739.
URL: http://dx.doi.org/10.1038/11197

Schwabe, L. & Obermayer, K. (2005), 'Adaptivity of tuning functions in a generic recurrent network model of a cortical hypercolumn.', *J Neurosci* 25(13), 3323–3332.
URL: http://dx.doi.org/10.1523/JNEUROSCI.4493-04.2005

Sclar, G., Lennie, P. & DePriest, D. D. (1989), 'Contrast adaptation in striate cortex of macaque.', *Vision Res* 29(7), 747–755.

Sclar, G., Maunsell, J. H. & Lennie, P. (1990), 'Coding of image contrast in central visual pathways of the macaque monkey.', *Vision Res* 30(1), 1–10.

Series, P., Latham, P.E. & Pouget, A. (2004), 'Tuning curve sharpening for orientation selectivity: coding efficiency and the impact of correlations.', *Nature Neurosci* 7:1129-1135

Seung, H.S. & Sompolinsky, H. (1993), 'Simple models for reading neuronal population codes.', *Proc Natl Acad Sci USA*, 90:10749-10753.

Shannon, C. (1948), '*A mathematical theory of communication*', Bell System Technical Journal 27, 379–423 & 623–656.

Shannon, C. & Weaver, W. (1949), *The Mathematical Theory of Communication*, Univ of Illinois Press,.

Shapley, R., P.-V. (1986), 'Cat and monkey retinal ganglion cells and their visual functional roles', *Trends in Neurosciences* 9, 229–235.

Shapley, R. & Victor, J. D. (1979), 'The contrast gain control of the cat retina.', *Vision Res* 19(4), 431–434.

Shepherd, G. M. G., Stepanyants, A., Bureau, I., Chklovskii, D. & Svoboda, K. (2005), 'Geometric and functional organization of cortical circuits.', *Nat Neurosci* 8(6), 782–790.
URL: http://dx.doi.org/10.1038/nn1447

Shou, T., Li, X., Zhou, Y. & Hu, B. (1996), 'Adaptation of visually evoked responses of relay cells in the dorsal lateral geniculate nucleus of the cat following prolonged exposure to drifting gratings.', *Vis Neurosci* 13(4), 605–613.

Simons, D. J. (1978), 'Response properties of vibrissa units in rat si somatosensory neocortex.', *J Neurophysiol* 41(3), 798–820.

Smith, E. C. & Lewicki, M. S. (2006), 'Efficient auditory coding.', *Nature* 439(7079), 978–982.
URL: http://dx.doi.org/10.1038/nature04485

Sobel, E. C. & Tank, D. W. (1994), 'In vivo ca2+ dynamics in a cricket auditory neuron: An example of chemical computation.', *Science* 263(5148), 823–826.
URL: http://dx.doi.org/10.1126/science.263.5148.823

Solomon, S. G., Peirce, J. W., Dhruv, N. T. & Lennie, P. (2004), 'Profound contrast adaptation early in the visual pathway.', *Neuron* 42(1), 155–162.

Sompolinsky, H., Yoon, H., Kang, K. & Shamir, M. (2001), 'Population Coding in Neuronal Systems with Correlated Noise.', *Physical Review E*, 64:051904.

Stratford, K. J., Tarczy-Hornoch, K., Martin, K. A., Bannister, N. J. & Jack, J. J. (1996), 'Excitatory synaptic inputs to spiny stellate cells in cat visual cortex.', *Nature* 382(6588), 258–261.
URL: http://dx.doi.org/10.1038/382258a0

Swadlow, H. A. (2002), 'Thalamocortical control of feed-forward inhibition in awake somatosensory 'barrel' cortex.', *Philos Trans R Soc Lond B Biol Sci* 357(1428), 717–1727.
URL: http://dx.doi.org/10.1098/rstb.2002.1156

Swadlow, H. A. (2003), 'Fast-spike interneurons and feedforward inhibition in awake sensory neocortex.', *Cereb Cortex* 13(1), 25–32.

Tadmor, Y. & Tolhurst, D. J. (2000), 'Calculating the contrasts that retinal ganglion cells and LGN neurones encounter in natural scenes', Vision Res 40(22), 3145-3157

Tolhurst, D. J., Movshon, J. A. & Dean, A. F. (1983), 'The statistical reliability of signals in single neurons in cat and monkey striate cortex.', *Vision Res* 23, 775-785

Thomson, A. M. (1997), 'Activity-dependent properties of synaptic transmission at two classes of connections made by rat neocortical pyramidal axons in vitro.', *J Physiol (Lond), 502:131-147.*

Traub, R. & Miles, R. (1991), *Neuronal Networks of the Hippocampus*, Cambridge Univ. Press, Cambridge.

Tsodyks, M. V. & Markram, H. (1997), 'The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability.', *Proc Natl Acad Sci* USA 94(2), 719–723.

Ulanovsky, N., Las, L., Farkas, D. & Nelken, I. (2004), 'Multiple time scales of adaptation in auditory cortex neurons.', *J Neurosci* 24(46), 10440–10453.
URL: http://dx.doi.org/10.1523/JNEUROSCI.1905-04.2004

Ulanovsky, N., Las, L. & Nelken, I. (2003), 'Processing of low-probability sounds by cortical neurons.', *Nat Neurosci* 6(4), 391–398.
URL: http://dx.doi.org/10.1038/nn1032

Varela, J. A., Sen, K., Gibson, J., Fost, J., Abbott, L. F. & Nelson, S. B. (1997), 'A quantitative description of short-term plasticity at excitatory synapses in layer 2/3 of rat primary visual cortex.', *J Neurosci* 17(20), 7926–7940.

Varela, J. A., Song, S., Turrigiano, G. G. & Nelson, S. B. (1999), 'Differential depression at excitatory and inhibitory synapses in visual cortex.', *J Neurosci* 19(11), 4293–4304.

Victor, J. D. & Shapley, R. M. (1979), 'Receptive field mechanisms of cat x and y retinal ganglion cells.', *J Gen Physiol* 74(2), 275–298.

Vidyasagar, T. R. (1990), 'Pattern adaptation in cat visual cortex is a co-operative phenomenon.', *Neuroscience* 36(1), 175–179.

von Heimendahl, M., Itskov, P. M., Arabzadeh, E. & Diamond, M. E. (2007), 'Neuronal activity in rat barrel cortex underlying texture discrimination.', *PLoS Biol* 5(11), e305.
URL: http://dx.doi.org/10.1371/journal.pbio.0050305

Wimmer, K., Hildebrandt, K. J., Hennig, R. M. & Obermayer, K. (2008), 'Adaptation and selective information transmission in the cricket auditory neuron an2.', *PLoS Comput Biol* 4(9), e1000182.
URL: http://dx.doi.org/10.1371/journal.pcbi.1000182

Woolsey, T. A. & der Loos, H. V. (1970), 'The structural organization of layer iv in the somatosensory region (si) of mouse cerebral cortex. the description of a cortical field composed of discrete cytoarchitectonic units.', *Brain Res* 17(2), 205–242.

Yamada, J., Furukawa, T., Ueno, S., Yamamoto, S. & Fukuda, A. (2007), 'Molecular basis for the gabaa receptor-mediated tonic inhibition in rat somatosensory cortex.', *Cereb Cortex* 17(8), 1782–1787.
URL: http://dx.doi.org/10.1093/cercor/bhl087

Yu, Y., Potetz, B. & Lee, T. S. (2005), 'The role of spiking nonlinearity in contrast gain control and information transmission.', *Vision Res* 45(5), 583–592.
URL: http://dx.doi.org/10.1016/j.visres.2004.09.024

Zhang, L. I., Tao, H. W., Holt, C. E., Harris, W. A. & Poo, M. (1998), 'A critical window for cooperation and competition among developing retinotectal synapses.', *Nature* 395(6697), 37–44.
URL: http://dx.doi.org/10.1038/25665

Zhaoping L. (2006), 'Theoretical Understanding of the early visual processes by data compression and data selection' in *Network: Computation in neural systems* 17(4):301-334.

Zohary, E., Shadlen, M. N. & Newsome, W. T. (1994), 'Correlated neuronal discharge rate and its implications for psychophysical performance.', *Nature* 370: 140-143.