DISS ETH NO. 16230

CONTRAST GAIN CONTROL, NOISE AND SPIKE THRESHOLD IN THE EARLY VISUAL SYSTEM

A dissertation submitted to the
SWISS FEDERAL INSTITUTE OF TECHNOLOGY ZURICH

for the degree of

Doctor of Natural Sciences

presented by

VINCENT BONIN

M.Sc., Technische Universität Hamburg-Harburg
born 6 October 1973
citizen of Canada

accepted on the recommendation of

Prof. Rodney Douglas, examiner
Prof. Kevan Martin, co-examiner
Dr. Daniel Kiper, co-examiner
Dr. Matteo Carandini, advisor

2005

To my family

Acknowledgements

Writing a dissertation is not a lone man's work. I am indebted to many wonderful people who have supported me during this exciting journey.

First and foremost, I would like to thank my advisor, Matteo Carandini, for his unconditional support, thorough guidance and infectious enthusiasm. Matteo's commitment to scientific clarity and elegance is a virtue I will always admire (and try to emulate).

I am also grateful to my thesis committee and to others who kindly offered their help at the examination. Thanks to my examiner Rodney Douglas who supported me throughout my doctoral studies. Thanks to Kevan Martin and to Daniel Kiper for critical reviews of the manuscript. Thanks to Danilo Pescia for overseeing the examination and to Simone Schumacher for organizing it.

This thesis is the fruit of four years of close collaboration with Valerio Mante who was also a graduate student in the laboratory and has become a dear friend. Valerio has contributed to nearly every stage of the work, from data collection to computational modeling. I have learned much from our daily interactions and will miss them enormously.

Other members of the laboratory have helped me a great deal. Thanks to Tobe Freeman and Severine Durand for introducing me to the physiological experiments. Severine also helped with collecting the data. I thank her for that and for her steady support. Thanks also to Andrea Benucci and to Robby Frazor for their support during thesis writing.

I am also indebted to institutions that have made this work possible. Thanks to the Neuroscience Center Zurich for putting together a truly multi-disciplinary PhD program. Thanks to the Institute of Neuroinformatics for hosting me while I initiated this work and to the Smith-Kettlewell Eye Research Institute for adopting our laboratory. I have met in these three institutions wonderful friends whose contribution is inestimable.

This thesis has also benefited from a number of ad hoc contributors. Thanks to Samuel Solomon for comments on an early version of Chapters 2 and 3. Thanks also to Constanze Hipp Hofstötter for helpful comments on the German translation of the summary. Thanks also to my wife Kati Mickley for proofreading critical parts of the manuscript.

Kati's contribution, however, goes well beyond proofreading. Her love and encouragement have ensured that I brought this project to a successful end.

Contents

Summary

The patterns of light falling onto our retinas vary from one instant to another. The magnitude of these variations surpasses the dynamic range of the neurons by several orders of magnitude. This is a fundamental problem, which the visual system resolves by rapidly adapting its responses to the prevailing image statistics. One such adaptation mechanism is contrast gain control, which adjusts responses to the local root-mean-square contrast of the stimulus, the local standard deviation of intensities relative to the mean. In this thesis we study the function and properties of this gain control at two consecutive stages of the visual system: the lateral geniculate nucleus (LGN) and the primary visual cortex (V1). Our approach is to record responses from LGN and V1 of anesthetized cats to a same stimulus set, and to design tractable computational models of the responses.

In Chapter 2 we devise an intuitive model of the effects of contrast gain control on LGN responses. The model consists of a suppressive field whose output divides the output of the classical center-surround receptive field. The suppressive field filters the retinal image and computes the standard deviation of the filters outputs (root-mean-square contrast). We test the model on responses to moving gratings of different contrasts, sizes and spatial configurations. The model accurately predicts the responses. It predicts (1) the saturation of responses with increasing contrast, (2) the selectivity of responses for stimulus size and (3) and the masking of the response to a given stimulus by superposition of a second.

In Chapter 3 we characterize the properties of gain control in LGN and discuss the possible sources of the underlying signals. We measure the spatial extent and visual preferences of the suppressive field. The suppressive field is no larger than the receptive field surround. It is not selective for stimulus orientation, and it responds to a wide range of frequencies, including very low spatial frequencies and high temporal frequencies. These properties suggest that the suppressive signals originate in retina or within LGN, but not in feedback from cortex.

In Chapter 4 we use white noise stimuli to study the statistical computation of gain control. We test whether the suppressive field estimates the true standard deviation of light intensities, or whether it also responds to higher-order luminance statistics such as

skewness and kurtosis. Neural gain is strongly affected by the standard deviation but is mostly invariant to skewness and kurtosis. These results validate the suppressive field model for a broad range of luminance statistics.

In Chapter 5 we devise a probabilistic model of the responses of LGN and V1 neurons. Two difficulties arise when comparing LGN and V1 responses. First, V1 neurons exhibit higher spike thresholds than LGN neurons, which might bias our assessment of gain control. Second, V1 responses exhibit more trial-to-trial variability than LGN responses, which can significantly alter the contrast response curves. We resolve these problems by adding a noise source and a rectification stage to the suppressive field model. We then test the model on responses to moving gratings of different contrasts and sizes. In both LGN and V1, the model accurately predicts the average responses and also describes the trial-to-trial variability of responses. The model provides a robust framework to compare responses across brain areas.

In Chapter 6 we compare the effects of gain control on the responses of LGN and V1 neurons. V1 responses saturate more with increasing contrast than LGN responses. They also exhibit distinct patterns of selectivity for stimulus size. We use the probabilistic model to investigate the factors underlying these differences. We consider the contributions of contrast gain control, trial-to-trial variability and spike threshold. We find the differences between LGN and V1 neurons to be explained by differences in gain control and spike threshold. The potential contribution of variability is much smaller. We also find that LGN and V1 neurons exhibit similar degree of size tuning when factoring out the effects of noise and of spike threshold. These results suggest that the LGN input likely contributes substantially to the effects of gain control that are observed in V1.

Zusammenfassung

Die Lichtmuster, die auf unsere Netzhäute (Retinae) fallen, ändern sich von einem Augenblick zum nächsten. Die Schwankung der Lichtintensitäten überschreitet den dynamischen Bereich der Neuronen um ein Vielfaches. Dies stellt ein grundlegendes Problem dar, welches das visuelle System durch schnelle Anpassung seiner neuronalen Antworten an die gegebene Bildstatistik löst. Ein solcher Anpassungsmechanismus ist die Kontrast-Verstärkungsregelung (contrast gain control), welche die Antworten an den lokalen Effektivkontrast (root-mean-square contrast) eines Stimulus anpasst.

In dieser Arbeit untersuchen wir Funktion und Eigenschaften dieses Mechanismus in zwei aufeinanderfolgenden Arealen des visuellen Systems von Katzen: dem Corpus Geniculatum Laterale (CGL) und dem primären visuellen Kortex (V1). In unserem experimentellen Ansatz leiten wir die Reizantworten im CGL und V1 anästhesierter Katzen ab, um darauf basierend einfache Modelle zu entwickeln.

In Kapitel 2 stellen wir ein einfaches Modell vor, das die Wirkung des Kontrast-Verstärkungsregelungsmechanismus auf die Antworten im CGL beschreibt. Das Modell besteht aus einem *suppressiven Feld*, dessen Ausgabe die Ausgabe des klassischen *center-surround* rezeptiven Feldes teilt. Das suppressive Feld entspricht der Standardabweichung von Filtern, die auf das Netzhautbild wirken. Wir testen das Modell an Antworten auf eine Reihe von sich bewegenden Gittermustern („moving gratings“) verschiedener Kontraste und Dimensionen. Das Model sagt die Antworten genau voraus. Es erklärt (1) die Sättigung von Antworten bei Zunahme des Kontrastes, (2) die Selektivität der Antworten bezüglich der Größe des Stimulus und (3) die Maskierung von Antworten auf einen Stimulus bei Überlagerung eines zweiten.

In Kapitel 3 charakterisieren wir die Eigenschaften der Kontrast-Verstärkungsregelung im CGL. Wir messen die räumliche Ausdehnung und die visuellen Präferenzen des suppressiven Feldes. Das suppressive Feld ist nicht größer als das rezeptive Feld. Es ist nicht selektiv bezüglich der Ausrichtung des Stimulus und es reagiert auf eine große Bandbreite von Frequenzen, einschließlich sehr niedriger räumlicher Frequenzen und hoher zeitlicher Frequenzen. Diese Eigenschaften suggerieren, dass die suppressiven Signale in der Retina oder im CGL entstehen, nicht aber durch Rückkopplung vom Kortex.

In Kapitel 4 benutzen wir stochastische Stimuli, um die Sensitivität des Verstärkungsregelungsmechanismus bezüglich der Statistik der Lichtintensitäten zu untersuchen. Wir testen, ob das suppressive Feld ausschiesslich die Standardabweichung der Lichtintensitäten misst oder zusätzlich auf Statistiken höherer Ordnung, wie Schiefe (skewness) und Wölbung (kurtosis), reagiert. Die neurale Verstärkung wird durch die Standardabweichung stark beeinflusst, nicht aber durch Schiefe und Wölbung. Diese Resultate bestätigen das Modell für eine Vielzahl von Lichtintensitätsstatistiken.

In Kapitel 5 weiten wir unsere Untersuchung des Kontrast-Verstärkungsregelungsmechanismus auf die Antworten im primären visuellen Kortex aus. Dabei bewältigen wir zwei Schwierigkeiten. Erstens haben V1-Neuronen hohe Schwellenwerte, was dem Effekt der Verstärkungsregelung entgegenwirken kann. Zweitens weisen V1-Antworten eine erhebliche Variabilität auf. Wir lösen diese Probleme, indem wir eine Gleichrichtungsstufe und eine Rauschquelle in das Modell einbauen. Wir testen das Modell anhand seiner Antworten auf bewegliche Gittermustern verschiedener Kontraste und Größen. Sowohl für CGL als auch für V1 sagt das Modell akkurat die durchschnittlichen neuronalen Antworten und deren Variabilität bei wiederholten Messungen voraus. Das Modell eignet sich somit dafür die Antworten in verschiedenen Gehirnbereichen zu vergleichen.

In Kapitel 6 vergleichen wir die Effekte der Verstärkungsregelung auf Antworten von CGL- und V1-Neuronen. V1-Antworten sättigen bei Erhöhung des Kontrasts stärker als CGL-Antworten. Beide weisen außerdem unterschiedliche Selektivitätsmuster bezüglich der Stimulusgröße auf. Anhand des stochastischen Modells untersuchen wir die Faktoren, die diesen Unterschieden zu Grunde liegen. Dabei ziehen wir die Beiträge des Verstärkungsregelungsmechanismus, der Streuung und des Schwellenwertes in Betracht. Wir zeigen, dass sich die Unterschiede zwischen CGL- und V1-Neuronen durch die Unterschiede in der Verstärkungsregelung und dem Schwellenwert erklären lassen. Der Effekt der Streuung ist wesentlich geringer. Wir finden außerdem, dass CGL- und V1-Neuronen eine ähnliche Selektivität bezüglich der Stimulusgröße aufweisen, wenn man die Effekte des Schwellenwertes berücksichtigt. Diese Resultate legen nahe, dass die Effekte der Verstärkungsregelung, die im CGL beobachtet wurden, wichtige Implikationen für die Antworten der V1-Neuronen haben.

Preface

This work has been performed in the laboratory of Matteo Carandini and is the fruit of a close collaboration with Valerio Mante. It was started in 2001 at the Institute of Neuroinformatics in Zurich. It was continued from December 2002 until September 2005 at the Smith-Kettlewell Eye Research Institute in San Francisco.

Part of this work has appeared in conference abstracts and proceedings. Chapter 2 was subject of a poster presentation at the 2002 Annual Meeting of the Society for Neuroscience (Bonin et al., 2002), and of an oral presentation at the 2003 Neural Information Processing Systems Conference (Bonin et al., 2004a). The results in Chapter 3 were presented orally at the Annual Meeting of the Vision Sciences Society (Bonin et al., 2003a). The results in Chapter 4 have been presented at the 2005 Computational and Systems Neuroscience conference (Bonin et al., 2005). The results in Chapters 5 and 6 were presented orally at the 2003 Annual Meeting of the Society for Neuroscience (Bonin et al., 2003b) and as a poster at the 2004 Computational and Systems Neuroscience conference (Bonin et al., 2004b).

Relevant to this work are contributions to a closely-related project led by Valerio Mante (Mante, 2005). These were also subject of several presentations in international conferences (Mante et al., 2002, 2003, 2004a, b; Mante et al., 2005a).

Chapter 1

General Introduction

A central goal of visual neuroscience is the development of functional models of neural responses to visual stimuli. Functional models summarize the processing performed by complex neural circuits. They provide a framework to study the encoding of visual information in the brain and can be used to investigate regularities in the visual world. They can also guide research in the underlying biophysics.

A brain area for which a functional model is of practical significance is the lateral geniculate nucleus (LGN). Neurons in LGN relay signals from retina to visual cortex. They are part of an intricate network whose basic circuitry has been laid out (Sherman and Guillery, 2004). Their visual responses have been subject of much investigation and are relatively well characterized (Hubel and Wiesel, 1961; McIlwain and Creutzfeldt, 1967; Sanderson, 1971; Levick et al., 1972; Cleland et al., 1983a; Troy, 1983; Derrington and Lennie, 1984; Sclar, 1987; Cai et al., 1997; Usrey and Reid, 2000).

The functional role of the LGN in visual processing not well understood. LGN responses resemble much the responses of retinal ganglion cells (Hubel and Wiesel, 1961; Levick et al., 1972; Cleland and Lee, 1985). But ganglion cells provide only 5-10% of the inputs to the LGN (Sherman, 2001a). The remaining inputs come from in feedback from neighboring nuclei and visual cortex (Budd, 2004; Sherman and Guillery, 2004). The function of this feedback remains subject of debate (for reviews see Sherman, 2001a; Alitto and Usrey, 2003). Some authors say that the influence of feedback is modulatory and has little effects on the visual properties of the neurons (Hubel and Wiesel, 1977; Sherman, 2001a). Other authors argue that visual properties of LGN responses are strongly shaped by cortical feedback (Sillito and Jones, 2002; Worgotter et al., 2002).

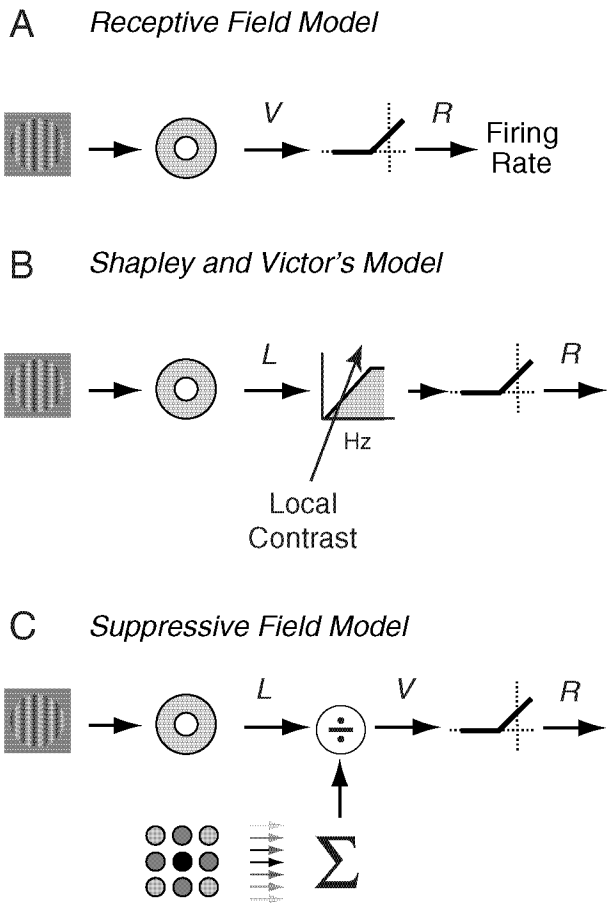

Figure 1.1. Deterministic models of LGN responses. **A**: Receptive field model. The stimulus is processed by a receptive field whose output is rectified to yield firing rates. **B**: Shapley and Victor's model of contrast gain control. The output of the receptive field is fed into a high pass stage whose time constant varies stimulus contrast. **C**: The suppressive field model. The output of the receptive field is divided by the output of a *suppressive field*. The suppressive field computes the standard deviation of light intensities falling in a Gaussian region of visual field.

The simplest yet most widely used functional model of visual responses is the receptive field. The receptive field specifies the weights that a cell applies to each retinal location when summing light intensities of the stimulus. LGN neurons have receptive fields with a center-surround organization. The receptive field center is a small circular region of visual field that provides most of the drive to the neuron. The receptive field surround is a larger concentric area that yields responses of opposite polarity and distinct time course.

The receptive field model is illustrated in Figure 1.1A. Let $S(x,y,t)$ be the pattern of light intensities falling on the retina at every point in space and time. The receptive field is a

linear filter $F(x,y,t)$ that maps the stimulus $S(x,y,t)$ onto a signed function $V(t)$, which we term generator potential, for it is an abstract representation of the membrane potential of the neuron. The generator potential must then be transformed into spikes. A reasonable approximation (Carandini and Ferster, 2000) states that the firing rate $R(t)$ of the neuron is linearly related to the subset of excursions of $V(t)$ that are above a certain threshold.

The receptive field model captures critical aspects of LGN responses. It predicts their selectivity for spatial and temporal frequencies (So and Shapley, 1981; Dawis et al., 1984; Soodak et al., 1987) and captures the main features of responses to complex video sequences (Dan et al., 1996; Mante, 2005). It is customary, therefore, to describe LGN neurons in terms of their center-surround receptive fields (e.g. Cai et al., 1997; Reid et al., 1997; Tavazoie and Reid, 2000; Keat et al., 2001; Einevoll and Plesser, 2002; Xu et al., 2002; Kilavik et al., 2003; Kremers et al., 2004; Lesica and Stanley, 2004; Alitto et al., 2005).

LGN responses exhibit suppressive phenomena that are not explained by the receptive field model. First, responses saturate when stimulus contrast increases (contrast saturation, Derrington and Lennie, 1984; Chino and Kaplan, 1988; Sclar et al., 1990; Kremers et al., 2001; Alitto and Usrey, 2004). Second, responses to a stimulus are reduced by superimposition of another (masking, Cudeiro and Sillito, 1996; Felisberti and Derrington, 1999; Freeman et al., 2002; Girardin et al., 2002; Solomon et al., 2002). Third, responses decrease when stimulus size is increased beyond an optimal value (size tuning, Cleland et al., 1983a; Murphy and Sillito, 1987; Jones and Sillito, 1991; Solomon et al., 2002; Ozeki et al., 2004). These three phenomena are suppressive, because responses are smaller than would be expected from the classical center-surround receptive field alone.

A potential explanation for the suppressive phenomena lies in the mechanism of contrast gain control operating in retina. Discovered by Shapley and Victor (Shapley and Victor, 1978a; Shapley et al., 1981), this mechanism adjusts responses based on the amount of stimulus contrast falling on the receptive field of the neuron. Increasing stimulus contrast reduces gain of responses so that stronger stimuli are required to obtain a given response. Increasing contrast also reduces integration time of the neurons so that responses to high temporal frequencies are enhanced.

Contrast gain control acts rapidly and operates within the integration time of the neurons. The changes in gain and integration time occur within 15 ms following the changes

in stimulus contrast (Victor, 1987; Albrecht et al., 2002). This delay is short considering that early visual responses take up to 100 ms to develop. There is certainly much evolutionary pressure for contrast gain control to operate rapidly. Eyes typically fixate a given location for only 200-300 ms, and typical eye movements will bring the receptive fields of neurons over image patches that differ substantially in contrast (Frazor and Geisler, 2005).

The work of Shapley and Victor yielded a powerful model of ganglion cells responses (Figure 1.1B). In the model, the output of the receptive field is fed into a high-pass stage whose time constant is set by a measure of local contrast (Shapley and Victor, 1981; Victor, 1987; Benardete et al., 1992; Benardete and Kaplan, 1999). The measure of local contrast is computed by an array of subunits distributed over the receptive field of the cell. The model predicts the responses to sine and square wave stimuli, and random steps in luminance of different amplitudes. Shapley and Victor's model might explain the suppressive phenomena because they all involve stimulus manipulation that increase local stimulus contrast.

Two issues must be dealt with before we can assess the contribution of contrast gain control to the suppressive phenomena.

First, we need to devise a model that is adapted to the responses of LGN neurons. Shapley and Victor's model has been thoroughly tested in retina (Shapley and Victor, 1981; Victor, 1987; Benardete et al., 1992; Benardete and Kaplan, 1999) but not in LGN. Retinal responses are generally faithfully relayed by geniculate neurons (Levick et al., 1972; Cleland et al., 1983b; Cleland and Lee, 1985) but there is also evidence for an intrageniculate contribution to gain control (Kaplan et al., 1987; Sclar, 1987; Cheng et al., 1995).

Second, we need to endow the model with a spatial footprint. The model by Shapley and Victor was designed for spatially restricted stimuli and was tested by stimulating the receptive field center and surround separately. Previous studies have suggested that varying the surface area of the stimulus has effects similar to those of varying stimulus amplitude (Shapley and Victor, 1981; Sclar et al., 1990; Benardete and Kaplan, 1999) but the model was not tested on stimuli of different sizes. The model does not account for the actual spatial spread of the subunits underlying gain control and therefore makes no predictions as to what responses should be for these stimuli.

We deal with these issues in Chapter 2. We consider Shapley and Victor's model and show that it can be reduced to a single divisive stage (Figure 1.1C), where local contrast appears in the denominator. We extend the model to the world of arbitrary spatial images by hypothesizing that local contrast is the standard deviation of light intensities computed over a Gaussian-shaped region of visual field, which we term the suppressive field (Levick et al., 1972). We design a set of experiments that taps into the suppressive field and receptive field of LGN neurons and teases out their relative contributions to the responses. We run these experiments on LGN neurons of anesthetized, paralyzed cats. We use the resulting measurements to constrain model parameters and to test model predictions. Finally we ask if contrast gain control can explain the suppressive phenomena.

The model and experiments above provide us with an assay to study fundamental properties of contrast gain control. Previous studies of gain control focused on the effects of the temporal waveform of the stimulus, but did not quantitatively address how stimulus contrast is integrated across space. For example, what is the spatial spread of the subunits performing gain control? How does this spread compare to the size of the receptive field? Moreover, what is the selectivity of the subunits for different stimulus attributes? Does it prefer high spatial frequencies or low temporal frequencies? A quantitative assessment of these properties is critical if we ought to use the model to predict responses to more complex stimuli such as encountered in the real world. A better understanding of these properties may also reveals constraints about the physiological origins of the suppressive phenomena.

We address these issues in Chapter 3. We use model fits to infer the spatial extent and the selectivity for spatial frequency of suppressive field. We present additional experiments to estimate its selectivity for temporal frequency and orientation. We relate the properties of the suppressive field to those of the classical center-surround receptive field. We then relate them to well-known properties of responses in retina and primary visual cortex. We discuss how these results put constraints on the physiological origins of the suppressive field.

Another important issue is how the responses of the subunits are pooled together to yield gain control. Shapley and Victor postulated that gain control estimate root-mean-square contrast by summing the rectified outputs of the subunits. While this hypothesis points to a plausible physiological substrate within the retinal circuit, it does not address the actual statistical processing performed by gain control. Does contrast gain control really compute root-mean-square contrast? Could it be preferentially sensitive to stimulus

fluctuations lying above or below the mean light level? Previous studies could not distinguish between these possibilities because they employed stimuli whose light intensities fluctuate symmetrically around the mean.

We address the statistical operation of the suppressive field in Chapter 4. In the previous chapters we had assumed the suppressive field to compute local root-mean-square contrast. To test this hypothesis we measure the gain of LGN neurons using white noise stimuli. We investigate how gain depends on the luminance statistics of the stimulus. We independently vary the standard deviation, the skewness and the kurtosis of the stimulus and ask these manipulations affect gain. These results complete the validation of the suppressive field model of LGN responses.

A longstanding issue regards the impact of retinal and thalamic gain control on responses at subsequent stages of the visual system. The suppressive phenomena are pronounced so they are likely to have important consequences on the responses of neurons in primary visual cortex (V1), the principal recipient of LGN input. In fact, V1 responses also exhibit the phenomena of contrast saturation (Albrecht and Hamilton, 1982; Sclar et al., 1990), size tuning (Sceniak et al., 1999; Sceniak et al., 2001; Cavanaugh et al., 2002a) and masking (Morrone et al., 1982; Bauman and Bonds, 1991; DeAngelis et al., 1992; Sengpiel and Blakemore, 1994; Carandini et al., 1997; Allison et al., 2001). These effects are thought to be more pronounced in V1 than in LGN (Sclar et al., 1990; Jones et al., 2000).

Surprisingly, the contribution of retinal and thalamic gain control to suppression in V1 has received little attention. The origins of V1 suppression are subject of debate (Angelucci et al., 2002; Freeman et al., 2002; Sengpiel and Vorobyov, 2005). Early explanations involved intracortical inhibition (Benevento et al., 1972; Blakemore and Tobin, 1972). Recent work from our laboratory suggested a potential role for synaptic depression (Carandini et al., 2002; Freeman et al., 2002). Most investigations of cortical gain control assume that V1 neurons receive a linear input and therefore do not consider the contribution of retinal and thalamic gain control.

We begin to address this issue in Chapter 5. We devise a model to compare the effects of gain control on the responses of LGN and V1 neurons. We overcome two important difficulties. A first difficulty is that V1 neurons have high spike thresholds, which can bias our assessment of gain control. A second difficulty is that V1 responses exhibit substantial

trial-to-trial variability (Tolhurst et al., 1981; Bradley et al., 1987; Vogels et al., 1989; Geisler and Albrecht, 1997; Gur et al., 1997; Reich et al., 1997; Buracas et al., 1998). This variability is more pronounced in V1 than in LGN (Kara et al., 2000) and can strongly affect the tuning of the neurons (Anderson et al., 2000; Chance and Abbott, 2002; Miller and Troyer, 2002).

The resulting model accounts for the effects of contrast gain control, noise in membrane potential and spike threshold (Figure 1.2). The model features (1) an integration field that captures the drive provided to the neuron, (2) a divisive suppressive field that describes the effects of contrast gain control, (3) a noise source that describes the trial-to-trial variability of responses and (4) a rectification stage that models the effects of spike threshold.

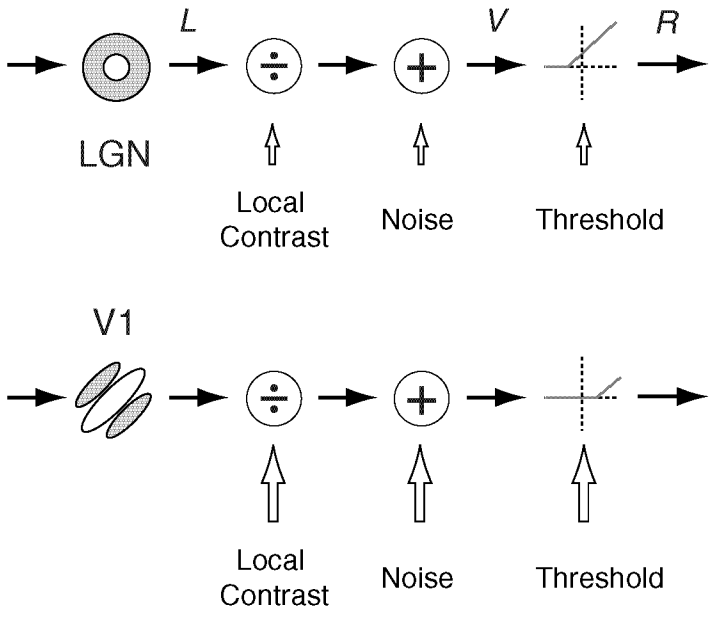

Figure 1.2. Probabilistic model of LGN and V1 responses.

We test the model on responses to stimuli of different contrast and sizes recorded in anesthetized, paralyzed cats. We ask whether the model predicts the average responses. We also ask whether it captures the trial-to-trial variability of responses. The model provides a framework to compare responses across brain areas.

We apply the model in Chapter 6 where we study the growth and saturation of responses with increasing contrast, and the tuning of responses for stimulus size. We compare LGN and V1 responses and ask whether or how they differ. We then use the model

to tease apart the contributions of contrast gain control, response trial-to-trial variability and spike threshold. We then discuss the potential contribution of the LGN input to the effects observed in V1.

Chapter 2

The suppressive field of lateral geniculate neurons

The responses of neurons in lateral geniculate nucleus (LGN) are shaped by powerful suppressive phenomena: contrast saturation, size tuning, and masking. These phenomena cannot be explained by the classical center-surround receptive field. We ask whether they can be explained by a simple mechanism of gain control similar to the one at work in retina. We devise an intuitive model of gain control where the output of the receptive field is divided by the output of a suppressive field. The suppressive field is a Gaussian-shaped region that computes the local root-mean-square contrast. We show that the suppressive field is entirely responsible for the suppressive phenomena. The prevailing view of LGN responses, based on the receptive field alone, is adequate only for small, low-contrast stimuli.

2.1 Introduction

To a first approximation, the responses of neurons in lateral geniculate nucleus (LGN) can be described by the classical center-surround receptive field. The receptive field specifies the weights that a cell applies to each retinal location when summing light intensities. It predicts the selectivity of responses measured with moving gratings (So and Shapley, 1981; Dawis et al., 1984; Soodak et al., 1987) and the main features of responses to complex video sequences (Dan et al., 1996). It is customary, therefore, to describe LGN responses in terms of the receptive field alone (e.g. Cai et al., 1997; Reid et al., 1997; Tavazoie and Reid, 2000; Keat et al., 2001; Einevoll and Plesser, 2002; Xu et al., 2002; Kilavik et al., 2003; Kremers et al., 2004; Lesica and Stanley, 2004; Alitto et al., 2005).

LGN responses, however, exhibit suppressive phenomena that cannot be explained by the receptive field alone. One of these phenomena is size tuning (Figure 2.1**A-B**, *bottom row*): enlarging a grating beyond an optimal size reduces responses (Cleland et al., 1983a; Jones and Sillito, 1991; Solomon et al., 2002). Another of these phenomena is contrast saturation (Figure 2.1 **A-B**, *last column*): responses tend to plateau as contrast is increased (Chino and Kaplan, 1988; Sclar et al., 1990; Kremers et al., 2001). The receptive field alone, by comparison, predicts that responses increase with stimulus diameter and grow linearly with stimulus contrast (Figure 2.1**B**, *dashed curves*).

These suppressive phenomena are related to one another. For instance, size tuning depends on contrast (Solomon et al., 2002; Ozeki et al., 2004), being strong at high contrast and absent at low contrast (Figure 2.1, compare *bottom* and *top row*). Similarly, contrast saturation depends on stimulus size, being prominent for large stimuli but minimal with smaller stimuli (Figure 2.1, compare *first* and *last column*).

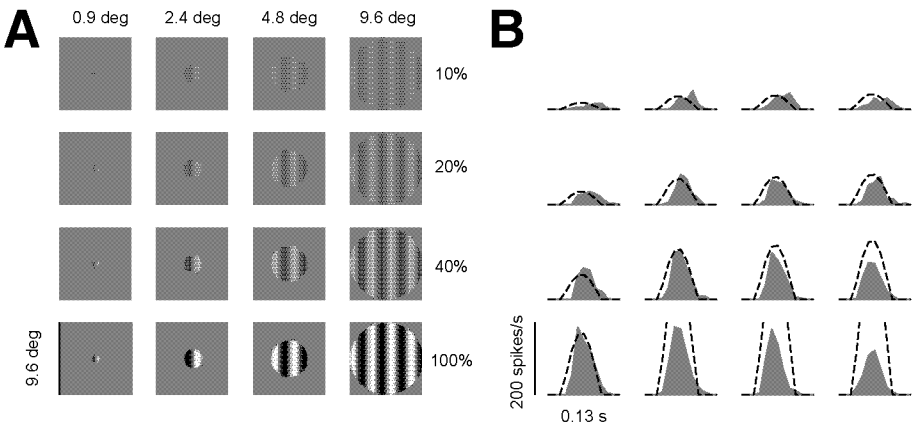

Figure 2.1. Contrast saturation and size tuning in an LGN neuron. **A**: Stimuli are drifting gratings varying in contrast and diameter. **B**: Firing rates averaged over one stimulus cycle. *Curves* are predictions of the receptive field followed by a rectification fitted to responses to low contrast stimuli. ON X-type cell (cell 33.1.3).

Previous investigations of the suppressive phenomena have led to models that account for one or another phenomenon individually. Some have invoked a compressive nonlinearity, which explains contrast saturation but not the effects of stimulus size (e.g. Derrington and Lennie, 1984; Sclar et al., 1990; Felisberti and Derrington, 1999; Przybyszewski et al., 2000; Kremers et al., 2001; Solomon et al., 2002; Webb et al., 2002; Alitto and Usrey, 2004). Others have postulated a large suppressive region whose output

subtracts the output of the receptive field (Solomon et al., 2002; Nolt et al., 2004; Ozeki et al., 2004). These models can explain size tuning but not the effects of contrast.

We hypothesized that the suppressive phenomena might be explained by the mechanism of contrast gain control discovered in retina by Shapley and Victor (1978b; 1981). This mechanism dictates that the responsiveness of the neurons is inversely related to a neural measure of local contrast (Victor, 1987). Gain control is thought to be particularly strong in LGN (Kaplan et al., 1987; Cheng et al., 1995). It might explain the suppressive phenomena, because these phenomena result from stimulus manipulations that increase local stimulus contrast.

The model by Shapley and Victor, however, has limitations. First, it does not provide an intuitive description of the effects of contrast on responses: to account for response dynamics, the model is expressed in terms of filters with variable time constants, not directly in terms of gain. Second, the model is not defined for arbitrary spatial images: it does not specify how local contrast should be integrated across space.

To overcome the first limitation, we simplify the Shapley and Victor model, and find that at heart it involves a simple division. In this division, the output of the receptive field is in the numerator, and local contrast is in the denominator. Seen this way, contrast gain control in retina is very similar to the gain control mechanisms seen in visual cortex (Albrecht and Geisler, 1991; Heeger, 1992; Carandini et al., 1997).

To overcome the second limitation, we hypothesize that local contrast is the standard deviation of the luminance falling under a Gaussian region concentric with the receptive field; we call this region "suppressive field". This term was coined by Levick et al. (1972), to describe a region that is engaged equally by light or dark stimuli, suppresses the cell's responses, and cannot elicit a response by itself.

We integrate these elements into a compact, intuitive and testable model of LGN responses. The model includes (1) the classical receptive field, (2) the suppressive field, and (3) divisive gain control. We ask whether the model captures LGN responses, and in particular if it explains the suppressive phenomena mentioned above.

2.2 Results

The model by Shapley and Victor can be simplified to a single divisive stage, in which local contrast appears in the denominator. In the original model, increases in stimulus contrast decrease the time constant of a high-pass filter, thus reducing response integration time and gain (Shapley and Victor, 1981; Victor, 1987; Benardete et al., 1992; Benardete and Kaplan, 1999). The model captures the observation that gain control is strongest for stimuli of low temporal frequency. For these stimuli, the model predicts that neural responses follow

$$R_{test} \approx \frac{L}{c_{50} + c_{local}}, \qquad 2.1$$

where L is the response of the receptive field (which grows linearly with contrast), c_{50} is a constant, and c_{local} is local contrast. The derivation of this expression is provided in Appendix 1.

Shapley and Victor did not specify how local contrast c_{local} should be computed, but provided important constraints on this computation. First, c_{local} is independent of stimulus position within the receptive field (Shapley and Victor, 1978b). Second, c_{local} is an "even function" of stimulus intensity, i.e. one that is the same for stimuli that are lighter or darker than the mean luminance (Victor, 1987). Third, c_{local} grows with stimulus area (Shapley and Victor, 1981). These observations led Shapley and Victor to suggest that local contrast is computed by summing the positive outputs of an array of subunits distributed over the receptive field, and perhaps beyond.

We hypothesize that local contrast c_{local} is the standard deviation of the luminance falling under a Gaussian region concentric with the receptive field, the suppressive field. The standard deviation is proportional to root-mean-square contrast (Peli, 1990), it fulfills the above constraints, and it can be applied to arbitrary images.

We integrate these concepts into a simple model, which contains two image-processing pathways, the receptive field and the suppressive field (Figure 2.2). The two pathways meet at a divisive stage, where the receptive field provides the numerator and the suppressive field, once its output is added to a constant c_{50}, provides the denominator. The resulting output V is then rectified to obtain a positive firing rate R.

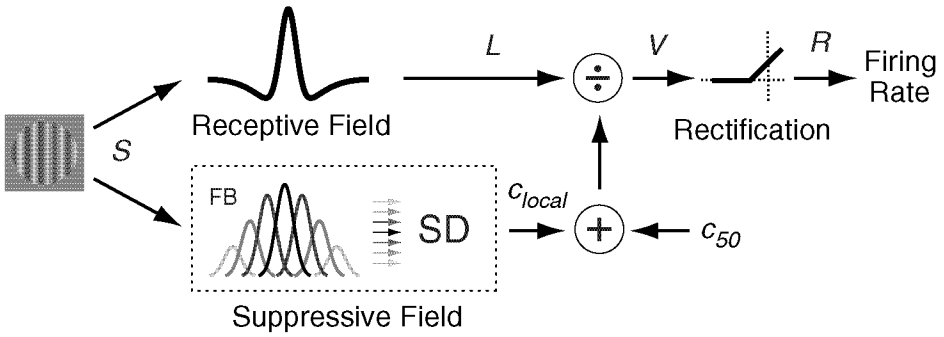

Figure 2.2. Model of LGN responses. The model includes a receptive field and a suppressive field. The receptive field has the classical center-surround organization (difference of Gaussians). The suppressive field computes the standard deviation (SD) of the outputs of a Gaussian-weighted bank of filters (FB), and sums the result to a constant, c_{50}. The signals from receptive field and suppressive field meet at a divisive stage. The output of the division is then rectified to yield positive firing rates.

The suppressive field computes c_{local}, the local standard deviation of the stimulus luminance. First, the image is processed through a bank of filters (subunits, in the terminology of Shapley and Victor, 1978). Second, the filters outputs are weighted by a 2-dimensional Gaussian, a window that represents the spatial extent of the suppressive field. Third, the standard deviation of the resulting signals is computed.

We test the validity of the model on responses from 34 isolated neurons in LGN of anesthetized, paralyzed cats.

The test proceeds in three steps. First, we estimate the receptive field and the strength of the suppressive field from a basic set of measurements. Second, we test model predictions on a set of masking experiments and thereby characterize the suppressive field. Third, we fix model parameters as estimated in the previous steps, and test model predictions on responses to drifting gratings varying in contrast and diameter. We illustrate the procedure with data from the same cell as in Figure 2.1.

2.2.1 Basic characterization

We characterize the receptive field by fitting responses to large drifting gratings varying in spatial frequency (Figure 2.3**A**). We model the receptive field as a difference of Gaussians (Rodieck, 1965; Enroth-Cugell and Robson, 1966). This model has three free parameters: the width of the center Gaussian, the width of the surround Gaussian, and the strength of the surround relative to the center. As expected from previous reports (So and Shapley, 1981;

Dawis et al., 1984; Cai et al., 1997; Kremers et al., 2001), this receptive field model provides excellent fits (Figure 2.3**A**, *curve*).

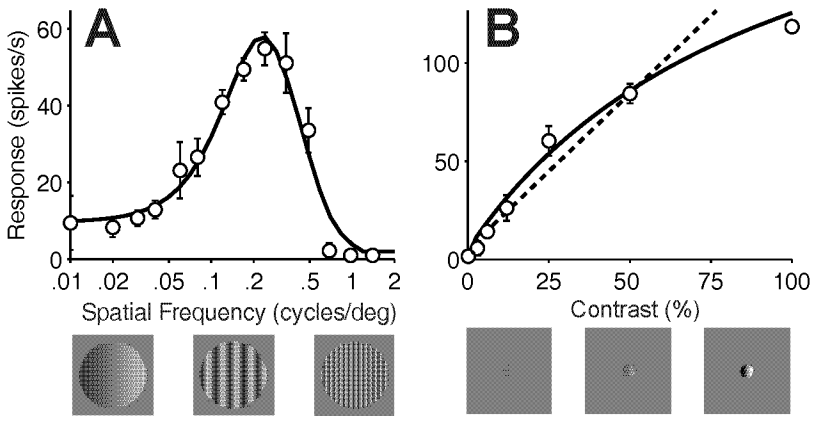

Figure 2.3. Characterizing the shape of the receptive field and the strength of the suppressive field. **A**: Selectivity for spatial frequency measured with large drifting gratings (0.24 cycles/deg, 7.8 Hz, 30 deg diameter). **B**: Effect of stimulus contrast measured with gratings of optimal diameter (1.4 deg). Responses (*symbols*) are the harmonic component of spike trains at the grating temporal frequency. Error bars indicate ± 1 s.d. *Dashed line* indicates predictions of the receptive field followed by rectification. *Curves* indicate fit of complete model. Model parameters are the same here, in Figure 4 and Figure 5 (see Table 1 in Methods for values). To account for spontaneous changes in responsiveness we allow two parameters, V_{max} and V_0, to vary across experiments. Their values are 167 and -3 spikes/s in **A**, and 273 and -6 spikes/s in **B.** Cell 33.1.3.

We then gauge the overall strength of the suppressive field by fitting responses to gratings of increasing contrast (Figure 2.3**B**). We stimulate with drifting gratings of optimal spatial and temporal frequency, contained in optimally-sized windows. For these stimuli, the outputs of both receptive field and suppressive field are simply proportional to grating contrast c, so the model predicts that firing rate obeys $R \approx c/(c_{50}+c)$. This expression is known to capture the contrast-responses of LGN neurons (usually contrast is elevated to a power of n, with n close to one, Derrington and Lennie, 1984; Sclar et al., 1990; Felisberti and Derrington, 1999; Kremers et al., 2001). Indeed, it provided excellent fits to our data (Figure 2.3**B**, *solid curve*), explaining >80% of the variance in 34/34 cells. By comparison, the rectified output of the receptive field alone provided inadequate fits (Figure 2.3**B**, *dashed line*).

2.2.2 Masking

To assess whether the model captures more complex properties of LGN responses, we turn to stimuli composed of two drifting gratings. With such stimuli we probe phenomena of *masking,* the suppression of responses that occurs when one grating extends in the region surrounding the other (Jones and Sillito, 1991; Solomon et al., 2002; Webb et al., 2002) or when the gratings are superimposed (Freeman et al., 2002). In our measurements, the two gratings are a fixed optimal *test* at 50% contrast and a *mask* whose attributes are varied (Figure 2.4). Test and mask have incommensurate temporal frequencies, which elicit distinct response components (Bonds, 1989): a *test response* that oscillates at the temporal frequency of the test, and a *mask response* that oscillates at the frequency of the mask. We first investigate how test responses are affected by mask attributes.

The model predicts that adding a mask should reduce test responses. The predicted test responses are given by Equation 2.1, where the mask affects only the output of the suppressive field (the term c_{local} in the denominator). Indeed, the mask does not affect the output of the receptive field at the frequency of the test (the term L in the numerator). Adding the mask, therefore, increases c_{local} and thus reduces test responses.

We test the model's predictions by varying mask contrast (Figure 2.4**A**), diameter (Figure 2.4**B**), and spatial frequency (Figure 2.4**C**), and find that the model yields good fits (Figure 2.4**A-C**, *solid curves*). (1) Increasing mask contrast progressively suppresses the response (Figure 2.4**A**). The model predicts this effect because increasing mask contrast increases the local contrast c_{local}. With masks of low contrast there is little suppression because the output of the suppressive field is dominated by the constant c_{50}. (2) Increasing mask diameter progressively suppresses test responses until responses reach a plateau (Figure 2.4**B**). The model captures this decrease because increasing mask diameter increases local contrast c_{local}. The plateau is reached once the mask extends beyond the suppressive field. (3) Suppression depends on mask spatial frequency (Figure 2.4**C**), being strong at low spatial frequencies and absent at high frequencies. The model captures this behavior thanks to the filters at the front-end of the suppressive field, which endow the suppressive field with frequency preferences.

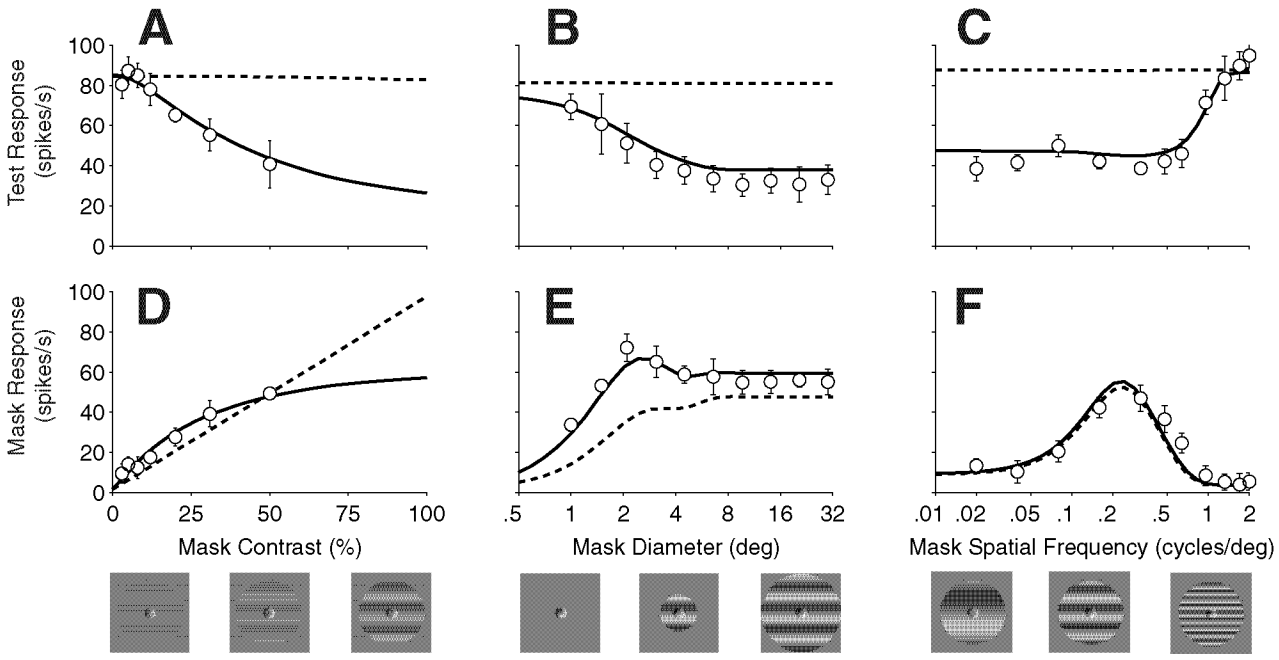

Figure 2.4. Masking. Stimuli are sums of a test grating and a mask grating drifting with incommensurate temporal frequencies. Responses are measured at the frequency of the test, 7.8 Hz (*test response*, **A-C**) and at the frequency of the mask, 12.5 Hz (*mask response*, **D-E**). *Curves* show model fit. *Dashed lines* indicate predictions of receptive field followed by rectification. **A,D**: Responses as function of mask contrast. (V_{max}=273,V_0=-6) **B,E**: Responses as function of mask diameter. (V_{max}=242,V_0=-6) **C,F**: Responses as function of mask spatial frequency. (V_{max}=275,V_0=-4) Test and mask had spatial frequency of 0.24 cycles/deg (unless varied), and diameter of 1.4 deg and 14.1 deg (unless varied). Cell 33.1.3.

The suppressive field plays a critical role in explaining these observations. The fits performed without it (Figure 2.4, *dashed curves*) are clearly inadequate, indicating that neither receptive field nor rectification stage can account for masking. The fits performed including the suppressive field, instead, have high quality: the model explains > 80% of the variance in 31/34 neurons.

These experiments allow us to constrain key model parameters, which are henceforth kept fixed. Fitting the effects of mask contrast we estimate the constant c_{50} (Figure 2.4**A**). Fitting the effects of mask diameter we estimate the size of the suppressive field (Figure 2.4**B**). Fitting the effects of mask spatial frequency we estimate the parameters of the filters at the front-end of the suppressive field (Figure 2.4**C**). A summary of model parameters is presented in Table 2.1.

Stage	Parameter	Symbol	Units	Example Cell		Population (N=34)	
				Min	*Max*	*Mean*	*S.d.*
Rectification	Responsiveness	V_{max}	spikes/s	167	275	212	206
	Threshold	V_0	spikes/s	-6	-2	-7.8	7.1
Receptive Field	Center size	σ_{ctr}	deg	0.5		0.6	0.4
	Surround size	σ_{srd}	deg	1.5		2.0	1.1
	Surround strength	k_{srd}	-	0.9		0.8	0.2
	Effective mask strength	α_{mask}	-	0.6		1.1	0.8
Suppressive Field	Size	σ_{SF}	deg	1.4		2.5	2.9
	Strength	c_{50}	-	0.1		0.3	0.2
	Filter center size	σ_u	deg	0.3		0.4	0.2
	Filter surround size	σ_d	deg	0.5		1.8	2.6
	Filter surround strength	k_d	-	0.5		0.4	0.3

Table 2.1. Summary of model parameters. Example cell: 33.1.3.

Having obtained model parameters from the test responses, we ask if the model can also predict the mask responses, the components of responses at the temporal frequency of the mask (Figure 2.4**D-F**). In this case varying mask attributes varies the outputs of both receptive field and suppressive field. Mask responses display strong saturation (Figure 2.4**D**, compare with Figure 2.3**B** where there is less overall contrast); they are tuned for stimulus size (Figure 2.4**E**); and their tuning for spatial frequency is bandpass (Figure 2.4**F**). To predict these responses, we need to estimate α_{mask}, the effectiveness of the mask in driving the receptive field relative to the test. We estimate α_{mask} from data in which mask contrast is varied (Figure 2.4**D**) and use the same value for the remaining experiments (Figure 2.4**E-F**). The resulting model predictions yield good fits to the mask responses (Figure 2.4, *solid curves*). Fits of similar quality were obtained in the remaining cells, with the model explaining > 80% of the variance in 28/34 cells.

2.2.3 Size tuning and contrast saturation

We now ask whether the model predicts the two related phenomena mentioned in the Introduction: size tuning and contrast saturation (Figure 2.1). We test the model on responses to single drifting gratings whose diameter and contrast are varied independently. These data provide a stringent test for the model, because they were not used to constrain the model parameters.

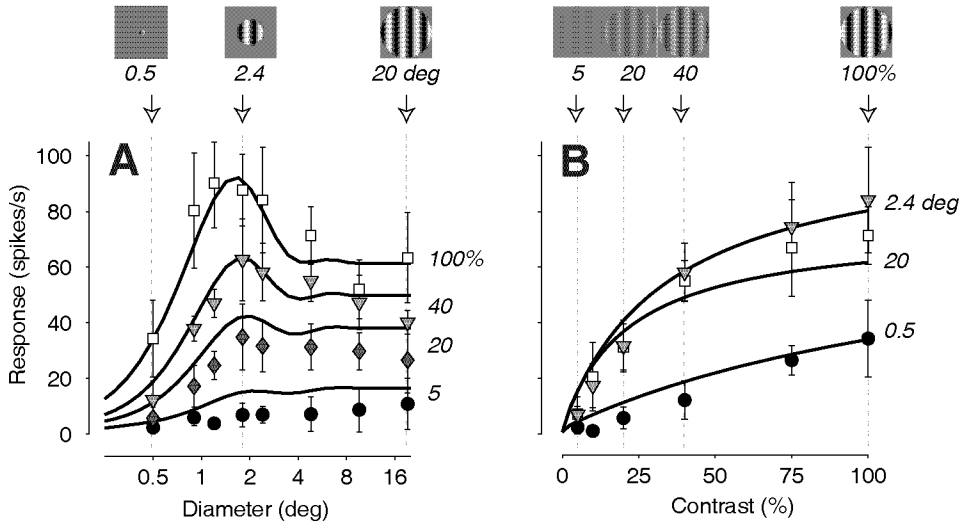

Figure 2.5. Size tuning and contrast saturation. Stimuli are gratings varying in diameter and contrast. *Curves* are predictions of model with parameters held fixed from previous measurements (Figure 2.3 and Figure 2.4, V_{max}=128,V_0=-2) **A**: Responses as a function of diameter, for selected contrasts **B**: Same data, plotted as a function of contrast, for selected diameters. Stimuli had optimal attributes: 0.24 cycles/deg and 7.8 Hz. Cell 33.1.3. (93.9% explained variance).

The model predicts size tuning, and how this tuning depends on contrast (Figure 2.5**A**). Size tuning is pronounced at high contrasts (Figure 2.5**A**, *lighter symbols*) but weak at low contrast (Figure 2.5**A**, *darker symbols*) (Solomon et al., 2002; Ozeki et al., 2004). The model captures these effects (Figure 2.5**A**, *curves*). At high contrasts, responses are tuned for size because increasing size increases the output of the suppressive field, c_{local}, while the output of the receptive field, instead, is roughly constant. At low contrasts, instead, c_{local} is much smaller than the constant c_{50}, and size tuning is not observed.

These effects were common in our sample (Figure 2.6). To quantify the degree of size tuning we estimate the amplitude of plateau responses, the responses to the largest stimulus tested, as a fraction of the peak responses. On average, at 10% contrast, plateau responses are 27±20% smaller than the peak responses (mean±s.d., N=34). The degree of size tuning increases with contrast and peaks at 100% contrast, where plateau responses are 39±18% below peak responses. The model predicts a similar trend, with values of 11±9% at 10% contrast, and 37±15% at 100% contrast.

The flip side of these effects is that preferred stimulus size decreases with contrast (Solomon et al., 2002). To quantify this effect we estimate the diameter of the stimulus

eliciting maximal response. This preferred stimulus diameter decreases as contrast is increased, ranging from 3.8±2.0 deg at 10% contrast to 2.4±1.7 deg at 100% contrast. The model captures this effect with predicted preferred diameters of 4.1±3.0 deg at 10% contrast and of 2.5±1.6 deg at 100% contrast.

The model also predicts a phenomenon that had hitherto gone unnoticed: that contrast saturation is pronounced only for large stimuli (Figure 2.5**B**). Responses to large stimuli show strong saturation (Figure 2.5**B**, *lighter symbols*). The model predicts saturation because large, high contrast stimuli elicit a large response from the suppressive field. Conversely, responses to small stimuli grow more linearly with contrast (Figure 2.5**B**, *darker symbols*). The model predicts linearity because small stimuli induce in the suppressive field a response that is negligible compared to the constant c_{50}.

These effects are common in our sample. In most cells, saturation is strong for large stimuli (Figure 2.6**B**, *white symbols*), weaker for optimal stimuli (Figure 2.6**B**, *gray symbols*), and largely absent for the smallest stimuli (Figure 2.6**B**, *black symbols*). To summarize the degree of saturation across cells, we fit a power-law to the contrast-response curve at each stimulus diameter. We perform this analysis on both data and model predictions. Power-law exponents close to 0 indicate strong saturation; exponents close to 1 indicate linear growth. With an average exponent of 0.78±0.41, saturation was weakest at the smallest stimulus diameter eliciting a reliable response. Saturation monotonically increases with stimulus diameter, with an average exponent of 0.46±0.18 at the largest diameters tested. The model predicts the trend yielding exponents of 0.61±0.11 for small diameters, and 0.40±0.16 for large diameters.

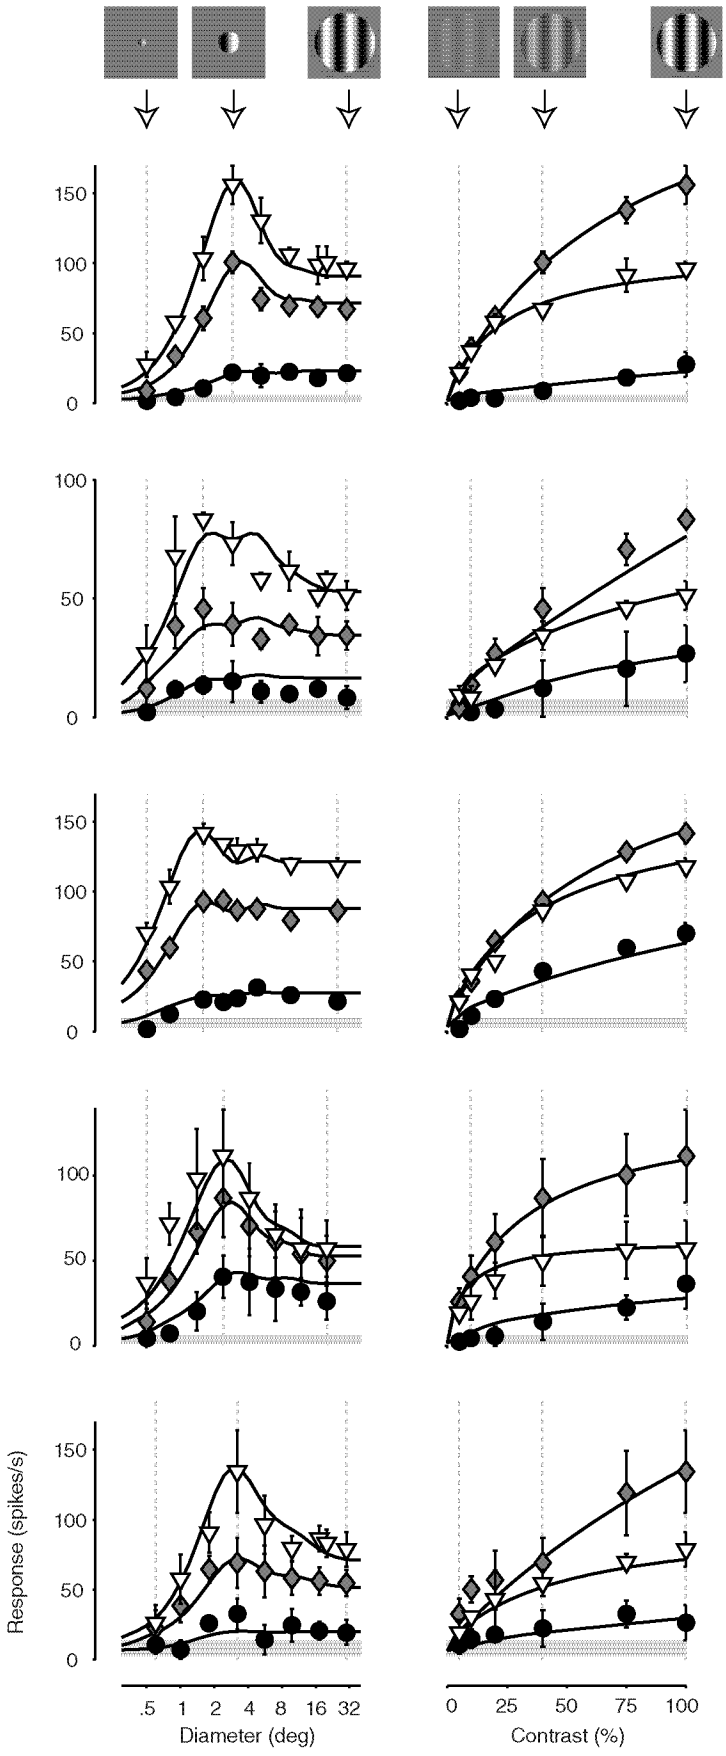
150
100
50
0
100
50
0
150
100
50
0
100
50
0
150
100
50
0
Response (spikes/s)
.5 1 2 4 8 16 32
Diameter (deg)
0 25 50 75 100
Contrast (%)

Figure 2.6. Size tuning and contrast saturation in five additional cells. Details as in Figure 2.5. Gray areas indicate spontaneous response ± 1 s.d. *Left column:* Responses as a function of diameter, shown here for two contrasts: the lowest one eliciting a reliable response (*black*), 40% (gray) and 100% (*white*). *Right column*: Responses expressed as a function of contrast, shown here for three diameters: the smallest diameter eliciting a reliable response (*black*), the optimal diameter (*gray*) and the largest diameter tested (*white*). *Dotted lines* indicate selected contrasts and diameters. *Top to bottom*: Cells 33.3.4 (OFF/X, 98.5% explained variance), 28.2.5(ON/X, 91.8%), 31.3.3 (OFF/X, 98.7%), 35.3.3 (OFF/X, 93.7%), 31.2.2 (OFF/Y, 93.3%).

2.2.4 Model performance

To measure the quality of model predictions, we calculate the percentage of the variance in the data that is explained by the model. Over the entire population, the model performs well in a vast majority of neurons, explaining > 80% of the variance in 29/34 neurons. This performance is remarkable, given that model parameters were held fixed from previous fits (those seen in Figure 2.3 and Figure 2.4).

These results were homogeneous across receptive field polarities and cell types. Explained variance averaged 89.8±6.0% in X cells (mean±s.d., N=28) and 81.5±12.5% in Y cells (N=6). Its median values were not significantly different (p>0.1, rank sum). Explained variance averaged 89.5+/-5.9% in ON cells (mean±s.d., N=28) and 86.8+/-10.0% in OFF cells (N=6).

To investigate model performance in individual stimulus conditions, we compute z-scores of the deviations between measured and predicted responses. The amplitudes of deviations are mostly below 1 ($|z|$=0.65, median, N=34) indicating that, for most stimuli, model predictions are as accurate as the trial-to-trial variability of responses.

The model predictions, however, are not always accurate, and show some significant biases. We estimate accuracy from the unsigned z-scores, conditioned on stimulus contrast or diameter. Predictions are most accurate for diameters slightly larger than the diameter eliciting maximal responses ($|z|$=0.52, median) and least accurate for small and large diameters ($|z|$ = 0.68 and 0.74). Similarly, predictions are most accurate at 40% contrast ($|z|$=0.52,) and least accurate at 5% and 100% contrast ($|z|$=0.86 and 0.64). We estimate bias from signed z-scores. The model overestimates the amplitudes of responses to large stimuli (z=-0.71), whereas predictions for small stimuli show little bias (z = -0.16). Likewise, the model overestimates the amplitudes of responses to low contrast stimuli (z = -0.89, 5% contrast) while predictions for high contrast stimuli show little bias (z = -0.06).

2.3 Discussion

We have devised a simple model of gain control and shown that it provides a parsimonious explanation of the powerful suppressive phenomena seen in LGN responses. The model extends an established model of retinal gain control to the domain of arbitrary spatial images. It specifies how local contrast is measured by a suppressive field, whose output divides the output of the classical receptive field. The model predicts the phenomena of masking, size tuning and contrast saturation and explains how these phenomena are interrelated. In particular it explains how size tuning is pronounced at high contrast and weak at low contrast, and how increasing stimulus contrast reduces the size of the preferred stimulus. Furthermore the gain control model predicts a phenomenon that had hitherto gone unnoticed: that contrast saturation is present only for large stimuli; responses to small stimuli grow linearly with contrast.

2.3.1 Comparison with previous models

Previous models of LGN responses were limited to explaining individual phenomena. The most widespread model includes only a receptive field, possibly followed by a rectification stage – the linear pathway in our model (Stevens and Gerstein, 1976; So and Shapley, 1981; Dawis et al., 1984; Soodak et al., 1987; Dan et al., 1996; Cai et al., 1997; Tavazoie and Reid, 2000; Keat et al., 2001; Einevoll and Plesser, 2002; Xu et al., 2002; Kilavik et al., 2003; Kremers et al., 2004). This model cannot explain any of the nonlinear behaviors studied here. Other models attempted to explain the nonlinear behaviors one by one. Contrast saturation has long been modeled with a divisive function similar to Equation 2.1 (e.g. Derrington and Lennie, 1984; Sclar et al., 1990; Felisberti and Derrington, 1999; Kremers et al., 2001); this model, however, does not account for the role of stimulus size in determining contrast saturation. Size tuning has recently been modeled with a subtractive interaction between two Gaussians, which in our model correspond loosely to the receptive field and the suppressive field (Solomon et al., 2002; Nolt et al., 2004; Ozeki et al., 2004); such a model, however, does not capture the effect of contrast on size tuning – it fits the data only when its parameters are allowed to vary with contrast – and misses the fact that suppression is divisive. A divisive model has been recently fitted to explain phenomena of surround suppression, whereby responses to contrast in a disk are reduced by contrast in a

surrounding annulus (Solomon et al., 2002); this divisive model, however, was not shown to explain other nonlinearities.

2.3.2 Limitations of the study

A limitation of our study is that we did not consider the nonlinear spatial summation seen both in retina and in LGN. We have modeled the component of the responses oscillating at the stimulus frequency. This component accounts for most of the variance in the responses of X cells. Y cells, however, also respond to moving gratings with an elevated mean firing rate (Enroth-Cugell and Robson, 1966; Hochstein and Shapley, 1976; Derrington et al., 1979; Demb et al., 1999). We did not consider the mean rate of responses, and therefore we did not test the suggestion that the subunits of Y cells are one and the same with the blank of filters in the suppressive field (Shapley and Victor, 1981; Victor, 1987).

Another example of nonlinear summation that we have not considered is the "shift effect" (also known as the "periphery effect"), whereby a neuron's firing rate can be increased by stimulus moving well outside the receptive field (McIlwain and Creutzfeldt, 1964; Levick et al., 1965; Ikeda and Wright, 1972; Kruger and Fischer, 1973; Derrington et al., 1979; Passaglia et al., 2001). Our model cannot explain this increase. Remote stimulation, however, can also induce suppressive effects that have been related to the shift effect (Felisberti and Derrington, 1999, 2001; Passaglia et al., 2001; Girardin et al., 2002). These effects resemble the effects of the suppressive field in many aspects. First, remote stimulation has a divisive effect on the responses to a central target (Felisberti and Derrington, 1999). Second, it affects the contrast gain of the neurons (Felisberti and Derrington, 1999; Girardin et al., 2002). Third, suppression can be induced by stimuli as distant as ~10 deg away from the receptive field center. This finding might be consistent with our finding that the suppressive field extends over a large region of visual field.

A major limitation of our work is that we have not considered the time course of the responses. Contrast gain control not only affects the gain of the responses but also their temporal dynamics (Shapley and Victor, 1978b; Shapley and Victor, 1981). Increasing contrast reduces the integration time of the neurons affecting the response's phase and selectivity for temporal frequency. The model does not capture these effects: in fact, we intentionally simplified the Shapley-Victor model so that it would predict only the

amplitude of the responses. Our model only works for stimuli of fixed temporal frequency for which the effects of gain control are constant.

Moreover, we have not considered the dynamics of contrast gain control, i.e. how a change in contrast at a given time can influence gain at a later time. We cannot study these dynamics because in our stimuli contrast was fixed through time. In particular, we cannot distinguish fast components of contrast gain control (Victor, 1987; Baccus and Meister, 2002; Zaghloul et al.) from slower components of contrast adaptation (Chander and Chichilnisky, 2001; Kim and Rieke, 2001; Baccus and Meister, 2002; Solomon et al., 2004b). We suspect that the effects are fast because slow adaptation in LGN responses is clear only for high-frequency stimuli (Solomon et al., 2004b), and is weak at lower frequencies (Ohzawa et al., 1985; Shou et al., 1996; Sanchez-Vives et al., 2000; Yang et al., 2003; Solomon et al., 2004b). Moreover, our stimuli were brief (typically 1-2 s, always <4 s), presented in randomized order, and separated by ~1 s of blank screen; these properties would minimize the effects of slow adaptation, which take seconds to develop and are strongest when low contrast stimuli are preceded by long periods of strong stimulation (Solomon et al., 2004b). Finally, in recent experiments we show that a step increase in contrast is met by an immediate reduction in gain (Mante et al., 2005b).

2.3.3 Implications for V1 responses

Our formulation of retinal contrast gain control in terms of division reveals close similarities with models of gain control or normalization in primary visual cortex (Albrecht and Geisler, 1991; Heeger, 1992; Carandini et al., 1997; Carandini et al., 2002). These models predict expressions similar to Equation 2.1, and thus similar divisive effects as those seen here. In fact, neurons in primary visual cortex exhibit contrast saturation (Albrecht and Hamilton, 1982; Sclar et al., 1990), masking (Bonds, 1989; Freeman et al., 2002) and size tuning (Sceniak et al., 1999; Sceniak et al., 2001; Cavanaugh et al., 2002a). The exact contribution of LGN suppression to these phenomena is not clear.

The fact that a model with similar structure can explain gain control in retina, in LGN, and in cortex is very good news. It paves the way for a long-due assessment of the relative roles of retina, LGN, and cortex to phenomena of suppression, gain control, and normalization seen in primary visual cortex. If these phenomena could be fitted with the same model in the different stages of visual processing, then the contribution of each stage

can be described by the parameters of the model measured at that stage (Sclar et al., 1990; Jones et al., 2000).

2.3.4 Function of the suppressive field

What use might the visual system make of divisive suppression? One possible use is to map the wide range of contrasts encountered in the natural environment onto the dynamic range of the neurons (Heeger, 1992). Suppression may also play a role in encoding of visual information by populations of neurons: a divisive gain control model applied to responses of primary visual cortex was shown to increase independence of the responses across neurons (Schwartz and Simoncelli, 2001; Valerio and Navarro, 2003).

We speculate that the suppressive field could also play a role in the joint encoding of spatial frequency and spatial position, which are subject to a stringent tradeoff (Daugman, 1985). In LGN, selectivity for spatial frequency results from the antagonistic effect of the receptive field center and surround. This selectivity comes, however, at the cost of lower spatial resolution because of the large size of the receptive field surround. The suppressive field counteracts this effect. The fall-off of responses in size tuning curves typically occurs when stimuli extend beyond the receptive field center (Figure 2.4**B**) while a linear receptive field model predicts (Figure 2.4**B**, *curves*) that responses to such stimuli should grow steadily with stimulus size. Therefore the suppressive field reduces the response of the receptive field surround to stimuli covering the receptive field center and thereby improves spatial localization. This improvement in encoding of spatial position would come at no cost because the suppressive field barely affects the tuning of a neuron for spatial frequency (Figure 2.4**F**).

2.4 Methods

Methods for physiological recordings, stimulation and data analysis are described in General Methods. We report here on recordings made in LGN of 8 adult anesthetized, paralyzed cats. We recorded from 34 well-isolated neurons (28 of type X and 6 of type Y, 19 ON and 15 OFF). These neurons were held long enough to perform a series of 7 experiments, which involved more than 180 stimuli. These experiments generally took more than two hours per neuron to complete. These cells have receptive fields with eccentricities

ranging from 2° to 45°, with an average value of 13.4±8.9° (s.d., N=34). We discuss results from 5 experiments in this chapter and address the remaining experiments in Chapter 3.

2.4.1 Model

The model contains two pathways, one linear and one nonlinear (Figure 2.2).

The linear pathway involves a classical center-surround receptive field, a difference-of-Gaussians, $RF = G_{ctr} - k_{srd} G_{srd}$, where G_{ctr}, G_{srd} are 2d Gaussian densities of width σ_{ctr}, σ_{srd} *and* k_{srd} denotes the relative strength of the surround. The receptive field is convolved with the stimulus $S(x,y,t)$ to produce a linear response $L(t) = [S * RF](0,0,t)$. For our stimuli, which consist of the sum of a *test* and a *mask* drifting gratings, $L(t)$ is the sum of the responses to the individual gratings:

$$L(t) = L_{test}(t) + \alpha_{mask} L_{mask}(t), \qquad 2.2$$

where α_{mask} denotes the relative effectiveness of the mask in driving the receptive field.

The nonlinear pathway involves a suppressive field which computes the standard deviation of local luminance c_{local}. Assuming that the stimulus $S(x,y,t)$ has a mean of zero, local contrast follows

$$c_{\text{local}} = \sqrt{\iiint \left(\hat{S}(x,y,t)\right)^2 G_{SF}(x,y)\,dx\,dy\,dt}, \qquad 2.3$$

where G_{SF} denotes the suppressive field, a Gaussian of width σ_{SF}, and $S'(x,y,t)$ denotes the stimulus processed through a bank of filters, $[S * H](x,y,t)$. We model each of these filters as a difference-of-Gaussians $H = G_u - k_d G_d$.

The suppressive field controls neural gain by dividing the output of the receptive field. The result is

$$V(t) = V_{\max} \frac{L(t)}{c_{50} + c_{\text{local}}}, \qquad 2.4$$

where c_{50} determines the strength of the suppressive field and V_{max} captures the overall responsiveness of the neuron.

Firing rate is a rectified version of V, with threshold V_0:

$$R(t) = \lfloor V(t) - V_0 \rfloor_+ . \qquad 2.5$$

2.4.2 Model Characterization

To estimate model parameters we fitted equations 2.2-2.5 to a sequence of four experiments, each constraining one or more parameters. The first experiment constrained the parameters of the receptive field (σ_{ctr}, σ_{srd} and k_{srd}), and involved drifting gratings varying in spatial frequency. Gratings had 50% contrast, optimal temporal frequency, >20 deg in diameter and one of >14 logarithmically spaced spatial frequencies. The subsequent three experiments constrained the parameters of the suppressive field (c_{50}, σ_{SF}, σ_u, σ_d, k_d, and α_{mask}), and involved sums of a *test* grating and a *mask* grating, where we varied mask contrast, mask diameter, and mask spatial frequency. Test diameter, spatial frequency and temporal frequency were optimized cell by cell to elicit maximal response. Unless varied, the attributes of the mask were also optimal for the cell. Each experiment included test gratings presented alone at >6 contrasts and sums of a 50% contrast test and a mask of variable attributes. Parameter estimates obtained when fitting one experiment were held fixed in fits to subsequent experiments. We repeated this sequence of fits until parameters changed by less than 1% since the previous sequence iteration. The result was a single set of parameters used to predict responses to all experiments.

Parameters V_{max} *and* V_0 were allowed to vary across experiments to account for slow changes in neural responsiveness and spontaneous activity. During the course of a recording session with a given neuron (which lasted on average 289 ± 72 minutes, s.d., N=34), the spontaneous firing rate varied over 18 ± 11 spikes/s, with variations > 25 spikes/s in 9 of 34 neurons. Similarly, the overall responsiveness varied substantially over time. The harmonic response to a 50% contrast stimulus varied on average by 25 ± 16 spikes/s, with variations of > 40 spikes/s in 4 of 34 neurons. These variations in spontaneous firing rate and responsiveness were not statistically correlated (p>0.2). We tracked them by letting V_{max} and V_0 vary from one experiment to another. In each experiment, we fixed V_0 to minus the mean spontaneous firing rate, which we obtained from the responses to blank stimuli, and we chose V_{max} to minimize square error between model and data. Had we fixed V_{max} and V_0 across experiments, the quality of predictions would have been compromised in many neurons.

The logic of the masking experiments is that the test elicits strong responses, and the mask probes how the suppressive field influences these responses (Figure 2.7). In general,

both test (Figure 2.7AB) and mask (Figure 2.7CD) elicit responses when presented alone. These responses oscillate at the drift frequency, and are well predicted by a simple model consisting of a receptive field followed by rectification (Figure 2.7A-D *dashed curves*). We chose the drift rates of test and mask to be incommensurate (e.g. 7.8 Hz and 12.5 Hz) so that even when test and mask are superimposed (Figure 2.7EF) we can distinguish a *test response* (the component of the response at the test frequency) and a *mask response* (the component at the mask frequency). The average temporal frequencies of test and mask were 7.3±2.5 and 12.6±2.9 Hz (N=34).

In the presence of the mask, the test response is smaller than predicted by the receptive field and rectification (Figure 2.7EF *dotted curves*). This *masking* effect is explained by the full model because adding the mask increases the output of the suppressive field. From the response suppression caused by the mask we deduce the degree to which it drives the suppressive field.

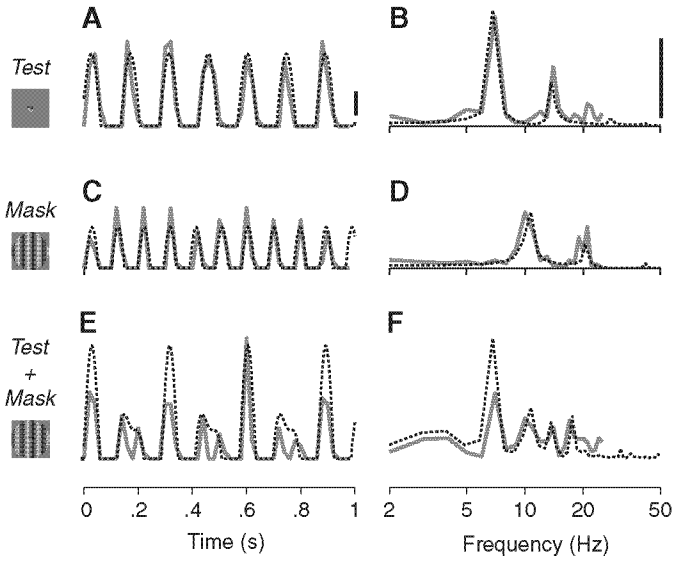

Figure 2.7. Masking experiments. Stimuli are sums of a test and a mask grating with incommensurate temporal frequencies. **A**: Firing rate responses of an idealized neuron (lacking gain control) as function of time. Model consists of a receptive field followed by a rectification. *Red curve* indicates response to test alone. *Blue curve* indicates response to mask alone. *Dashed curve* indicates response to test + mask. **B**: Same, but showing the amplitude spectrum of responses. Highest two peaks are harmonic responses of test and mask, which are unaffected by summing test and mask. **C**: *Solid curves* are firing responses of a real LGN neuron. **D**: Same, but showing the amplitude spectrum of responses. Harmonic responses in test + mask condition (*black curve*) have lower amplitudes than in the remaining conditions (*shaded areas*) indicative of gain control.

We obtained one parameter of the receptive field α_{mask}, from the mask responses, and all parameters of the suppressive field from the test responses. We estimated the strength of the suppressive field c_{50} by varying mask contrast over >6 logarithmically spaced values. We estimated the size of the suppressive field σ_{SF} by varying mask diameter over >10 logarithmically spaced values. We estimated the parameters that describe the filter bank (σ_u, σ_d and k_d) by varying mask spatial frequency over >10 logarithmically spaced values. Finally we used the mask responses to constrain α_{mask}, the relative effectiveness of the mask in driving the receptive field. We estimate α_{mask} from data in which mask contrast is varied and use the same value for the remaining experiments.

Stimuli sometimes fell off the exact receptive field center of the neurons. We estimated this stimulus offset Δd by fitting the model onto responses to stimuli of different diameters. This estimate held in the remaining experiments.

2.4.3 Testing the Model

We tested model predictions with optimal drifting gratings varying in contrast and diameter (Figure 2.5). This experiment included > 50 stimuli covering a full matrix of >6 contrasts and >8 diameters distributed on logarithmic scale. Parameters V_{max}, V_0 were left free to account for changes in neural responsiveness. The remaining parameters were held fixed to their previously estimated values.

To investigate the model's ability to predict the effects of contrast on the size tuning curves we fitted a difference-of-Gaussians to the tuning curves at each contrast. We performed this analysis on the data and on the model predictions. From these fits we calculated the preferred stimulus diameter and the degree of size tuning expressed as a fraction of the peak responses.

To investigate the model's ability to predict the effects of stimulus size on contrast saturation, we fitted a power-law to the contrast-response curves at each diameter. The fits thus obtained were excellent. The power-law exponent β captures the degree of linearity of the growth of responses with contrast: a value close to 0 indicates strong saturation, and a value close to 1 indicates linearity.

To measure the precision of model predictions in individual stimulus conditions we calculated z-scores of deviations between measured and predicted responses. The z-score of

deviations is given by $z = (\bar{r} - m)/\hat{\nu}$ where $\hat{\nu}$ is a robust estimate of the standard deviation of responses calculated across trials. To obtain $\hat{\nu}$ we fitted a power-law relating standard deviations ν to the mean $\bar{r}$ and used the resulting predictions as estimate of standard deviation.

Chapter 3

Origins of the suppressive field

The responses of neurons in lateral geniculate nucleus (LGN) are shaped by powerful suppressive phenomena. These phenomena have been ascribed to a variety of mechanisms. We have explained them with a gain control model where the output of the receptive field is divided by the output of a suppressive field. To constrain the possible origins of the suppressive field we measured its spatial extent and preferences for stimulus attributes. The suppressive field is no larger than the receptive field surround. It is not selective for stimulus orientation, and it responds to a wide range of frequencies, including very low spatial frequencies and high temporal frequencies. These properties suggest that the suppressive signals originate in retina or within LGN, but not in feedback from cortex. Together our results provide a thorough characterization of gain control in LGN.

3.1 Introduction

The visual responses of neurons in lateral geniculate nucleus (LGN) reveal suppressive phenomena that cannot be explained by the classical center-surround receptive field. First, responses saturate when stimulus contrast increases (contrast saturation, Derrington and Lennie, 1984; Chino and Kaplan, 1988; Sclar et al., 1990; Kremers et al., 2001; Alitto and Usrey, 2004). Second, responses to a stimulus are reduced by superimposition of another (masking, Cudeiro and Sillito, 1996; Felisberti and Derrington, 1999; Freeman et al., 2002;

Solomon et al., 2002). Third, responses decrease when stimulus size is increased beyond an optimal value (size tuning, Cleland et al., 1983a; Murphy and Sillito, 1987; Jones and Sillito, 1991; Solomon et al., 2002; Ozeki et al., 2004). These three phenomena are suppressive, because responses are smaller than would be expected from the classical center-surround receptive field alone.

There is little agreement as to the physiological basis of these suppressive phenomena. Some have ascribed suppression to inhibitory signals within thalamus (Hubel and Wiesel, 1961; Levick et al., 1972; Cleland et al., 1983a; Sclar, 1987; Felisberti and Derrington, 1999; Blitz and Regehr, 2005). Others attributed it to cortical feedback (Murphy and Sillito, 1987; Przybyszewski et al., 2000; Webb et al., 2002; Alitto and Usrey, 2003).

In the previous chapter we have shown that the suppressive phenomena can be explained by a simple mechanism of gain control. We devised a model that consists of a nonlinear suppressive field whose output divides the output of the classical center-surround receptive field.

In this chapter we probe the possible origins of the suppressive phenomena through the properties of the suppressive field. We measure its spatial extent comparing it to that of the receptive field. We study its visual preferences and ask whether these preferences are more consistent with its originating in retina, in thalamus, or in cortex.

3.2 Results

Having validated the model with a number of tests in the previous chapter, we now turn to its key component, the suppressive field. We study the suppressive field's preferences for visual stimulus attributes, and ask whether the suppressive field extends well beyond the receptive field. These measurements constrain the possible sources of suppressive signals, and the possible functional roles of suppression.

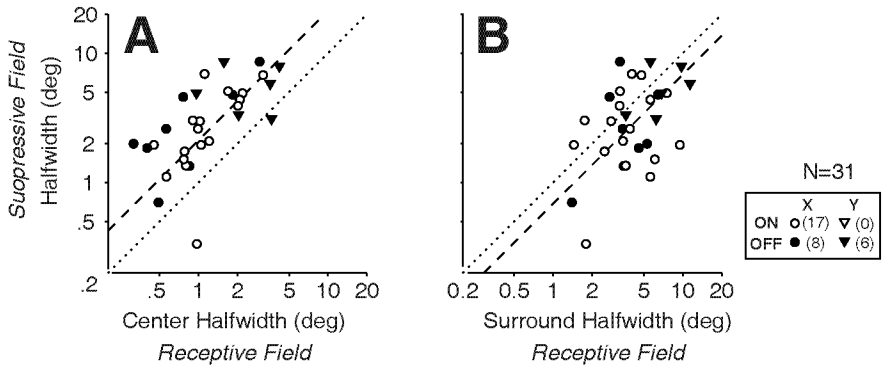

Figure 3.1. Spatial extent of receptive field and suppressive field obtained from model fits. Extent is measured as half-width (the width at half of the maximum output). **A**: Comparison of extents of suppressive field and receptive field center. *Dashed line* indicates linear regression. **B**: Comparison of extents of suppressive field and receptive field surround.

3.2.1 Spatial Extent

The suppressive field is much larger than the receptive field center, but generally comparable in size to the receptive field surround. We estimate the spatial extent of receptive field and suppressive field from the model parameters (obtained in the fit of Figure 2.3**A** and Figure 2.4**B**) for the 31/34 cells which showed a measurable receptive field surround (those with surround gain k_s>0.2, see Experimental Procedures). In agreement with Solomon et al. (2002) we find that suppressive field is larger than the receptive field center (Figure 3.1**A**). On average, the suppressive field is 215±24 % (bootstrap estimates) larger. The size of the suppressive field increases with that of the receptive field center, with a correlation (r=0.67±0.12, bootstrap estimates) comparable to the correlation between the sizes of receptive field center and surround (r=0.53±0.16). The sizes of suppressive field and receptive field surround were also correlated (r=0.38±0.16). Perhaps surprisingly, the size of the suppressive field is on average only 69±8 % of that of the receptive field surround (Figure 3.1**B**). The suppressive field is thus generally not larger than the surround of the classical receptive field.

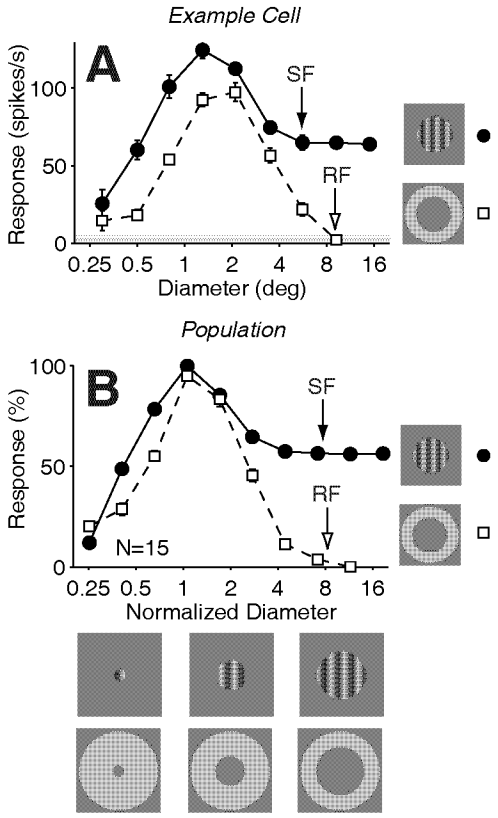

Figure 3.2. Control measurements of spatial extent of receptive field surround and suppressive field. *Closed symbols* indicate responses to *disks* containing drifting gratings. *Open symbols* indicate responses to *annuli* containing a uniform field whose contrast is modulated in time. *Arrows* indicate estimated extents of receptive field surround and suppressive field. **A**: Responses of an example neuron (cell 44.4.1). Error bars indicate ± 1 s.d. **B**: Responses averaged across cells. We normalized the abscissa and the ordinate of each curve so that the size tuning curves peak at 1 with a value of 1. Error bars (invisible because smaller than symbols) indicate ± 1 s.e. (N=15).

To confirm these results, we collected data that measure the extent of receptive field surround and suppressive field, without the need for model fits (Figure 3.2). First, we recorded from 15 cells in one additional animal responses to contrast-modulated circular annuli, and we varied annulus inner diameter (Figure 3.2**A**, *open symbols*). The annulus tuning curve thus obtained reflects the center-surround organization of the receptive field (Kremers et al., 2004). The inner diameter where this curve reaches the floor (within ± 1 s.d. of the response to a blank stimulus) provides an estimate of the size of the receptive field surround (Figure 3.2**A**, *open arrow*). Second, we recorded responses to optimal drifting gratings, and varied the diameter of their disk window (Figure 3.2**A**, *closed symbols*). The diameter for which these curves reach a plateau (within 1 s.d. of the response to the largest

stimulus) provides an estimate of the size of the suppressive field (Figure 3.2**A**, *closed arrow*). For the cell in Figure 3.2**A**, the suppressive field is contained within the bounds of the receptive field and does not extend beyond it. Similar conclusions were reached in the remaining cells (Figure 3.2**B**). On average, the estimated receptive field and suppressive field sizes were undistinguishable (8.2 ± 0.9 and 7.9 ± 1.9 times larger than the optimal disk diameter, s.e.m., N=15). Receptive field and suppressive field, therefore, have nearly identical extent.

A consequence of these results is that a commonly employed method to estimate receptive field size consistently underestimates true size. Receptive field size is often defined as the size of the high-contrast disk eliciting maximal response (Felisberti and Derrington, 1999; Jones et al., 2000; Sillito and Jones, 2002; Solomon et al., 2002; Webb et al., 2002; Alitto and Usrey, 2003; Sun et al., 2004). This measure underestimates true size by more than a factor of eight (8.2 ± 0.9, Figure 3.2**B**).

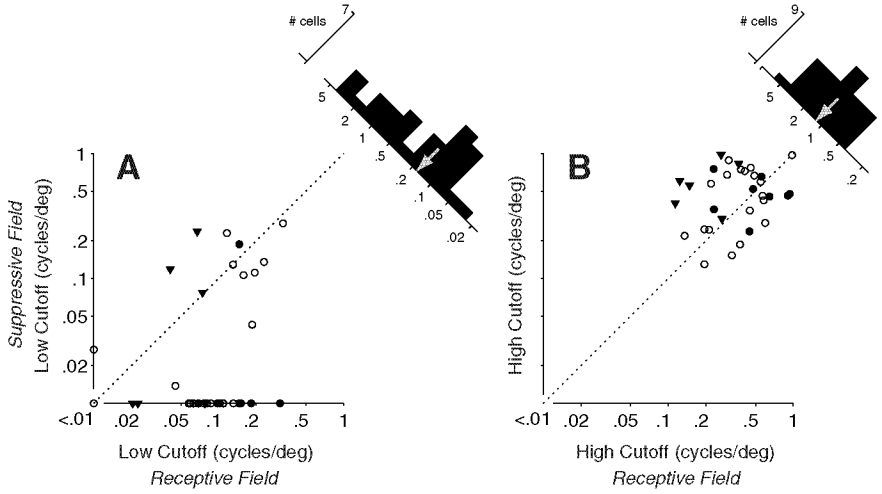

Figure 3.3. Selectivity for spatial frequency of receptive field and suppressive field. Cutoffs are frequencies at which outputs of receptive field (L in Figure 2.2) and of suppressive field (c_{local} in Figure 2.2) reach their half-maximum. **A**: Scatter plot shows low cutoffs of receptive field and suppressive field for population of cells (N=34). Data points that lie on abscissa and ordinates indicate cutoffs below the lowest frequency tested. Histogram depicts distribution of cutoffs ratios. Arrow indicates median of ratios. **B**: Same, for high cutoffs.

3.2.2 Selectivity for Spatial Frequency

The suppressive field is more broadly selective for spatial frequency than the receptive field (Figure 3.3). We infer the selectivity for spatial frequency from the model parameters (obtained in the fit of Figure 2.3**A** and Figure 2.4**C**). The suppressive field responds well to low spatial frequencies, frequencies that elicit only weak responses in the receptive field (Figure 3.3**A**). The suppressive field's low cutoffs mostly lie below 0.01 cycles/deg, i.e. the lowest spatial frequency tested. By comparison, the low cutoffs of the receptive field typically lie in the range of 0.05 to 0.2 cycles/deg (Maffei and Fiorentini, 1973; Dawis et al., 1982; Cheng et al., 1995). In 27/34 neurons, the suppressive field has low cutoffs below those of the receptive field. On average, the suppressive field's low cutoff is 0.23±0.12 times that of the receptive field (median, bootstrap estimates). The suppressive field responds to a broad range of frequencies, which often extends beyond the resolution limit of the receptive field (Figure 3.3**B**). In 22/34 neurons the suppressive has high cutoff frequencies above those of the receptive field. On average, its high cutoff is 1.27±0.19 times that of the receptive field.

Being broadly tuned, the suppressive field should contribute only minimally to the overall spatial frequency tuning of the neurons. This prediction is borne out by our data, as the spatial frequency tuning of our neurons is largely independent of contrast (compare Figure 2.3**A** with Figure 2.4**F**). We measured spatial frequency tuning curves for stimuli of low (<25%) and high contrast (100%) and calculated the high frequency cut-off of each curve. In all cells the cut-off frequencies did not change with contrast (N=7).

3.2.3 Selectivity for Orientation

The suppressive field has no preferences for orientation: its bias for orientation is much weaker than that of the receptive field (Figure 3.4). To measure the preferences of suppressive field and receptive field for orientation, we recorded responses to sums of test and mask gratings, and varied mask orientation. We obtain these measurements in 30 the 34 neurons in the main sample. We estimate the orientation preference of the suppressive field (Figure 3.4**B**) from the test responses (the component of the responses at the temporal frequency of the test). We subtract the test responses measured in the presence of the mask (Figure 3.4**A**, *symbols*) from the test response measured without mask (Figure 3.4**A**, *dashed*

line). The resulting profiles (Figure 3.4**B**, *thin lines*) reflect the orientation tuning of the suppressive field. We find this tuning to be weak or nonexistent, reaching significance (ANOVA, $p<0.05$) in only 7/30 neurons (Figure 3.4**C**). By comparison the receptive field often shows consistent biases for orientation (Figure 3.4**D-F**) (Vidyasagar and Urbas, 1982; Soodak et al., 1987; Shou and Leventhal, 1989; Zhou et al., 1995; Sun et al., 2004). We measure the orientation tuning of the receptive field from mask responses (the component of the responses at the frequency of the mask, Figure 3.4**D**, *symbols*). The profile of receptive field responses thus obtained (Figure 3.4**E**, *thin lines*) show clear orientation biases (ANOVA, $p<0.05$) in 20/30 neurons (Figure 3.4**F**). This bias had no consistent relationship with that seen occasionally in the suppressive field: aligning responses to center around 0 deg the orientation that gave the peak receptive field response (Figure 3.4**E**, *thin lines*) and averaging across neurons yields a curve that is essentially flat (Figure 3.4**B**, *thick line*). The suppressive field, therefore, is largely insensitive to stimulus orientation.

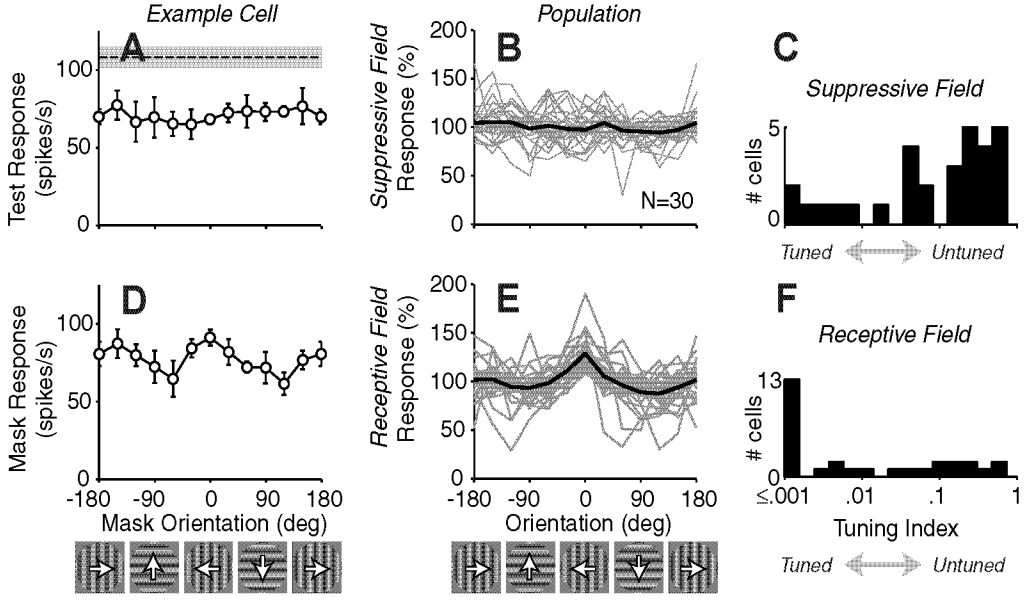

Figure 3.4. Selectivity for orientation of receptive field and suppressive field. Stimuli are sums of test and mask gratings with incommensurate temporal frequencies (see details in Figure 2.4). **A**: Test responses (*symbols*) as function of mask orientation. *Dashed line* indicates response to test alone; *shaded area* denotes ± 1 s.d. **B**: *Thin lines* are estimated suppressive field responses as function of orientation, in population of neurons (N=30). *Thick line* is average across cells. Responses are normalized to a mean of 100%, and aligned so that maximum receptive field response (**E**) centered on 0 deg. **C**: Distribution of p-values for the hypothesis that there is no orientation tuning (low p-values are associated with clear tuning). **D:** Same as **A** for mask responses. **E**, **F**: Same

as **B**, **C** for estimated receptive field responses. For the example cell (33.3.4), test attributes were 0.12 cycles/deg, 10.4 Hz, 2.3 deg and mask attributes were .12 cycles/deg, 15.6 Hz and 21 deg.

3.2.4 Selectivity for Temporal Frequency

The suppressive field is selective for temporal frequency, and prefers frequencies of 10-20 Hz or higher (Figure 3.5). To measure the selectivity for temporal frequency, we recorded responses to sums of test and mask gratings, and varied mask temporal frequency. We obtained these measurements in 40 neurons (which include the 34 of our main sample). We estimate the selectivity of the suppressive field by subtracting the test responses measured with mask (Figure 3.5**A**, *symbols*) from the test response measured without mask (Figure 3.5**A**, *dashed line*), and averaging across cells. The resulting profile (Figure 3.5**B**, *symbols*) shows weak suppression at low frequencies and strong suppression at high frequencies. The selectivity of the suppressive field resembles the that of the receptive field as estimated from the mask responses (Figure 3.5**C-D**, *symbols*), or from single drifting gratings varying in temporal frequency (Figure 3.5**B-D**, *gray curves*) (see also Dawis et al., 1984; Saul and Humphrey, 1990; Kilavik et al., 2003; Alitto and Usrey, 2004). Because test and mask stimuli were interleaved, we could not explore frequencies above 21 Hz with these stimuli (1/6 of 124 Hz, the refresh rate of the monitor) so we do not know how the suppressive field responds to higher frequencies.

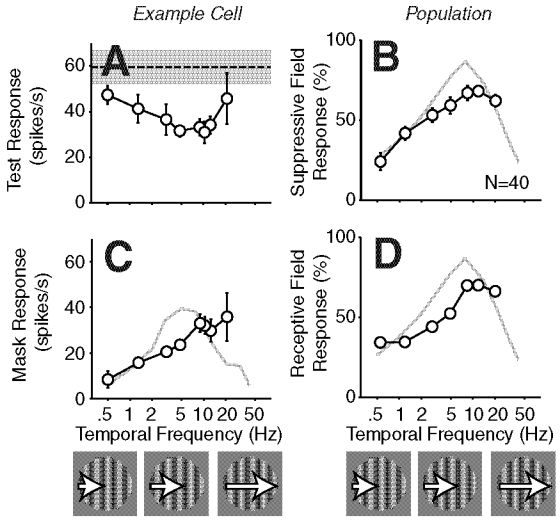

Figure 3.5. Selectivity for temporal frequency of receptive field (RF) and suppressive field (SF). *Symbols* are responses to sums of test and mask gratings in which mask temporal frequency is varied (details as in Figure

2.4). *Gray lines* are control: responses to single gratings varying in temporal frequency. **A**: Test responses of example neuron. *Dashed line* indicates response to test alone; shaded area denote ± 1 s.d. **B**: Estimated suppressive field responses, normalized to a peak of 100% and averaged across cells. Error bars indicate s.e. (N=40). **C**: Mask responses of example neuron. **D**: Estimated receptive field responses, normalized to a peak of 100% and averaged across neurons. For the example cell (39.3.4), test attributes were 0.6 cycles/deg, 7.8 Hz, 2 deg and mask attributes were .6 cycles/deg, and 20 deg.

3.3 Discussion

We characterized the visual preferences of the suppressive field and found that it is generally (1) not larger than the receptive field surround, (2) broadly tuned for spatial frequency, (3) not selective for orientation, and (4) broadly selective for high temporal frequencies.

3.3.1 Comparison with previous measurements

Our finding that suppression is confined to the receptive field surround appears to be at odds with previous anatomical and physiological studies which describe suppression as originating from beyond the classical receptive field (Sanderson et al., 1971; Levick et al., 1972; Sherman and Koch, 1986; Murphy and Sillito, 1987; Cucchiaro et al., 1988; Uhlrich et al., 1991; Cudeiro and Sillito, 1996; Felisberti and Derrington, 1999; Przybyszewski et al., 2000; Kaplan and Benardete, 2001; Sillito and Jones, 2002; Solomon et al., 2002; Webb et al., 2002; Alitto and Usrey, 2003). Anatomical studies have linked thalamic inhibitory circuits to these effects. Our stimuli might possibly fail to reveal suppression from beyond the receptive field. Alternatively, previous studies might simply have underestimated the actual extent of the receptive field. The receptive field surround only provides a weak drive to the cell, which could easily go unnoticed. Furthermore, receptive field size is often defined as the size of a high-contrast disk eliciting maximal response (Felisberti and Derrington, 1999; Jones et al., 2000; Sillito and Jones, 2002; Solomon et al., 2002; Webb et al., 2002; Alitto and Usrey, 2003; Sun et al., 2004), a measure that severely underestimates true size (by a factor of eight, Figure 3.2**B**).

We observed that the suppressive field is tuned much more broadly than the receptive field. These results agree with previous qualitative assessments of the selectivity of suppression (Shapley and Victor, 1979; Cudeiro and Sillito, 1996; Sun et al., 2004). A

consequence of this observation is that the spatial frequency tuning of LGN neurons is largely independent of contrast (see discussion of Figure 9). In fact, this behavior is consistent with the data – if not with the interpretation – of a recent study of spatial frequency tuning at different contrasts (Nolt et al., 2004).

Our observation that the suppressive field has no preferences for orientation agrees with some reports (Solomon et al., 2002; Webb et al., 2002; Sun et al., 2004) but not with others (Sillito et al., 1993; Cudeiro and Sillito, 1996). The latter studies used gratings enclosed in a disk and an annulus, and found that the orientation of the grating in the annulus affected the responses to the disk. This result might be due to interactions between receptive field responses to the disk and to the annulus. The annulus in these studies was placed just outside the diameter that yields the maximal response for high-contrast disk stimuli. Our data (Figure 8**B**) indicates that this diameter is about 8 times smaller than the actual extent of the receptive field surround. It is thus reasonable to assume that the annulus stimulated the receptive field surround, and elicited some response. This response would oscillate at the same frequency as the response to the disk. Depending on the details of the stimulus and of the receptive field, the two responses could interact destructively or constructively. The result could be an apparent effect of mask orientation.

Finally, our assessment of the temporal frequency preferences of the suppressive field complements previous quantitative studies of gain control (Shapley and Victor, 1978b; Shapley and Victor, 1981; Benardete and Kaplan, 1999). These studies measured how gain control affects responses to stimuli of different temporal frequencies. They found that gain control mostly affects responses to low frequencies, and barely affects responses to high frequencies. Our experiments elucidate a different aspect of gain control, namely the selectivity of the signal driving these changes in dynamics. Consistent with an early anecdotal report (Shapley and Victor, 1979), we found this signal to be strong at high temporal frequencies but weak at low frequencies.

3.3.2 Origins of the suppressive field

The suppressive field is distinct from the receptive field in many aspects. First, it is responsible for suppressive phenomena that can be explained by neither the receptive field alone nor the rectification of the spike threshold (Figure 2.1, Figure 2.3 and Figure 2.4, *dashed lines*). Second, it responds well not only to high spatial frequencies, but also to low

spatial frequencies, which elicit poor responses in the receptive field (Figure 3.3**A**). Third, it computes the standard deviation of local luminance thus responding equally to light increments and decrements (3.4).

The suppressive signals affecting LGN neurons have been ascribed to one or more of three mechanisms: (1) Mechanisms of suppression or gain control operating in retina (Shapley and Victor, 1978b; Shapley and Victor, 1981; Cleland et al., 1983a; Cleland and Lee, 1985; Victor, 1987; Nolt et al., 2004; Solomon et al., 2004a); (2) Inhibitory circuitry in thalamus (Hubel and Wiesel, 1961; McIlwain and Creutzfeldt, 1967; Singer and Creutzfeldt, 1970; Levick et al., 1972; Singer et al., 1972; Cleland et al., 1983b; Sherman and Koch, 1986; Kaplan et al., 1987; Funke and Eysel, 1998); (3) Negative feedback from visual cortex (Ahlsén and Lindström, 1983; Sherman and Koch, 1986; Murphy and Sillito, 1987; Sillito et al., 1993; Cudeiro and Sillito, 1996; Sillito and Jones, 2002; Webb et al., 2002; Worgotter et al., 2002; Alitto and Usrey, 2003).

The visual preferences of the suppressive field are entirely consistent with an origin in retina: just as the suppressive field, retinal ganglion cells respond as well to low spatial frequencies as to high spatial frequencies, follow high temporal frequencies, and are barely tuned for stimulus orientation.

Our measurements, however, cannot exclude a contribution from thalamic inhibition. There is evidence that gain control is stronger in LGN than in retinal ganglion cells (Kaplan et al., 1987; Cheng et al., 1995). LGN cells receive inhibition from two sources (Dubin and Cleland, 1977; Ahlsén et al., 1982; Sherman and Koch, 1986): feedforward from local interneurons and feedback from the perigeniculate nucleus (PGN). Interneurons have visual properties that closely resemble the properties of relay cells (Dubin and Cleland, 1977). To explain its broad tuning, therefore, the suppressive field should pool the outputs of many interneurons responding to a broad range of spatial frequencies. A more likely source of suppression lies in PGN neurons, whose visual preferences resemble those of the suppressive field. First, PGN neurons have large receptive fields (Sanderson, 1971; Uhlrich et al., 1991; Wrobel and Bekisz, 1994; Funke and Eysel, 1998), certainly larger than the receptive field centers of LGN neurons; Second, PGN neurons respond well to low spatial frequencies (Price and Morgan, 1987; Xue et al., 1988; Murphy et al., 1994; Wrobel and Bekisz, 1994). Third, they respond equally to both light increments and decrements (ON/OFF responses, Sanderson, 1971; Dubin and Cleland, 1977; So and Shapley, 1981; Xue

et al., 1988; Uhlrich et al., 1991; Wrobel and Bekisz, 1994; Funke and Eysel, 1998); Fourth, they show poor selectivity for stimulus size (Jones and Sillito, 1994). The last property is consistent with the observation that suppression is maintained for large diameters (Figure 2.4**B**). To test for a role of the PGN one may want to perform experiments with a test grating in one eye and a mask grating in the other eye. Indeed, PGN neurons can be driven from both eyes (Sanderson, 1971; Ahlsén et al., 1983; Xue et al., 1988; Uhlrich et al., 1991; Funke and Eysel, 1998), and have been linked to binocular suppressive effects in LGN (Sanderson et al., 1971; Pape and Eysel, 1986; Wang et al., 1994; Funke and Eysel, 1998).

Finally, the visual preferences of the suppressive field are mostly inconsistent with a cortical origin. The suppressive field has the same extent as the receptive field surround, while V1 neurons have larger receptive fields (Jones et al., 2000). In addition, the suppressive field responds strongly to low spatial frequencies (Figure 3.3) and high temporal frequencies (Figure 3.5), which elicit little response in V1 (Ikeda and Wright, 1975; Movshon et al., 1978; Saul and Humphrey, 1992; DeAngelis et al., 1993a; Freeman et al., 2002).

Cortical feedback, however, does affect LGN responses. Cortical ablation or inactivation modifies the gain of the neurons (Przybyszewski et al., 2000; Webb et al., 2002) and reduces size tuning (Murphy and Sillito, 1987; Alitto et al., 2002). These effects might possibly be explained if cortex modulated the gain of the suppressive field, for example by controlling the value of the parameter c_{50}, which we have considered to be constant (Figure 2).

3.4 Methods

Methods for physiological recordings, stimulation and data analysis are described in General Methods. We report on the same 34 LGN neurons presented in Chapter 2 (28 of type X and 6 of type Y, 19 ON and 15 OFF). The model and the experiments used to constrain the model were described in Chapter 2. We also present additional data regarding the spatial extent of the suppressive field (15 cells in two additional animals).

3.4.1 Measuring spatial extent

We estimated the spatial extents of receptive field and suppressive field from model fits. We defined spatial extent as the width of the stimulus eliciting half the maximum response

(*half-width*). For the receptive field half-width is given by $2\sigma_{srd}\sqrt{2\log 2}$. For the suppressive field half-width follows $2\sigma_{SF}\sqrt{2\log 4/3}$. These expressions differ because the Gaussians underlying the receptive field operates on stimulus intensity while that underlying the suppressive field operates on stimulus energy.

To obtain independent measure of receptive field size, we recorded the responses to contrast-modulated circular annuli and varied annulus inner diameter (Figure 3.2, *symbols*). The harmonic responses yield an annulus tuning curve that reflects the center-surround organization of the receptive field (Kilavik et al., 2003; Kremers et al., 2004). Responses are small for small inner diameter, because the annulus covers both the receptive field center and surround. Responses increase as the overlap between the annulus and the receptive field center decreases and reach a maximum when the annulus fills the receptive field surround. Responses then fall off as the annulus exits the receptive field surround. The inner diameter where responses reach the floor (within ± 1 s.d. of the response to a blank stimulus) indicates the size of the receptive field surround (Figure 3.2, *open arrow*). We compared the estimates of receptive field size obtained with this method to those obtained from the spatial frequency tuning of responses and found them consistent.

3.4.2 Measuring selectivity

We used model fits to compare the selectivity of receptive field and suppressive field for spatial frequency. We calculated the low- and high-frequency cutoffs of the outputs of receptive field and suppressive field. We defined cutoff as the frequency at which the outputs reach 50% of their maximum power (71% of the maximum amplitude).

We used masking experiments to measure the selectivity of suppressive field for temporal frequency and orientation (Figure 3.4 and Figure 3.5). Experiments were similar to those described in Chapter 2 but with masks of either >7 incommensurate temporal frequencies or 12 orientations. We estimated the selectivity of the suppressive field by subtracting the test responses measured with mask from the test response measured without mask, normalizing and averaging across cells. We estimated the preferences the receptive field by normalizing and averaging mask responses.

Chapter 4

Statistical operation of the suppressive field

Contrast gain control is thought to sum the responses of several spatially distributed filters. We have described this pooling by a suppressive field, which computes the standard deviation the filters' outputs. Here we use white noise stimuli to probe the statistical operation of the suppressive field. We measure the gain of neurons in lateral geniculate nucleus and ask how gain is affected by the luminance statistics of the stimulus. We use an orthogonal basis set to independently vary the standard deviation, the skewness and the kurtosis of the stimulus. We find that gain is strongly influenced by the standard deviation but mostly invariant to skewness and kurtosis. These results indicate that gain control estimates the true standard deviation (or root-mean-square contrast) of the stimulus. They show the suppressive field model to hold for a broad range of luminance statistics.

4.1 Introduction

Two mechanisms control the gain of early visual system responses to the prevailing image luminance statistics. Light adaptation adjusts responses to the local mean intensity (Shapley and Enroth-Cugell, 1984). Contrast gain control has been suggested to be sensitive to local root-mean-square contrast (Shapley and Victor, 1978a; Shapley and Kaplan, 1981; Victor, 1987).

In the previous chapters, we have devised a simple model for the effects of contrast gain control in lateral geniculate nucleus. In the model, gain is set by a suppressive field whose output divides the output of the receptive field. The suppressive field filters the retinal image, and pools the resulting signals. We postulated that the suppressive field computes

the standard deviation of the filters outputs. We have shown that the model predicts the responses to moving gratings and sums of gratings of different sizes and contrasts.

Grating stimuli, however, represents a weak test for the suppressive field model because they provide little control over the statistics of light intensities of the stimulus. To fully validate the model we need to test explicitly whether the suppressive field computes the standard deviation of light intensities, or whether it might also be sensitive to higher-order statistics of the stimulus. One such statistics is *skewness*, which describes the degree of asymmetry in the luminance distribution. Grating stimuli are not skewed; every deflection above the mean is accompanied by a symmetrical deflection downward. Experiments with gratings, therefore, cannot probe the sensitivity of the suppressive field to skewness. Another statistics is *kurtosis,* which describes the relative probability of stimulus deflections far away from the mean. Gratings have a fixed distribution of light intensities and thus cannot probe sensitivity of the suppressive field to kurtosis.

There is strong psychophysical evidence for a mechanism sensitive to higher-order luminance statistics of the stimulus. Chubb and colleagues (Chubb et al., 1994; Chubb et al., 2004) investigated the mechanisms underlying the discrimination of random texture stimuli (Figure 4.1). They found that humans could not only discriminate stimuli of different standard deviations (Figure 4.1A), but also stimuli of different skewness (Figure 4.1B) and stimuli of different kurtosis (Figure 4.1C). Their analysis also revealed that human performance could be accounted by postulating the existence of three unitary mechanisms: one sensitive to mean light intensity, one sensitive to the standard deviation and a last mechanism is sensitive to amount blackness in the visual scene, coined *blackshot*.

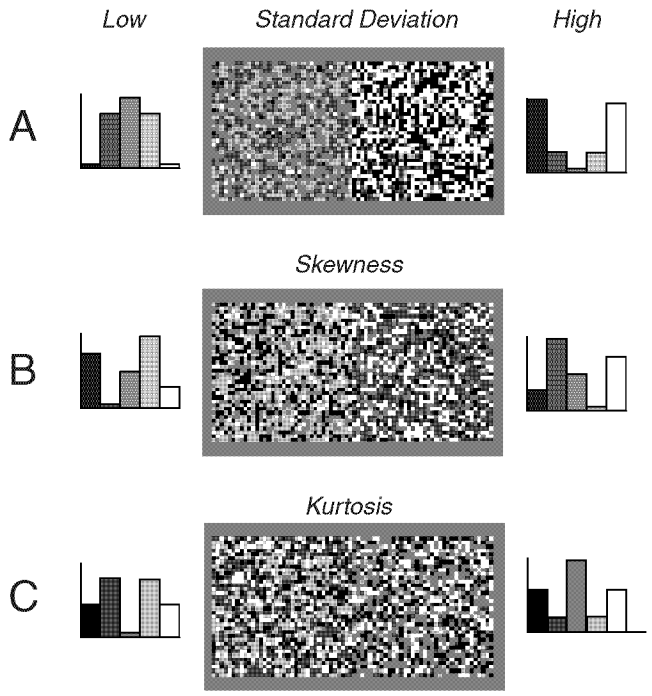

Figure 4.1. Illustration of Chubb et al (1994)'s texture discrimination task. Stimuli are static textures whose light intensities are drawn at random. Task is to detect texture boundary. **A**: Boundary defined by difference in standard deviation (root-mean-square contrast). **B**: Boundary defined by difference in skewness. **C**: Boundary defined by difference in kurtosis.

The dynamic equivalent to the textures employed by Chubb and colleagues is white noise. White noise consists of random images whose pixels are drawn independently from a fixed distribution. Recent work has employed such stimuli to study the effects of gain control in retina (Chander and Chichilnisky, 2001; Chichilnisky, 2001; Kim and Rieke, 2001; Baccus and Meister, 2002; Zaghloul et al., 2005). These studies used responses to white noise to estimate the receptive field and the gain of the neurons. They varied noise amplitude, which effectively changed the standard deviation of the stimuli. However, they kept fixed the shape of the luminance distribution (typically Gaussian or uniform), so that skewness and kurtosis was held constant.

In this chapter we use white noise stimuli with diverse luminance statistics to study the operation of the suppressive field. We measure the gain of neurons in lateral geniculate nucleus and investigate how gain depends on the standard deviation, skewness and kurtosis of the stimulus.

To estimate gain, we model responses by a linear filter followed by a static nonlinearity. The model predicts the main features of the responses. Its performance is largely

independent of the statistics of the stimulus. We infer gain from the amplitudes of the linear filters.

We find that gain strongly depends on the standard deviation of the stimulus but is largely invariant to skewness and kurtosis. These results confirm that gain control computes the true root-mean-square contrast of the stimulus. They also show that the suppressive field model holds for a broad range of light intensity distributions.

4.2 Results

We recorded responses from single neurons in lateral geniculate nucleus (LGN) of anesthetized, paralyzed cats. Stimuli consisted of grids of uniform squares whose light intensities were drawn at random (Figure 4.2A). Grids were updated at the frame rate of the monitor (124 Hz). We obtained a complete set of responses for a population of 25 neurons recorded in 3 animals.

We model responses with a linear filter followed by a static nonlinearity (Figure 4.2B). We estimate a separate filter for each stimulus and hold the nonlinearity constant across conditions (Chichilnisky, 2001). The linear filter captures the selectivity of responses for stimulus attributes. The static nonlinearity approximates the rectification of the spike threshold.

We use model fits to estimate the gain of the neurons (Figure 4.2**C**). The linear filter can be viewed as a collection of impulse response, one for each square in the stimulus. We estimate gain from the impulse response with maximum amplitude. We define gain as the standard deviation of the impulse response calculated across time.

We investigate how gain is affected by the statistics of the stimulus (Figure 4.3). We stimulated neurons with white noise whose distribution of light intensities varied from one condition to another. In the control condition (Figure 4.3A) light intensities were drawn from a uniform distribution. In the test conditions (Figure 4.3B), the standard deviation, skewness and kurtosis of the distribution were varied separately (Chubb et al., 1994). Mean luminance was held constant across all conditions (32 cd/m^2). Example frames of the test stimuli are shown in Figure 4.3C.

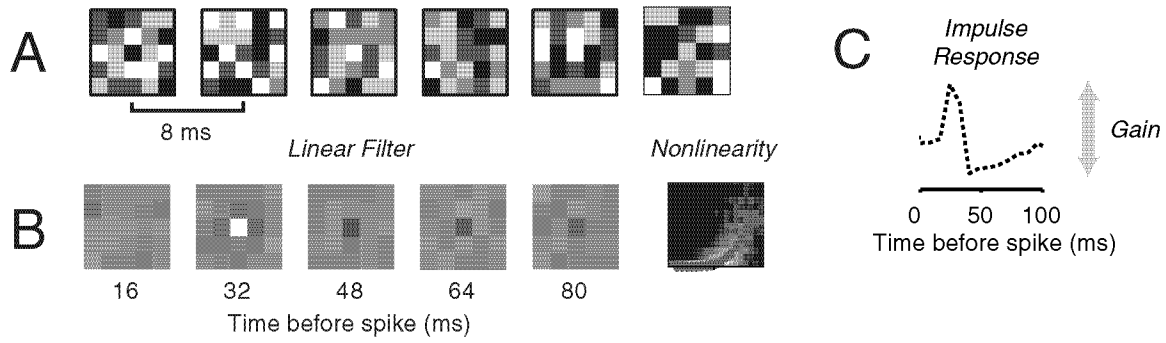

Figure 4.2. Estimating gain from responses to white noise stimuli. **A**: Example frames of the control stimulus. **B**: Example model fit (cell 51.3.6). Model consists of a linear filter with a static nonlinearity. **C**: Gain is defined as standard deviation of the impulse response of the square eliciting maximum response.

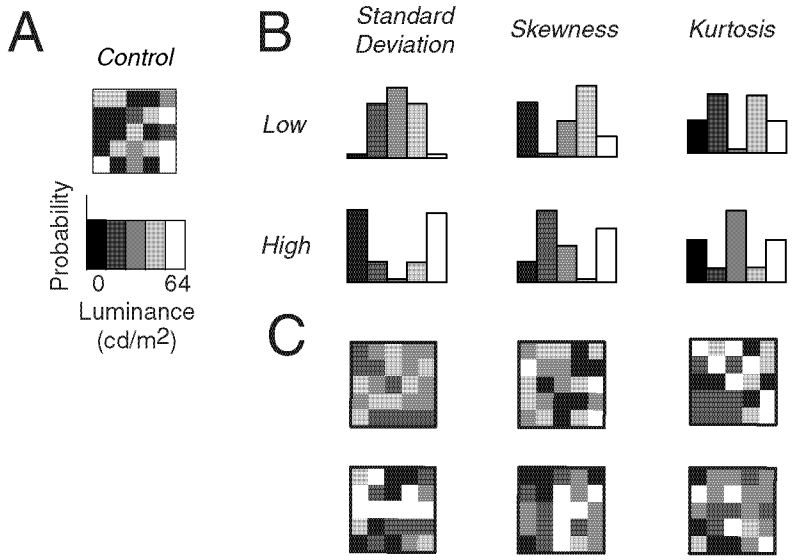

Figure 4.3. Stimulus conditions. A: Control condition. Light intensities are uniformly distributed. Histogram of light intensities (*bottom*) along with example frames (*top*). B: Histogram of light intensities for test conditions. The standard deviation, skewness and kurtosis of light intensities are varied. C: Example frames for test conditions.

4.2.1 Model Performance

Before we turn to study the model fits we test whether the model provides an adequate description of the responses. Any model will yield predictions that deviate in one way or another from the measured responses so it is important to define clear requirements. In the context of this work it suffices to show that the quality of model predictions is high, and that it is largely independent of the actual stimulus presented.

The model captures the main features of the responses. An example of model predictions for a typical cell is shown in Figure 4.4. The graphs show the measured (Figure 4.4, *shaded areas*) and the predicted responses (Figure 4.4, *solid curves*) for an interval of 2.5 s. The model predicts the timing of the firing events but is less accurate in predicting the magnitudes of these events. This behavior is observed for the low standard deviation condition (Figure 4.4A), and for the high standard deviation condition (Figure 4.4B).

The quality of model predictions is high and nearly independent of the actual stimulus presented. To quantify the quality of the predictions we calculate the percentage of stimulus-driven variance in the data that is explained by the model (Sahani and Linden, 2003; Machens et al., 2004) (see General Methods). Model predictions are good. Across neurons and stimulus conditions the model predicts 76.9±1.8% of the variance (mean±s.e, N=25). Moreover, the quality of predictions depends only weakly on the actual stimulus presented. The model explains 79.3±2.0% of the variance in the low standard deviation condition (12.8 cd/m^2). Explained variance decreases slightly to 77.6±1.7% and 75.4±1.8% as the standard deviation of the stimulus is increased (22.4 and 28.8 cd/m^2). A similar behavior was observed in the remaining conditions. These measures are summarized in Table 4.1.

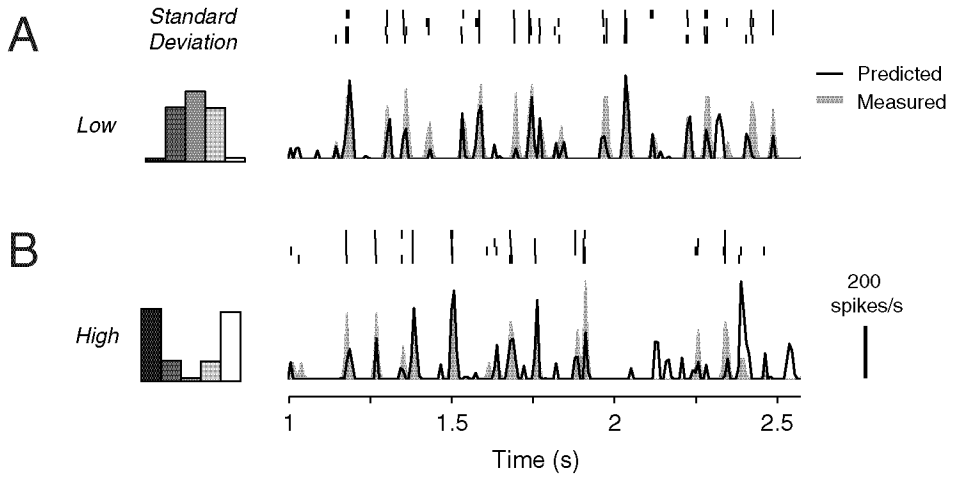

Figure 4.4. Example of model predictions. *Shaded area* indicate measured responses. *Solid line* indicate model predictions. **A**: Low standard deviation condition. **B**: High standard deviation condition.

	Stimulus			Example Cell		Population	
	Standard Deviation	Skewness	Kurtosis	Explained Variance	Gain	Explained Variance	Gain
	(cd/m²)	-	-	(%)	-	*mean±s.e.* (%)	-
Control	22.4	0.0	1.7	83.8	1.00	77.6±1.7	1.00
Test	12.8	0.0	1.7	77.7	1.46	79.3±2.0	1.38±0.16
	28.8	0.0	1.7	82.7	0.77	75.4±1.8	0.79±0.08
	22.4	-0.4	1.7	83.6	0.80	77.2±1.9	0.97±0.08
	22.4	0.4	1.7	78.7	1.00	77.2±1.8	0.97±0.11
	22.4	0.0	1.5	80.8	0.95	77.3±1.7	1.01±0.07
	22.4	0.0	1.9	77.4	0.89	74.5±2.0	0.93±0.04

Table 4.1. Summary of stimulus statistics and results. Example cell: 51.3.6.

A better way to assess the adequacy of model fits is to consider the full distribution of deviations. The analysis above depends on the variance of deviations and therefore does not include other aspects of the deviations that might correlate more with the statistics of the stimulus. To address this issue, we calculated the distributions of deviations and used quantile analysis to compare the distributions.

We find that model deviations are unlikely to reflect stimulus-dependent adaptations because the distribution of variations barely varies with the statistics of the stimulus. In Figure 4.5 we plot the quantiles of deviations for the test conditions as a function of quantiles obtained for the control condition. The resulting plots provide a measure of similarity between the distributions. For all plots, the data points fall near the diagonal line indicating that deviations in test and control conditions have nearly identical distributions.

In summary, the model captures the main features of the responses and responses vary with the statistics of the stimulus.

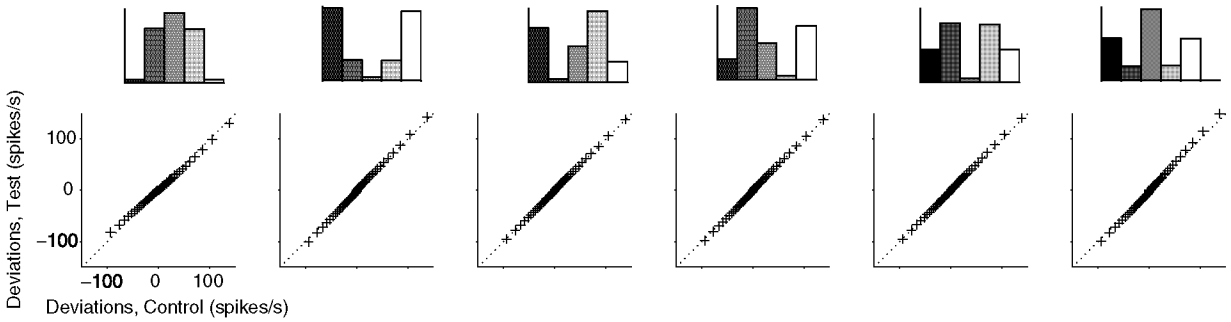

Figure 4.5. Quantile analysis of the distribution of deviations. Data points indicate quantiles of the distribution of deviations for the test conditions as a function of the quantiles obtained for the control condition.

4.2.2 Effects of standard deviation

We first investigate how varying the standard deviation of the stimulus affects the gain of the neurons. In the control stimulus the standard deviation of the stimulus was 22.4 cd/m^2, amounting to about a third of the dynamic range of the monitor (64 cd/m^2). In one condition we reduced standard deviation to 12.8 cd/m^2. In another condition we increased it to 28.8 cd/m^2. We held the mean, skewness and kurtosis constant throughout.

We find that varying standard deviation strong affects neural gain. Results for an example neuron are shown in Figure 4.6 (*middle*). For this cell, the impulse response obtained in the low standard deviation condition (Figure 4.6**A,** *gray curve*) is significantly larger than that in the control condition (Figure 4.6**A,** *dotted curve*). This larger impulse response reflects a 46% increase in gain. Likewise the impulse response in the high standard deviation condition (Figure 4.6**B,** *gray curve*) is smaller than in the control condition, reflected in a 23% reduction in gain.

Similar effects are observed for our population of neurons (Figure 4.6, *right*). Gains in the test conditions (Figure 4.6, *histogram*) are expressed relative to the gain observed in control condition, which is normalized to 1 (Figure 4.6, *dotted line*). Overall, gain is 40±3% larger (median, bootstrap estimates, N=25) in the low standard deviation condition (Figure 4.6**A**, *right*) than in the control condition. Similarly, gain is 32±1% smaller in the high standard deviation condition (Figure 4.6**B**) than in the control condition.

Interestingly, the effects of gain control increase as the strength of the stimulus increases. For low amplitude stimuli, the effects of gain control are weak. The 40±3% increase in gain observed when reducing standard deviation is well below the 75%

necessary to fully compensate for the changes in stimulus strength (Figure 4.6**A**, *dashed line*). For strong stimuli, by comparison, gain control nearly perfectly compensates for the changes in stimulus amplitude. The 32±1% decrease in gain observed when increasing standard deviation matches the 29% expected from the changes in signal amplitude (Figure 4.6**B**, *dashed line*).

In summary we have shown that the effects of gain control strongly depend on the exact standard deviation of the stimulus and that these effects are weak for weak stimuli and pronounced for strong stimuli.

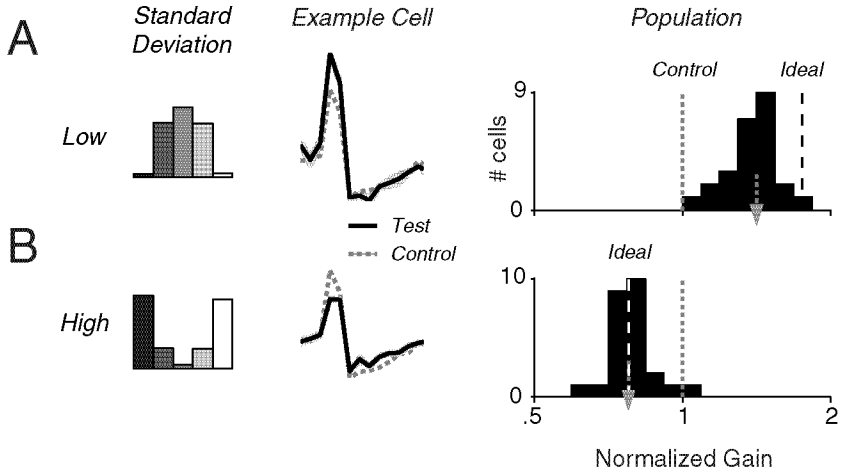

Figure 4.6. Effects of standard deviation. *Left*: Luminance histogram of test conditions. *Middle*: Results for an example cell (51.3.6). *Dashed curves* indicate the impulse response for the control condition. *Gray curves* indicate impulse responses for test conditions. *Right*: Results for population of neurons. Histogram of gains measured in test conditions relative to control. *Dotted lines* are gains in control condition. *Dashed lines* are gain changes necessary to fully compensate for changes in stimulus strength. *Arrows* are medians. **A**: Low standard deviation condition **B**: High standard deviation condition.

4.2.3 Invariance to skewness

We next ask whether the degree of asymmetry in the luminance distribution, or skewness, also affects the gain of the neuron. If the changes in gain we investigate indeed reflect a contrast gain control mechanism then stimulus skewness should have no effect. In the control condition the distribution is not skewed (skewness=0). In one test condition we skewed towards dark values (skewness = -0.4). In another we skewed it towards bright

values (skewness = 0.4). We held fixed the mean, the standard deviation and the kurtosis of the stimulus.

We find that skewness only has weak effects on the impulse responses (Figure 4.7). In the example cell, the impulse response of the low skewness condition (Figure 4.7**A**, *gray curve*) closely resembles that of the control condition (Figure 4.7**A**, *dotted curve*). Similarly, the impulse response of the high skewness condition (Figure 4.7**B**, *gray curve*) is undistinguishable from that of the control condition.

We find the effects of skewness on gain to be weak throughout our population of neurons. The gains observed for the low skewness condition (Figure 4.6, *histogram*) and for the high skewness condition (Figure 4.6, *histogram*) differ little from those in the control condition (Figure 4.6, *dotted line*). On average, gains differ from the control by 2±2 % and -6±4% (median, bootstrap estimates, N=25).

Skewness has distinct albeit small effects on the gain of ON and OFF cells. Skewing the stimulus towards dark values has a small but significant effect in ON cells (-6±4%, N=12) but has no effect in OFF cells (0±2%, N=13). Skewing the stimulus towards bright values has similar effects in both ON and OFF cells (-7±2% and -8±2%).

In summary, gain control depends only weakly on the skewness of the stimulus and is therefore insensitive to asymmetries in the light intensity distribution.

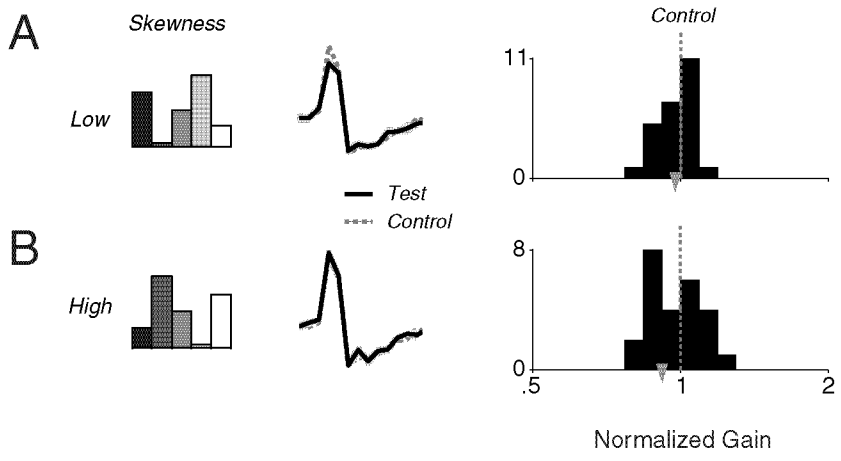

Figure 4.7. Invariance to skewness. Format as in Figure 4.6. **A**: Low skewness condition. **B**: High skewness condition.

4.2.4 Invariance to kurtosis

Finally we ask whether varying stimulus kurtosis has any effect on neural gain. The kurtosis describes the relative probability of stimulus deflections far away from the mean. The Gaussian distribution, which is often employed as a reference, has a kurtosis of 3, largely due to its long tails. Our stimuli had lower kurtosis because of the finite nature of the gray levels. In the control condition the stimulus had a kurtosis of 1.7. In one test condition we reduced kurtosis to 1.5. In another condition we increased it to 1.9. We held the mean, standard deviation and skewness of the stimulus constant.

We find that varying kurtosis has little effects on the impulse responses (Figure 4.8). In the example cell, the impulse responses for the low kurtosis condition (Figure 4.8**A**, *gray curve*) and for the control condition are undistinguishable (Figure 4.8**A**, *dotted curve*). Likewise the impulse response for the high kurtosis condition (Figure 4.8**B**, *gray curve*) resembles much that of control.

Across cells, kurtosis has only weak effects on gain. The gains for the low kurtosis condition is centered on control (Figure 4.8**A**, *histogram*) leaving the average gain unchanged (0±2%). The results for the high kurtosis condition (Figure 4.8**B**, *histogram*) reveals effects on gain (8±1%) but again these effects are small.

In summary we have shown that gain control is highly sensitive to the standard deviation of the stimulus but mostly invariant to skewness and kurtosis.

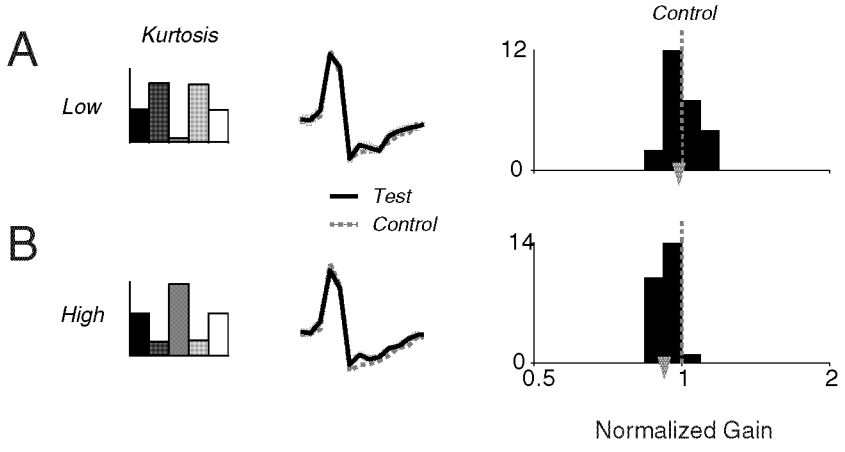

Figure 4.8. Invariance to kurtosis. Format as in Figure 4.6. **A**: Low kurtosis condition. **B**: High kurtosis condition.

4.3 Discussion

We used white noise stimuli to investigate the statistical operation of gain control. We studied how the statistics of the stimulus determine the gain of LGN neurons. We reasoned that if gain control computes root-mean-square contrast then it should show strong sensitivity to the standard deviation of the stimulus and be invariant to higher-order statistics of the stimulus.

Our results indicate that gain control in LGN estimate the true root-mean-square contrast of the stimulus. Whereas neural gain strongly depends on the standard deviation of light intensities, it is largely invariant to changes in stimulus skewness and kurtosis.

For strong stimuli, the effects of gain control perfectly compensate for changes in stimulus amplitude. For weak stimuli, however, gain control does not fully compensate these changes. This fall off in the strength of gain control with stimulus amplitude might help limit the effects of noise on responses to weak stimuli. Indeed gain signal likely becomes increasingly noisy as stimulus amplitude is decreased, and the fall off we observed might reflect a mechanism to counteract this reduction in signal-to-noise ratio.

Our results suggest that the suppressive field model holds for a broad range of stimuli. By manipulating the standard deviation, skewness and kurtosis of the stimulus we generated stimuli whose distributions of light intensities differ widely. Our results support the assumption that the suppressive field computes the standard deviation of light intensities for all these stimuli.

One potential problem in our approach is the method of reverse correlation to estimate the linear filters. This method was shown to provides unbiased estimates of the linear filter under the assumption of Gaussian stimuli and Poisson-generated spikes (Chichilnisky, 2001; Paninski, 2003). It can show significant biases when the light intensities of the stimulus are drawn from other distributions (Simoncelli et al., 2004) or when spike times are not independent (Pillow and Simoncelli, 2003).

These limitations, however, are unlikely to affect our conclusions. We measured the percentage of explained variance and found only weak dependence on the stimulus condition (Table 4.1). We also compared the full distribution of model deviations in the test conditions to that obtained in the control condition and found the distributions to be very similar

(Figure 4.5). Thus if there are biases in our estimates of neural gain, they are likely to be small.

Our finding that the processing in LGN is invariant to stimulus skewness is somewhat at odds with the psychophysical results of Chubb and colleagues (Chubb et al., 1994; Chubb et al., 2004). We have found no evidence in the responses of LGN neurons for a mechanism sensitive to asymmetries in the distribution of light intensities such as blackshot. A possibility is that a neural correlate for such a mechanism is not visible at the level of single cell responses. Another possibility is that it is only visible in higher visual areas.

A related question is whether the gain of LGN neurons is invariant to changes in the phase spectrum of the stimulus. Studies in retina (Smirnakis et al., 1997) and primary visual cortex suggest (Mechler et al., 2002) the existence of nonlinear mechanisms sensitive to spatial phase. We did not directly address this question but our stimuli did involve changes the phase spectrum. The amplitude spectrum of an image solely depends on its second-order luminance statistics such as the standard deviation. The phase spectrum, however, depends on higher-order statistics such as skewness and kurtosis. In our experiments we have manipulated the skewness and kurtosis of stimulus while holding standard deviation constant. Consequently we have varied the phase spectrum of the stimulus while holding its amplitude spectrum constant. Because these manipulations had only weak effects on the estimated linear filters, we can conclude that the processing of LGN neurons is invariant to these manipulations of the phase spectra.

4.4 Methods

Methods for physiological recordings, stimulation and data analysis are described in General Methods. We report here on recordings made in LGN of 3 adult anesthetized, paralyzed cats. We recorded from 25 well-isolated neurons (12 ON and 13 OFF). There is no overlap between this population of neurons and those presented in the remaining chapters.

Stimuli were grids of uniform squares whose light intensities were drawn at random. Stimuli had 5-10 elements in width, covered the receptive field center and surround, were presented at a rate of 124 Hz and lasted 20-30 s. Stimulus conditions were presented in randomized order and repeated 5-10 times. We varied the distribution of light intensities across stimulus conditions.

We used the method developed by Chubb et al (1994) to synthesize the distribution of light intensities. Briefly, we devised a set of orthogonal basis functions to modulate the standard deviation, the skewness and the kurtosis of the distribution. Let $v = 0,1,...,M$ denote the M different gray levels the stimuli can take. We computed the basis vector $f_i(v) = v^i$ where $i = 0,1,...,4$ denotes the order of the basis function. We then orthogonalized $f_i(v)$ using the Gram-Schmidt algorithm to obtain the set $F_i(v)$. Finally we normalized each vector $F_i(v)$ to an absolute maximum of $1/M$. We used the resulting vectors λ_i to synthesize the distributions. Translating a distribution alone λ_2 modulates its standard deviation while hold remaining statistics constant. Likewise translations along λ_3 and along λ_4 independently vary the skewness and the kurtosis of the stimulus.

Each experiment had 7 conditions. In the control condition, gray levels were drawn from the uniform distribution U (with probability $1/M$). In the low and high standard deviation conditions intensities were drawn from $U - k\lambda_2$ and $U + k\lambda_2$. In the low and high skewness conditions the stimuli followed $U - k\lambda_3$ and $U + k\lambda_3$. In the low and high kurtosis conditions the stimuli followed $U - k\lambda_4$ and $U + k\lambda_4$. We used $k = 0.9$ to ensure that all gray levels have nonzero probabilities. We used $M = 5$ gray levels to maximize the differences between the low and high conditions.

We used standard forms for expressing the standard deviation, skewness and kurtosis of the distributions. Let the random variable x_i denote the light intensity of any square in a stimulus. Standard deviation σ is the square root of second central moment or variance $\sigma^2 = E(x_i - \bar{x})^2$, where $E(\ )$ denotes the expected value operator and $\bar{x}$ is mean light intensity. Skewness is the third central moment: $E(x_i - \bar{x})^3 / \sigma^3$ and kurtosis is the fourth central moment $E(x_i - \bar{x})^4 / \sigma^4$.

We modeled responses with a linear filter followed by a static nonlinearity (Chichilnisky, 2001). First we estimated a separate filter for each stimulus by computing the average stimulus preceding a spike. We then estimated a separate nonlinearity for each stimulus by convolving the filter with the stimulus and scattering the resulting linear response against the measured response. We fixed the nonlinearity to the average measured response for a given linear response. At this stage we had one linear filter and one nonlinear function for

each stimulus. In a final step we rescaled the linear filters obtained for the test condition so as to minimize the mean square error between these filters and the filter obtained in the control condition. The resulting linear filters consist of one impulse response for each square in the stimulus. To estimate gain we considered the impulse response with maximum amplitude and calculated the standard deviation of the impulse response across time. Other measures of gain yielded similar results.

Chapter 5

Gain control, noise and spike threshold

We devise a probabilistic model of responses in lateral geniculate nucleus (LGN) and primary visual cortex (V1). The model includes (1) an integration field that captures the drive provided to the neuron, (2) a divisive suppressive field that describes the effects of contrast gain control, (3) a noise source that describes the trial-to-trial variability of responses and (4) a rectification stage that models the effects of spike threshold. We test the model on responses to stimuli of different contrasts and sizes recorded in anesthetized, paralyzed cats. In both LGN and V1, the model predicts the average responses, capturing the response saturation with increasing contrast and their selectivity for stimulus size. We also assess the model's ability to describe the standard deviation of responses calculated across trials. Model fits capture the distribution of standard deviation observed in LGN responses but underestimate the variability of V1 responses by approximately 40%.

5.1 Introduction

In the previous chapters we have studied the effects contrast gain control on the responses of neurons in lateral geniculate nucleus (LGN). Responses saturate with increasing contrast (*contrast saturation*); they are suppressed by superposition of a second stimulus (*masking*); they exhibit selectivity for stimulus size (*size tuning*). These effects are prominent and are likely have important implications for responses at subsequent stages of the visual system.

Indeed, neural responses in primary visual cortex (V1), the main recipient of LGN input, also exhibit pronounced effects of gain control. V1 responses show contrast saturation

(Maffei and Fiorentini, 1973; Dean, 1981; Albrecht and Hamilton, 1982; Ohzawa et al., 1982; Li and Creutzfeldt, 1984; Sclar et al., 1990), masking (Bonds, 1989; Freeman et al., 2002) and size tuning (Sceniak et al., 1999; Jones et al., 2000; Cavanaugh et al., 2002a). Can these effects be explained by those observed in the LGN input?

The contribution of the thalamic input to V1 suppression can be inferred by comparing responses to a single stimulus set. Saturation is more pronounced in V1 than in LGN (Sclar et al., 1990; Albrecht, 1995; O'Keefe et al., 1998; Freeman et al., 2002; Alitto and Usrey, 2004), suggesting that an additional mechanism might be at work. Studies of size tuning suggest that tuning LGN could contribute significantly to that observed in V1 (Jones et al., 2000; Ozeki et al., 2004).

At number of difficulties arise when comparing LGN and V1 responses. First, neurons in V1 exhibit higher firing thresholds and lower spontaneous activity than in LGN. High spike thresholds lead to strong rectification effects, which can significantly influence the shape of the contrast response curves. Second, V1 responses exhibit considerable trial-to-trial variability (Tolhurst et al., 1981; Bradley et al., 1987; Vogels et al., 1989; Geisler and Albrecht, 1997; Gur et al., 1997; Reich et al., 1997; Buracas et al., 1998), substantially more than LGN responses (Kara et al., 2000). This variability can strongly influence the tuning curves of the neurons (Anderson et al., 2000; Chance and Abbott, 2002; Miller and Troyer, 2002). Third, LGN and V1 responses are driven by distinct aspects of the stimulus. To fairly compare responses we need to estimate the actual drive provided to the neuron.

We overcome these difficulties by devising a simple model of the responses. The model includes (1) an integration field to describe the drive provided to the neuron, (2) a suppressive field and a divisive stage to account for the effects of contrast gain control, (3) a noise source to account for the trial-to-trial variability of responses, and (4) a rectification stage to account for differences in spike threshold. In Chapter 2, we have shown that the suppressive field captures the responses of LGN neurons of a variety of stimuli. We ask here whether it can also predict V1 responses. The combination of a noise source and a rectification stage has been shown to capture the mean and the trial-to-trial variability of V1 responses (Anderson et al., 2000; Hansel and van Vreeswijk, 2002; Miller and Troyer, 2002; Carandini, 2004). We ask here whether it can effectively describe the firing rates of both LGN and V1 neurons.

We test the model on responses to stimuli of different sizes and contrast recorded in LGN and V1 of anesthetized cats. The model captures the average responses, predicting the saturation of responses with increasing contrast and the tuning of responses for stimulus size. The model also describes the trial-to-trial variability of responses, explaining the growth in variability with response mean.

The model therefore provides a promising tool for comparing responses across brain areas.

5.2 Results

The results are organized as follows. We devise a simple probabilistic model of the responses to stimuli of different contrasts and sizes and test the result on recordings made in LGN and V1 of anesthetized cats.

The model is logically divided into two parts (Figure 5.1). First, the stimulus S is mapped onto a potential response V_G by processing it through an integration field and a suppressive field (Figure 5.1A). Second, the potential V_G is summed to noise and converted into firing rates (Figure 5.1B).

The integration field provides the drive to the neuron. The integration field is a Gaussian profile that describes the sensitivity of responses to a given stimulus falling at a particular location of visual field. The integration field differs from the receptive field in that it does not contain information about that the preferred stimulus is. This description is necessary because LGN and V1 neurons prefer different stimulus features. It is only valid if the frequency or the orientation of the stimulus is held constant. The output of the integration field O_I grows linearly with stimulus contrast. It grows as the overlap between the stimulus and the integration field increases.

The suppressive field captures the effects of contrast gain control. The suppressive field represents the Gaussian region of visual field over which gain control operates. Its output is a measure of local stimulus contrast, which divides the output of the integration field thereby controlling the gain of responses. The result of division is the generator potential V_G.

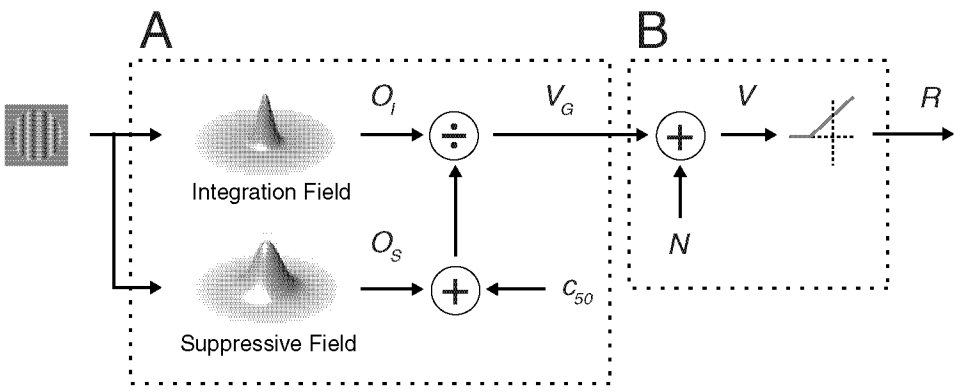

Figure 5.1. Probabilistic model of LGN and V1 responses. The stimulus is processed by an integration field and a larger overlapping suppressive field. The output of the suppressive field O_S is summed to a constant c_{50} to divide the output of the integration field O_I. The resulting generator potential V_G is summed to Gaussian noise and rectified yielding firing rates R.

We use the Gaussian-rectification model introduced by Carandini (2004) to account for the effects of noise in membrane potential and spike threshold. To allow for random fluctuations in the responses, the generator potential V_G is summed to independent Gaussian noise. The resulting potential V is compared to threshold V_T and rectified to obtain firing rates R.

The model describes both the mean and the trial-to-trial variability of responses (Figure 5.2). Increasing generator potential V_G increases mean firing rate because it progressively raises V above threshold. This increase in response mean is accompanied by an increase in response standard deviation because the distribution of R simultaneously grows in width. The model therefore predicts the well-known growth of response standard deviation with the mean (Tolhurst et al., 1981; Carandini, 2004).

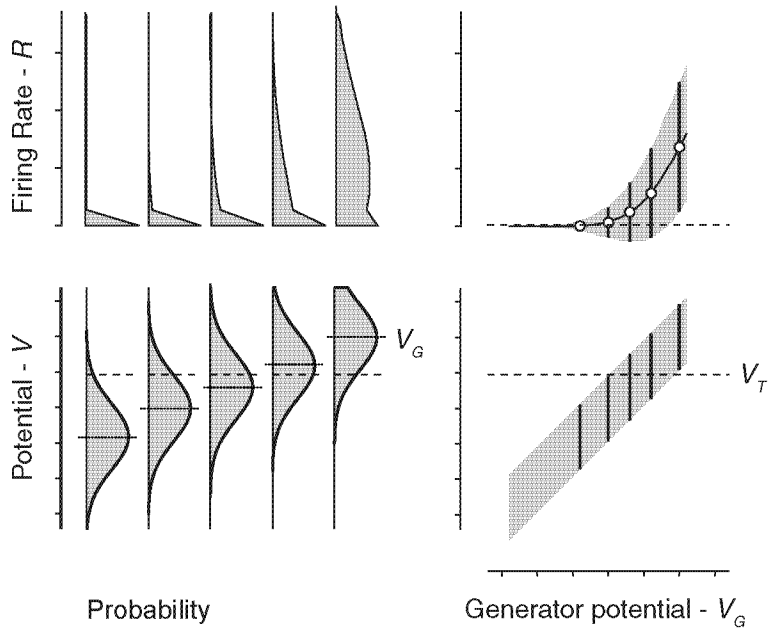

Figure 5.2. Gaussian-rectification model of firing rate encoding.

5.2.1 Predicting response mean

To test the model we recorded extracellular responses from neurons in LGN and V1 of anesthetized, paralyzed cats. We presented gratings enclosed in circular apertures and varied their contrast and diameter. Each experiment consisted of ≥ 50 stimuli including ≥ 2 blank stimuli to estimate the spontaneous activity of the neuron. Stimuli were repeated 3-7 times to estimate the trial-to-trial variability of responses. We consider the average firing rate measured during stimulus presentation.

These measurements fully constrain the six parameters of the model. We estimate the model parameters that maximize the likelihood of observing the data. The overall amplitudes of the responses set the responsiveness v_{MAX}. The tuning of responses for stimulus size constrains the width of integration field w_I and that of suppressive field w_S. The saturation of responses with contrast constrains the strength of gain control c_{50}. The responses to blank stimuli, along with the trial-to-trial variability of responses, determine threshold v_T and noise amplitude σ.

The model predicts the phenomena of contrast saturation and size tuning. The graphs in Figure 5.3 show the responses obtained for four LGN neurons (Figure 5.3A-D) and four V1 neurons (Figure 5.3E-G). In all neurons, size tuning is weak at low contrast (Figure 5.3,

dark symbols) and pronounced at high contrast (Figure 5.3, *light symbols*). The flip side of these observations is that saturation depends on stimulus size. In all neurons, saturation is weak for small stimuli and pronounced for large stimuli. The model captures these observations (Figure 5.3, *curves*). The model explains size tuning because extending stimulus beyond the confine of the integration field increases O_S while O_I is approximately constant. At low contrast, responses depend linearly on stimulus contrast and size because the denominator in divisive stage is dominated by the constant c_{50}. Responses to large, high contrast stimuli saturate because c_{50} becomes much smaller than the output of suppressive field O_S.

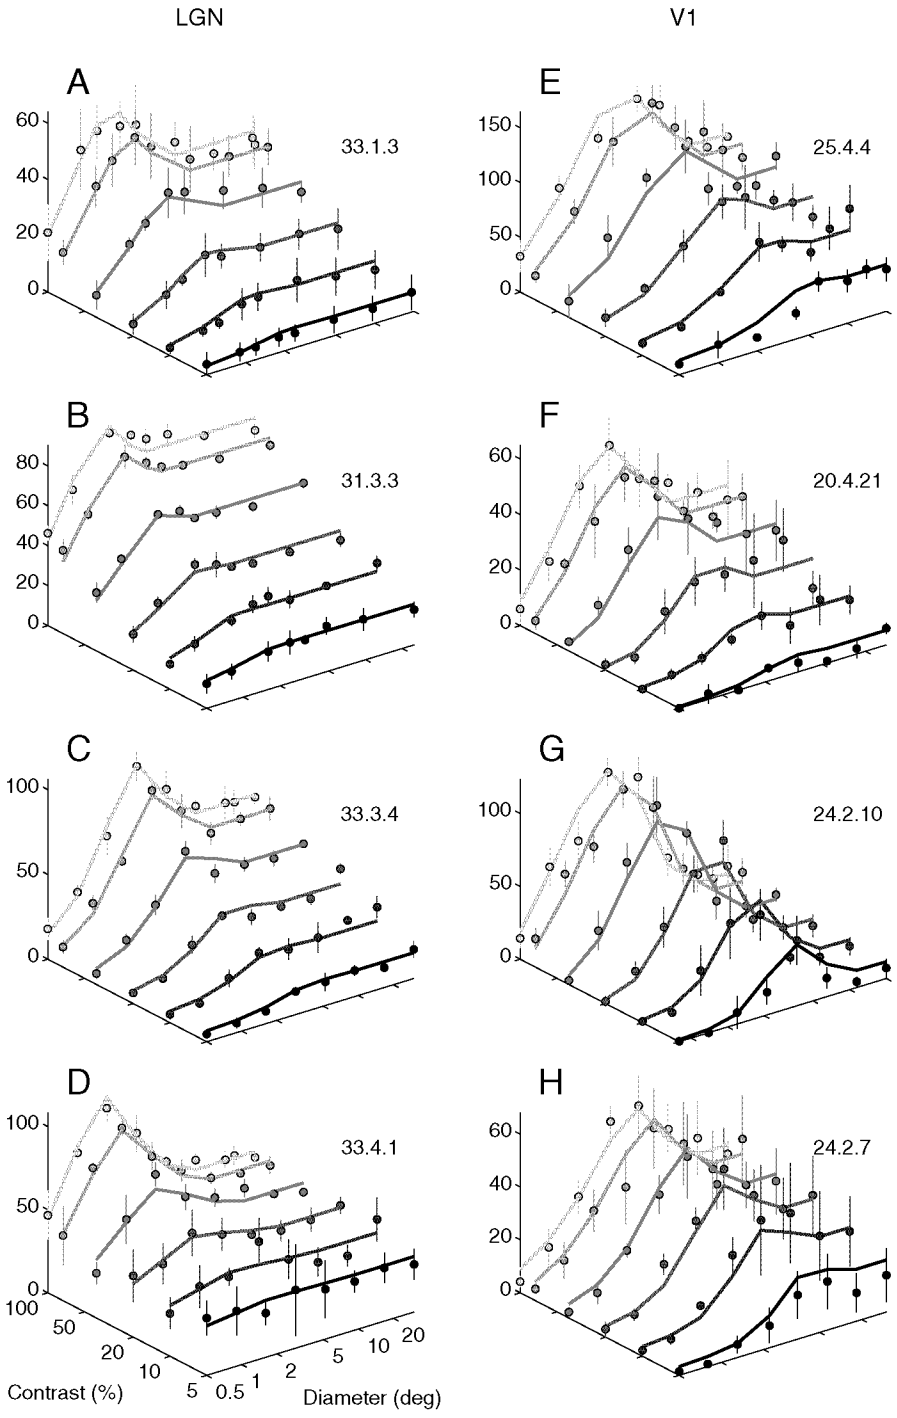

Figure 5.3. Example fits for several LGN and V1 neurons. *Symbols* are measured responses as function of contrast and diameter. *Curves* indicate model predictions. *Symbols* and *curves* are assigned one gray level for each diameter. **A**: Responses from the same LGN neuron as in Figure 6.1AB. **E**: Responses from the same V1 neuron as in Figure 6.1CD. **B-D**: Responses from more LGN neurons. **F-H**: Responses from more V1 neurons.

5.2.2 Predicting response variability

In addition to describing the response mean, the model also describes their trial-to-trial variability. This is illustrated in Figure 5.4 where we plot the response mean (Figure 5.4, *symbols*) and standard deviation (Figure 5.4, *error bars*) as a function of contrast (for the same LGN neuron as in Figure 5.3B and V1 neuron in Figure 5.3F). The model not only predicts the response mean (Figure 5.4, *curves*) but also their standard deviation (Figure 5.4, *shaded areas*). It captures the low variability of the example LGN neuron (Figure 5.4A) and the pronounced variability of the example V1 neuron (Figure 5.4B). The model also correctly predicts that variability is lower for weak responses (Figure 5.4B, *top axes*) and pronounced for large responses (Figure 5.4B, *middle axes*).

The model therefore explains the growth of response variability with the mean often observed in the responses of V1 neurons (Tolhurst et al., 1981; Carandini, 2004). Variability is low for small, low contrast stimuli because for these stimuli potential V generally lies below threshold V_T. Increases stimulus contrast or size are increases variability because it progressively pushes V above threshold V_T.

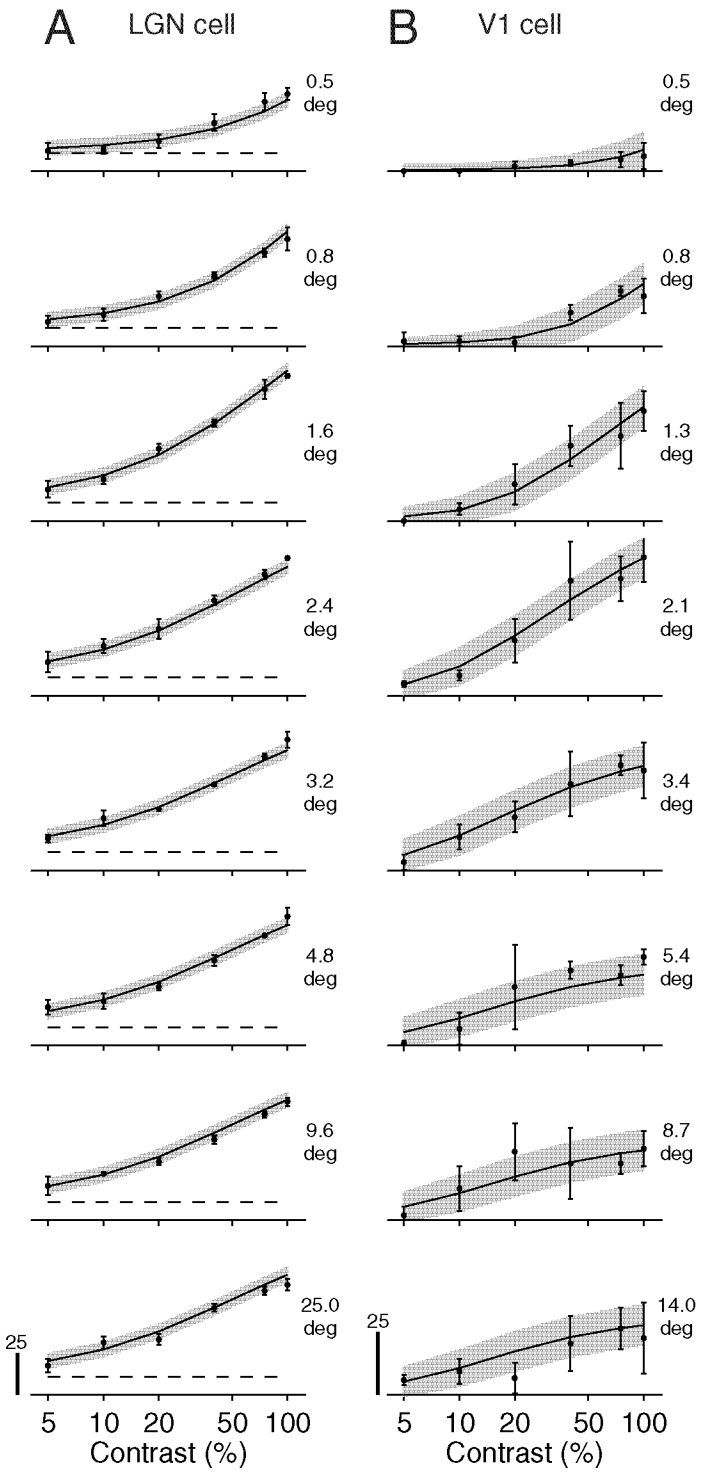

Figure 5.4. Predicting response mean and variability. *Symbols* and *error bars* indicate mean and standard deviation of measured responses. *Solid curves* and *shaded areas* indicate model predictions. *Dashed lines* indicate spontaneous firing rates. Each axis corresponds to a distinct diameter. **A**: Responses of the LGN neuron as in Figure 5.3B (cell 31.3.3). **B**: Responses of same V1 neuron as in Figure 5.3F (cell 20.4.21).

We can better visualize the model's ability to predict the variability of responses using mean-variance plots (Figure 5.5). We plot the trial-to-trial variance of responses as function of the mean for the example LGN neuron (Figure 5.5A) and for the example V1 neuron (Figure 5.5B). For the LGN neuron, data points mostly fall below the unity line indicating that variability is lower than that predicted by a Poisson process. For the V1 neuron, data points are scattered around the unity line indicating a more pronounced trial-to-trial variability. The model can predict the distribution of response variance (Figure 5.5AB, *shaded areas*). According to the model, 50% of the data points should fall within the dark gray region, 75% of them should fall within the intermediate gray region, and 95% of them should fall the light gray region.

The quality of these predictions can be assessed using quantile analysis (Figure 5.5CD). We plot the quantiles of observed standard deviation versus those predicted by the model. The data points for the example LGN neuron (Figure 5.5C) fall on the unity line indicating that model captures the variability of responses. By comparison, the data points for the example V1 neuron (Figure 5.5D) fall below the unity line indicating that the model underestimates variability.

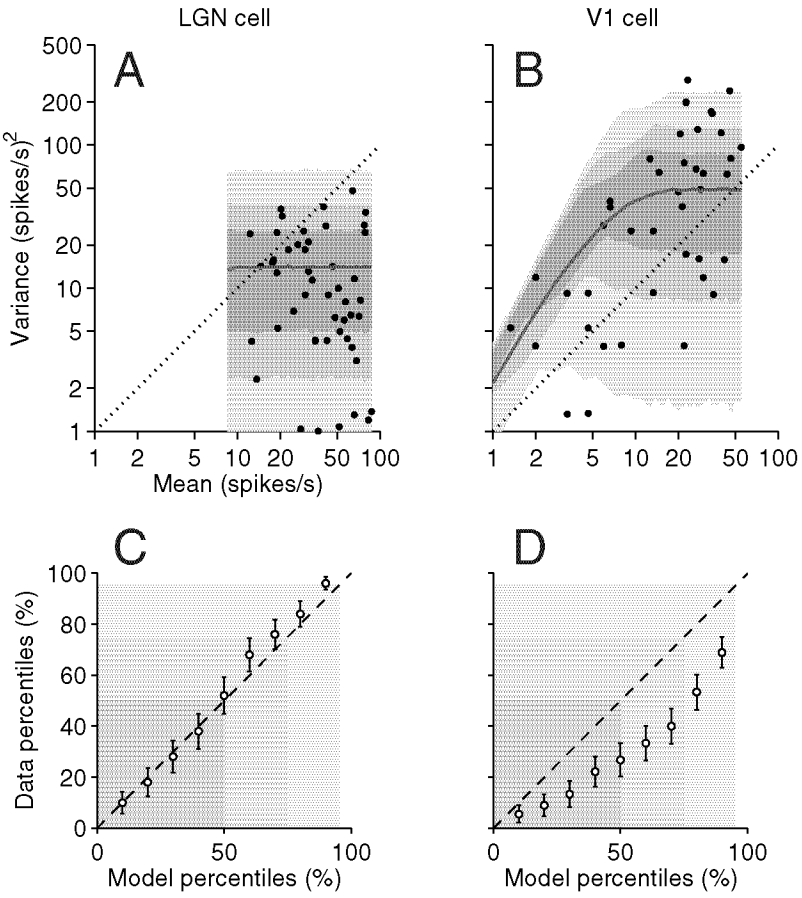

Figure 5.5. Predicting response variability. **A**: Responses of an LGN neuron. *Data points* indicate measured mean versus variance for all stimuli. *Shaded areas* indicate model predictions. These areas delineate the boundaries of the 50th percentile (*dark*), of the 75th percentile (*gray*) and of the 95th percentile (*light*). *Curve* indicates the expected value of the predicted mean-variance functions. **B**: Same for a V1 neuron. **C**: Quantile analysis of response variability in the LGN neuron. *Symbols* indicate quantiles of measured standard deviation as function of 10% quantiles of model predictions. *Error bars* indicate ±1 s.d. of measured quantiles (bootstrap estimates). *Dashed line* indicates expected result for identical distributions. *Shaded areas* indicate 50th percentile (*dark*), 75th (*gray*) and 95th (*light*) percentiles **D**: Same, for the V1 neuron.

5.2.3 Model performance

We turn to quantify model performance for our populations of neurons. We evaluate model predictions for the mean responses and floor the standard deviation of responses. We ask whether the model predicts equally well the responses of LGN and V1 neurons.

The model performs well in predicting the average responses. We calculate for each neuron the percentage of variance in the average data that is explained by the model. In

LGN, the model performs relatively well, explaining >90% of the variance in 30/39 neurons. In V1, however, the quality of fits is lower, with >90% of the variance explained in 28/60 neurons. This lower explained variance might indicate lower model performance, or it might simply reflect the pronounced trial-to-trial variability of V1 responses. Stimuli were only presented a limited number of times so our estimates of the mean can be noisy. To test for that possibility, we use Machens' method (2004) to estimate the percentage of stimulus-driven variance in the data that is explained by the model (see General Methods). We find that the model explains >95% of the variance in nearly all the LGN and V1 neurons (N=38/39 and 55/60). The model performs equally well in ON and OFF cells. The median explained variance for these groups did not differ significantly (Wilcoxon rank sum, $p>0.2$). We obtained similar results for simple and complex cells.

Model fits capture the variability of LGN responses but tend to underestimate that of V1 responses. As explained in the previous section, the model has the ability to predict the standard deviation of responses. We can quantify this ability by comparing the distribution of observed variance to that predicted by the model. For LGN neurons, the model captures the distribution of standard deviation in 25/39 neurons (i.e. measured quantiles within $\pm$ 2 s.d. of predicted quantiles). For V1 neurons, fits are less accurate with predictions being satisfactory in only 8/60 neurons. Overall, model fits tend to underestimate the actual variability of responses. For LGN neurons, predicted standard deviation amounts for 91$\pm$17% of the observed scatter in standard deviation (N=39, mean$\pm$s.d.). For V1 neurons, the predicted standard deviation amounts for 60$\pm$20% of the observed scatter in standard deviation (N=60).

In summary, the model provides a compact description of response mean and standard deviation. In LGN, model predictions are accurate for the mean and standard deviation. In V1, model predictions are accurate for the mean but the fits underestimate the trial-to-trial variability of responses.

5.3 Discussion

We have devised and validated a simple probabilistic model of LGN and V1 responses. The model predicts responses to stimuli of different contrasts and sizes. The model includes (1) an integration field to describe the drive provided to the neuron, (2) a suppressive field and a divisive stage to capture the effects of contrast gain control, (3) a noise source to

capture the trial-to-trial variability of responses, (4) a rectification stage to capture the effect of the spike threshold.

The model provides a unified description of LGN and V1 responses to stimuli of different contrasts and sizes. The model can be simplified to the model of LGN responses introduced and validated in Chapters 2-4. The model also resembles models of cortical gain control previous applied to the responses of V1 neurons (Sceniak et al., 2001; Cavanaugh et al., 2002a). The model improves on previous efforts in several ways.

First, the model explicitly describes the effects of contrast on the responses. Previous models required measuring empirically these effects, which introduced several additional parameters and endowed individual neurons with their own definition of contrast. Our model postulates that gain control computes local root-mean-square contrast, a measure that applies to arbitrary images.

Second, the model uses a linear encoder of firing rates while previous models assumed an expansive nonlinearity in the input-output relationship of the neuron. The contrasts-response curves of V1 neurons is generally described by adding an expansive nonlinearity, or an exponent, to the contrast-response curves (Sclar et al., 1990; Albrecht, 1995; Sceniak et al., 2001; Cavanaugh et al., 2002a). There are, however, several problems to this approach. One problem is that firing rate encoding is thought to be linear (Carandini and Ferster, 2000). Another problem is that different nonlinearities are needed to explain the different degrees of steepness observed in LGN and V1. Finally the effects of the nonlinearity on the response curves can trade with those of gain control (Chapter 6). We found that a linear encoder accurately predicts the responses, without the need for an expansive nonlinearity. One possible explanation for this success is that our model accounts for the effects of noise in membrane potential, which can produce such nonlinearity (Anderson et al., 2000; Hansel and van Vreeswijk, 2002; Miller and Troyer, 2002). Another possible explanation is the effect of spike threshold, which can increase the steepness of the response curves. We address these issues in Chapter 6.

Third, the model describes the trial-to-trial variability of responses. Variability differs substantially from one brain area to another (Kara et al., 2000). The model is a promising framework to describe these differences. Model fits capture the variability of LGN responses but they underestimate that of V1 responses by ~40%. The reasons for this bias are unclear.

It might be an artifact of the fitting procedure. Or it might indicate a genuine failure of the model. Our study differs from Carandini (2004) in the sampling rate of responses. We considered here the mean firing rate calculated over the entire duration of the stimulus (>1s). Carandini (2004) used a much shorter sampling interval (10 ms).

Central to the model is also the assumption that noise in membrane potential is additive and follows a Gaussian distribution. This assumption is appropriate for the responses of neurons in cat primary visual cortex (Carandini, 2004). A recent study of retinal responses has shown that their trial-to-trial statistics are entirely consistent with the hypothesis of additive Gaussian noise (Passaglia and Troy, 2004). We do not know of equivalent data in the lateral geniculate nucleus but the assumption seems reasonable considering (1) that responses were measured over long time scales (i.e. >1s, the duration of the stimulus), and (2) that stimuli were presented in random order.

Finally, the model describes the effects of the spike threshold. LGN and V1 neurons exhibit different spike threshold. LGN neurons have low thresholds, which are reflected in their high spontaneous firing rates. In these neurons, threshold can be estimated from the spontaneous firing rate. By comparison, V1 neurons have higher thresholds, resulting in low spontaneous firing rates. In these neurons, threshold cannot be directly inferred from the spontaneous firing rate. The model provides a solution to this problem. It uses information from the statistics of the responses to infer estimates of spike threshold and noise in membrane potential. While the model may not reveal the true membrane potential of the neuron, it reveals subthreshold information contained in the extracellular responses. We show an application of this concept in Chapter 6.

Our approach also has limitations. For instance, we did not consider the phase of the responses, which is strongly influenced by gain control (Dean and Tolhurst, 1986; Carandini and Heeger, 1994). Also, suppression in V1 is known to be selective for stimulus orientation (Blakemore and Tobin, 1972; DeAngelis et al., 1994; Li and Li, 1994; Jones et al., 2001; Cavanaugh et al., 2002b), being more pronounced at the preferred orientation of the neuron. Models of contrast gain control cannot explain this selectivity. There is also evidence that suppression is concentrated into hot spots around the receptive field (Walker et al., 1999; Cavanaugh et al., 2002b). Our model does not account for possible asymmetries in the sources of suppression.

5.4 Methods

Methods for physiological recordings, stimulation and data analysis are described in General Methods. We report here on recordings made in lateral geniculate nucleus and primary visual cortex (area 17/18) of adult anesthetized, paralyzed cats. We present data from 20 animals, 7 in LGN and 13 in V1.

We stimulated neurons with optimal drifting gratings enclosed in circular apertures. Each experiment included > 50 stimuli covering a full matrix of >6 contrasts and >8 diameters distributed on logarithmic scale. Stimuli were repeated 3-7 times.

We ran this experiment on 206 well-isolated neurons, 56 in LGN and 150 in V1. We selected the subset of cells that gave a reliable responses to these stimuli (mean rate of >5 spike/s for at least one stimulus) and for which we had complete size tuning curves (with the smallest diameter tested well below optimal and the largest chosen so that responses reached a plateau). The resulting populations of LGN (N=39) and V1 neurons (N=60) were further analyzed. The LGN neurons included 22 ON cells and 17 OFF cell and had receptive fields with an average eccentricity of 13.8±10.8 deg (mean±s.d.). Some of this data has already been presented in Chapter 2 (12 of 39 cells). The V1 neurons included 28 simple cells and 32 complex cells and had receptive fields with an average eccentricity of 5.0±3.7 deg.

We described the responses with a simple probabilistic model (Figure 5.1). In the model, the stimulus is processed by two Gaussian fields: an integration field describing the drive provided by the stimulus to the neuron, and a larger overlapping suppressive field capturing the spatial footprint of contrast gain control. The outputs of integration field and suppressive field O_I and O_S meet at divisive stage yielding the deterministic generator potential

$$V_G = V_{\max} \frac{O_I}{c_{50} + O_S}, \qquad 4.1$$

where V_{max} denotes the overall responsiveness of the neuron and c_{50} determines the strength of gain control. For a grating stimulus of contrast c and diameter d, the outputs of integration field and suppressive field are given by $O_I = c\left(1 - e^{-2d^2/w_I^2}\right)$ and $O_S = c\left(1 - e^{-2d^2/w_S^2}\right)$ where w_I and w_S are their widths.

To account for the trial-to-trial variability of the responses the generator potential is summed to Gaussian noise (Carandini, 2004). The resulting potential V is a Gaussian random variable of mean V_G and width σ:

$$P(V) = N[V_G, \sigma](V) \qquad 4.2$$

To obtain firing rates potential V is 4.3compared to threshold V_T and the result is rectified: $R = [V - V_T]_+$. Firing rate R follows a truncated Gaussian distribution. For positive rates (R>0), the probability of firing follows

$$P(R) = N[V_G - V_T, \sigma](R). \qquad 4.4$$

The probably of silence is the probability that $V<V_T$:

$$P(0) = \int_{-\infty}^{V_T} N[V_G - V_T, \sigma](V)\,dV = \operatorname{erf}\left(\frac{V_G - V_T}{\sigma}\right). \qquad 4.5$$

Let R_{ij} denote the observed response to stimulus i at trial j. We estimated the set of model parameters that maximize the probability of observing the data:

$$\arg\max \sum_{i,j} P(R_{ij}). \qquad 4.6$$

Chapter 6

Contrast saturation and size tuning in LGN and V1

We compare the effects of contrast gain control at two consecutive stages of the visual system, the lateral geniculate nucleus (LGN) and the primary visual cortex

(V1). We study the growth and saturation of responses with increasing contrast, and the tuning of responses for stimulus size. We find that V1 responses saturate more than LGN responses and that LGN and V1 neurons exhibit distinct patterns of size tuning. We use a probabilistic model to tease apart the contributions of contrast gain control, trial-to-trial variability of responses and spike threshold to these observations. We find that differences in gain control and spike threshold explain the differences in saturation. By comparison, the potential contribution of trial-to-trial variability is small. We also find that LGN and V1 neurons exhibit similar degree of size tuning when factoring out the effects of noise and of spike threshold. These results suggest that the LGN input likely contributes substantially to the effects of gain control that are observed in V1.

6.1 Introduction

Responses of neurons in lateral geniculate nucleus (LGN) and primary visual cortex (V1) are shaped by profound suppressive phenomena. Responses show saturation with increasing contrast (Sclar et al., 1990) (Figure 6.1AC). They also show selectivity for stimulus size (Jones et al., 2000; Ozeki et al., 2004) (Figure 6.1BD).

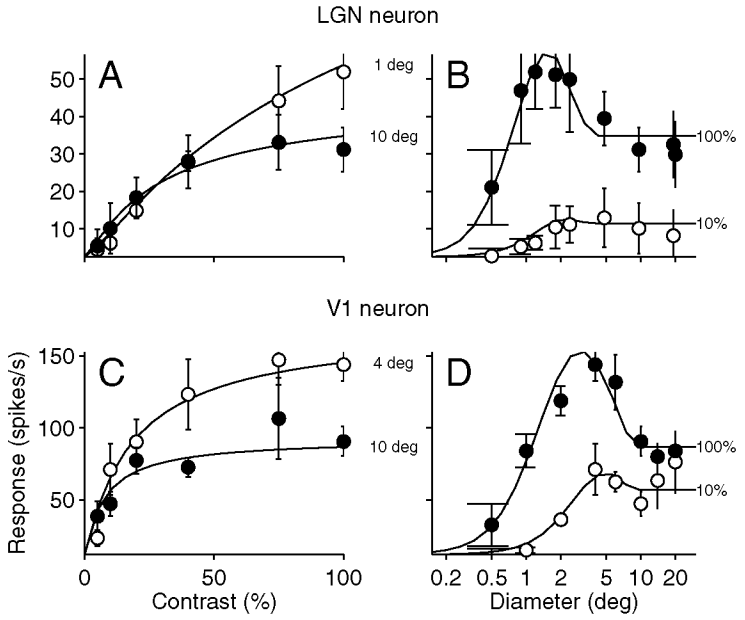

Figure 6.1. Contrast saturation and size tuning. Stimuli are moving gratings in circular apertures of different contrast and diameters. *Symbols* indicate mean firing rates. Error bars indicate standard deviation of firing rates. *Curves* are model predictions (see Results). A: Responses of an example LGN neuron (cell 33.1.3) as function of contrast for two selected diameters. B: Responses of the same LGN neuron as function of diameter for two selected contrasts. C: Responses of example V1 neuron (cell 25.4.4) as function contrast for two selected diameters. D: Response of the same V1 neuron as function of diameter for two selected contrasts.

The origins of the suppressive phenomena as observed in the responses of V1 neurons are unclear (Angelucci et al., 2002; Freeman et al., 2002; Sengpiel and Vorobyov, 2005). Early explanations involved intracortical inhibition (Benevento et al., 1972; Blakemore and Tobin, 1972). Models of V1 responses generally assume a linear input (Albrecht and Geisler, 1991; DeAngelis et al., 1992; Heeger, 1992; Carandini et al., 1997; Sceniak et al., 1999; Sceniak et al., 2001; Cavanaugh et al., 2002a; Sceniak et al., 2002) and therefore also invoke a cortical mechanism. Recent work suggests, however, that the thalamic input may also play an important role (Carandini et al., 2002; Freeman et al., 2002; Solomon et al., 2002; Ozeki et al., 2004).

In the previous chapter, we have devised a simple model to describe the suppressive phenomena (Figure 6.1, *curves*). The model includes (1) an integration field that captures the drive provided to the neuron, (2) a suppressive field and a divisive stage that accounts for the effects of contrast gain control, (3) a noise source that describes the trial-to-trial

variability of responses and (4) a rectification stage that models the effects of the spike threshold. Here we study how these components act together to yield a response.

6.2 Results

We report here on the same recordings and model fits presented in Chapter 5. These involved recordings made in LGN and V1 of adult anesthetized, paralyzed cats. We present data from 39 LGN neurons and 60 V1 neurons recorded from 20 animals (7 in LGN and 13 in V1). Stimuli were moving gratings of different contrasts and diameters.

Having validated the model in the previous chapter, we now investigate the factors determining the contrast-response and size-tuning curves of the neurons. We study the growth of responses with contrast at low contrast and their degree of saturation at high contrast. We also investigate the strength and spatial extent of size tuning. Throughout our analysis we assess the contributions of contrast gain control, noise in membrane potential and spike threshold to these phenomena.

6.2.1 Contrast-response curves

We first consider the average behavior of our populations of neurons. To summarize the data, we calculate average contrast-response curves by pooling together responses to stimuli of similar diameters, relative to the preferred diameter of the neurons. We do so for the measured responses and for the model predictions.

V1 responses grow more steeply with contrast and saturate more than the responses of LGN neurons (Figure 6.2). V1 responses (Figure 6.2A, *closed symbols*) differ from LGN responses (Figure 6.2A, *open symbols*) in a number of aspects: (1) they have relatively low amplitudes; (2) they show less spontaneous activity; (3) they grow more steeply with stimulus contrast at low contrast and (4) they show stronger saturation at high contrast. These differences are marked for stimuli of optimal diameter (Figure 6.2A) but are less pronounced for large stimuli (Figure 6.2B).

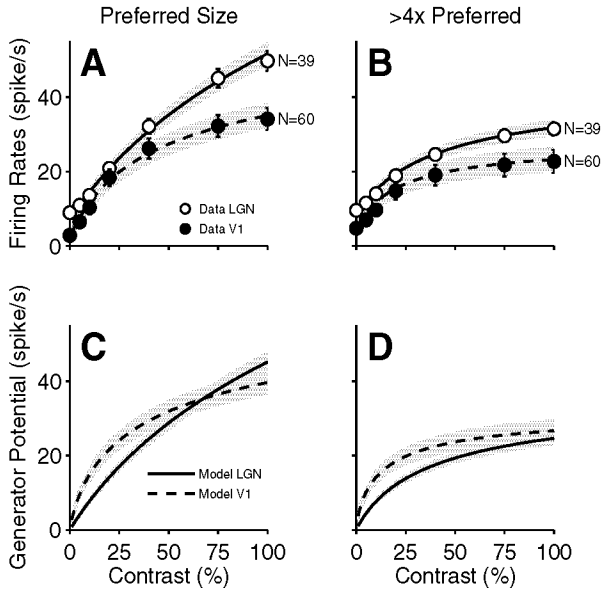

Figure 6.2. Average contrast-response curve for populations of LGN and V1 neurons. *Symbols* are average of measured responses. *Error bars* are ±1 s.e. of measured responses. *Curves* are average model predictions. Shaded areas are ±1 s.e. of model predictions. **A**: Firing rate responses to stimuli near optimal diameter. **B**: Firing rate responses to stimuli >4x the preferred diameter. **C**: Estimated generator potential in response to stimuli near optimal diameter. **D**: Estimated generator potential in response to stimuli >4x the preferred diameter.

The model precisely captures these observations. The average model predictions (Figure 6.2, *curves*) are accurate and show little overall bias. The model captures critical aspects of the responses. It captures (1) the level of spontaneous activity, (2) the saturation of responses at high contrast, (3) the dependence of the responses on the size of the stimulus, and (4) the growth of responses at low contrasts. The latter point is best seen by re-plotting the response curves on a logarithmic axis (Figure 6.3).

We next investigate the factors determining the shape of the contrast-response curves. Does the more pronounced saturation observed in V1 responses enhanced gain control? What is the role of noise in shaping the responses? The model can help us address these questions.

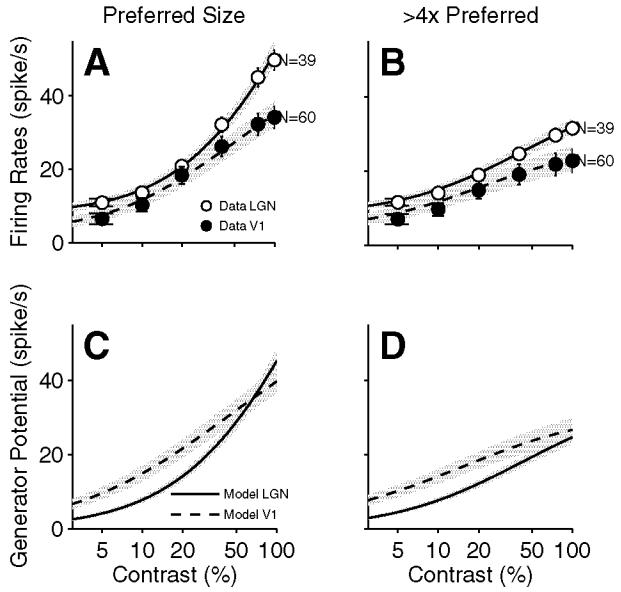

Figure 6.3. Same data as in Figure 6.2 but with contrast-response curves plotted on a logarithmic abscissa.

The effects of gain control are four times stronger in V1 neurons than in LGN neurons. We can isolate the effects of gain control from those of noise and spike threshold by looking at the generator potential responses in the model fits. Generator potential responses show stronger saturation in V1 neurons (Figure 6.2CD, *dashed curves*) than in LGN responses (Figure 6.2CD, *solid curves*). These differences are pronounced for stimuli of preferred size (Figure 6.2C) but less so for large stimuli (Figure 6.2D). For larger stimuli, the strength of gain control is determined by the constant c_{50}. In our population of LGN neurons the average c_{50} is 28±4% (median, bootstrap estimates, N=39). In V1 gain control is four times stronger with an average c_{50} of 7±2% (N=60).

What are the effects of noise and spike threshold on the contrast-response curves? To address this question we inspect model predictions and vary the strength of gain control c_{50}, the spike threshold V_T and the noise amplitude σ. We set the remaining parameters to the median obtained for our population of V1 neurons (similar results were obtained using the parameters from our population of LGN neurons). We normalize the response curves so that the spontaneous firing rate equals 0 and the peak response equals 100%. The results of this analysis are shown in Figure 6.4. As already mentioned varying c_{50} strongly affects the response curves. Increasing its value reduces the strength of saturation (Figure 6.4A). Increasing c_{50} results in a shift of the response curves to the right (Figure 6.4B) (Heeger, 1992).

The effects of threshold V_T counteract those of contrast gain control. Reducing V_T below 0 has little effect on the curves (Figure 6.4C, *thin curves*). Increasing V_T above 0 reduces the saturation of responses at high contrast (Figure 6.4C, *thick curves*). On a logarithmic scale increasing V_T shifts the response curves to the right (Figure 6.4D).

The spike threshold plays a significant role in shaping the response curves of LGN and V1 neurons. From the model fits the threshold V_T has an average value of -4.3±1.5 spike/s (median, bootstrap estimate, N=39) in LGN versus 3.9±0.7 spikes/s (N=60) in V1. Varying V_T over that range has pronounced effects on the response curves (Figure 6.4CD).

By comparison, the potential role of noise in shaping the contrast-response curves of LGN and V1 neurons is small. Increasing noise amplitude σ threefold has little effects on the strength of saturation at high contrast (Figure 6.4E) and on the growth of responses at low contrast (Figure 6.4F).

The responses of V1 neurons are on average about twice as variable as responses in LGN. The model fits provide us with such estimates. On average, noise amplitude is σ = 5.4±0.5 spikes/s (N=39) in LGN and 6.7±0.7 spikes/s (median±s.e., N=60) in V1. In the previous section we have shown that the model fits underestimate the standard deviation of the responses. In LGN the amplitude of variability predicted by the model is 9±17% lower than in the actual responses. Correcting for this bias yields effective noise amplitude $\hat{\sigma}$ of about 5.9 spikes/s. In V1 model fits underestimate variability by 40±20% so the effective noise amplitude $\hat{\sigma}$ is 9.4 spikes/s.

In summary, we have shown that enhanced gain control and the rectification of the spike threshold can explain the differences observed between the contrast-response curves of LGN and V1 neurons. The potential effect of noise amplitude is small.

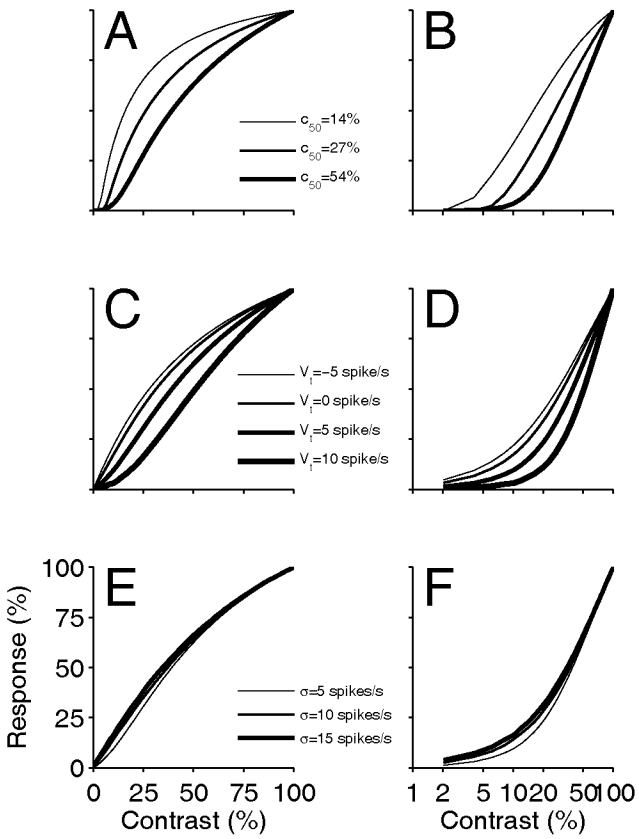

Figure 6.4. Effects of model parameters on contrast-response curves.

6.2.2 Strength of size tuning

We now turn to study the selectivity responses for stimulus size. An important question is whether the tuning of V1 responses could potentially be explained by the tuning of the LGN input. One way to tell these alternatives apart is to compare the strength of tuning in both areas. Size tuning in V1 may be too strong to be explained by the properties of the LGN input. The model can help us estimating the strength of tuning while factoring out the effects of the spike threshold.

First, we look at the average behavior of the neurons. We consider the size tuning curves of firing rates for our populations of neurons. We normalize the abscissa of each curve by the preferred stimulus diameter measured at 100% contrast and average across neurons. We compute one average curve for each contrast tested.

LGN and V1 responses show similar degree of size tuning. The results of this analysis are shown in Figure 6.5. We plot the tuning curves of firing rates for stimuli of 20% contrast

(Figure 6.5A) and 100% contrast (Figure 6.5B). At 20% contrast, V1 responses (Figure 6.5A, *dashed curve*) are slightly more selective for size than LGN responses (Figure 6.5A, *solid curve*). At 100% contrast, size tuning is more pronounced in both areas.

In fact, factoring out the effects noise and spike threshold yields normalized tuning curves for LGN and V1 that are nearly undistinguishable. We calculate average tuning curves from the generator potential inferred by the model. This measure isolates the contribution of gain control from that of spike threshold. At low contrast (Figure 6.5C), V1 responses are slightly more tuned than LGN responses. At 100% contrast (Figure 6.5D), however, the curves are nearly identical indicating that V1 neurons exhibit on average the same strength of tuning as LGN neurons.

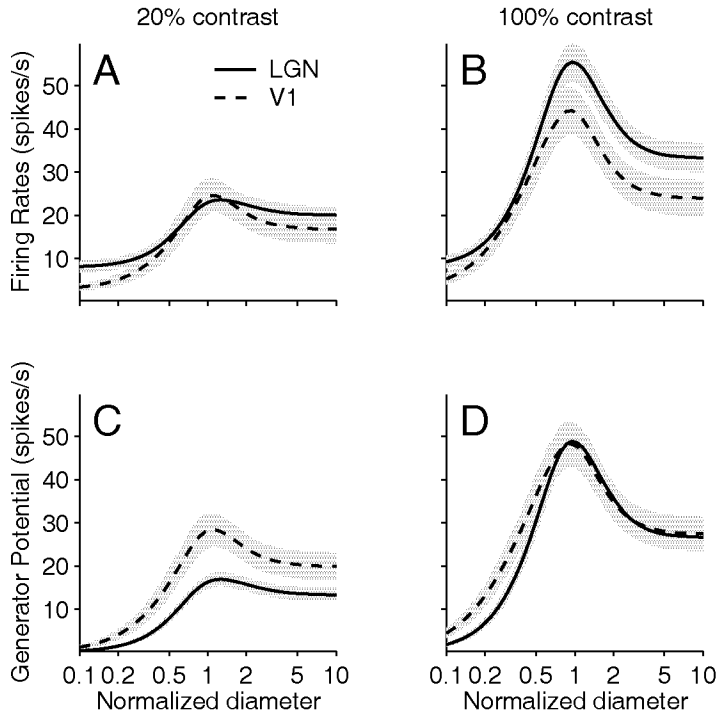

Figure 6.5. Average size tuning curves of LGN and V1 neurons. *Solid curves* indicate average predicted response of LGN neurons as function of stimulus diameter. *Dashed curves* are average curves for V1 neurons. *Shaded areas* indicated ±1 s.e. Diameters are normalized so that responses of individual neurons peak at value of 1. Data points (which would fall on the curves) cannot be shown because normalization results in uneven sampling of the diameter axis. **A**: Predicted firing rate responses to stimuli of 20% contrast. **B**: Predicted firing rate responses to stimuli of 100% contrast. **C**: Generator potential response at 20% contrast. **D**: Generator potential response at 100% contrast.

Next we consider how size tuning is expressed in the responses of individual neurons. Although LGN and V1 neurons exhibit on average a similar degree of size tuning, the strength of size tuning can vary substantially from one neuron to another (Jones et al., 2000; Ozeki et al., 2004). These variations might be more pronounced in V1 than in LGN. We quantify the tuning of individual neurons with a size tuning index. We calculate the difference between peak and plateau responses and express the result as a percentage of the peak response. An index of 100% indicates complete suppression of responses for stimuli of large diameters. An index of 0% indicates no selectivity for stimulus size. We calculate indices for all contrast levels. We do so for the measured responses and for the model predictions.

Consistent with the average tuning curves described above, size tuning is weak at low contrast and pronounced at high contrast. The histograms in Figure 6.6 show the distribution of tuning indices obtained for the measured responses. We plot the tuning indices for the 20% contrast stimuli (Figure 6.6AC), and for the 100% contrast stimuli (Figure 6.6BD). At 20% contrast, size tuning is weak in both LGN (Figure 6.6A) and V1 neurons (Figure 6.6C) with average tuning indices of 23±26% and 40±34% (mean±s.d.). At 100% contrast, LGN neurons show a moderate degree of size tuning (Figure 6.6B) with an average index of 44±18%. By comparison, V1 neurons show more pronounced tuning (Figure 6.6D) with an average index of 60±25%.

The main difference between the tuning of LGN and V1 neurons, however, lies in the spread of the tuning indices for the high contrast condition. In LGN neurons, the distribution of tuning indices is narrow and shows a sharp peak at ~38% (Figure 6.6B). In V1 neurons, indices are nearly uniformly distributed over the range 25-100% (Figure 6.6D). This difference is visible at high contrast (Figure 6.6BD) but not at low contrast (Figure 6.6AC).

The model captures these effects. The scatter plots in Figure 6.6 show measured versus predicted tuning indices for two contrast conditions. At low contrast (Figure 6.6E), the data points show much scatter, presumably reflecting the low reliability of responses. They also often fall below the diagonal indicating that the model tends to underestimate the degree of size tuning at low contrast. At high contrast (Figure 6.6F), the data points fall near the diagonal indicating that the model captures the strength of size tuning.

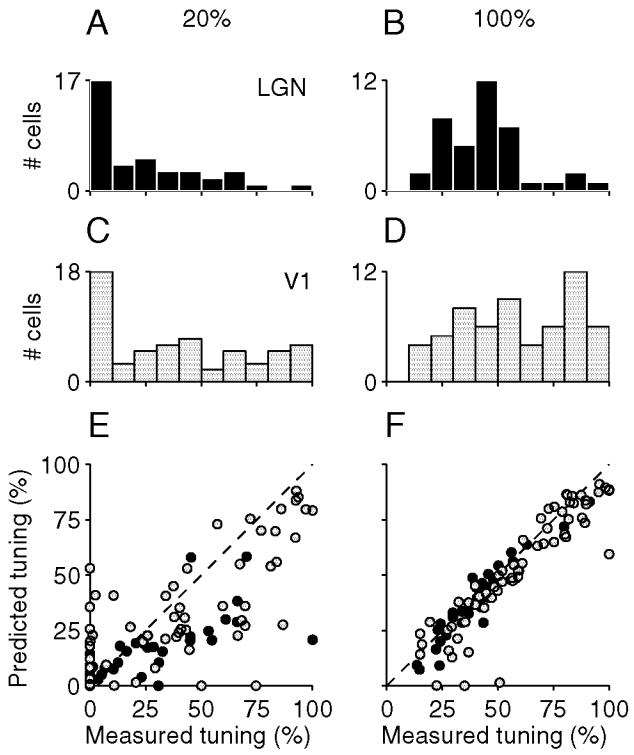

Figure 6.6. Size tuning in individual LGN and V1 neurons. *Histograms* show distribution of tuning indices estimated from the measures responses. *Scatters plots* show correspondence between measured and predicted tuning indices. **A**: Measured indices in LGN neurons for stimuli of 20% contrast. **B**: Measured indices in same population for the 100% contrast condition. **C-D**: Same as A-B for population of V1 neurons. **E**: Measured versus predicted indices at 20% contrast. **F**: Same for the 100% contrast condition.

According to the model, the strength of size tuning may be influenced by two factors. The first factor is contrast gain control, which give the generator potential responses their tuning. The second factor is the spike threshold, which hides the responses of lower amplitudes.

Two quantities together determine the size tuning of potentials (Figure 6.7). These are the strength of gain control c_{50} and the relative size of the suppressive field w_S/w_I (i.e. the ratio of the size of suppressive field over that of integration field). The contour plot in Figure 6.7A depicts their effect on the predicted size tuning index. Decreasing c_{50} enhances size tuning because it causes the output of the suppressive O_S to dominate the divisive term. Increasing w_S/w_I enhances size tuning because it decreases the overlap between integration field and suppressive field.

Where do LGN and V1 neurons lie in the plane spanned by c_{50} and w_S/w_I? The values from model fits are shown in Figure 6.7B and Figure 6.7C (*symbols*). The relative size of the suppressive field is similar in both areas (medians = 3.1±0.6 in LGN and 2.9±0.6 in V1, bootstrap estimates). But V1 neurons exhibit stronger gain control than LGN neurons. They have low c_{50}, typically in the range 1-20%, significantly lower than the 10-100% of LGN neurons. Should not the enhanced gain control of V1 neurons result in stronger size tuning of responses? Not necessarily. For small values of c_{50} (strong gain control) the strength of size tuning is largely determined by w_S/w_I. This dependence falls off for large values of c_{50} (weak gain control) where the degree of size tuning depends increasingly on c_{50}. But there are few neurons in this regime.

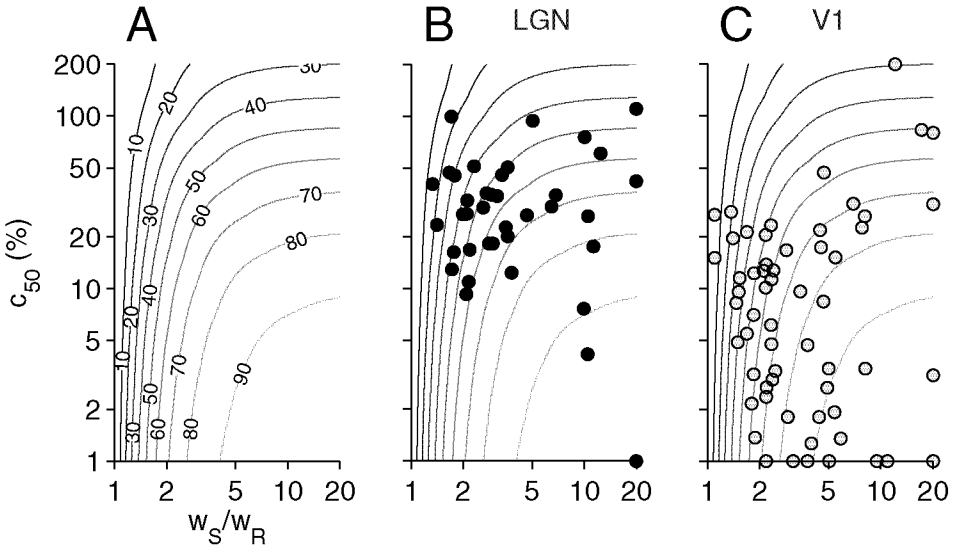

Figure 6.7. Model parameters determining the tuning of potentials. **A**: *Contour plot* shows predicted size tuning index as function of strength of gain control c_{50} and relative size of suppressive field w_S/w_I. **B**: Same, with *closed symbols* indicating the location of individual LGN neurons in the space spanned by c_{50} and w_S/w_I. **C**: Same, with *open symbols* indicating the location of V1 neurons in the space.

The second factor determining the tuning of firing rates is the spike threshold. As we pointed out earlier, V1 neurons have higher threshold than LGN neurons. This difference in threshold might explain the differences observed in the distribution of tuning indices (Figure 6.6BD).

In fact, the model predicts that threshold should strongly affect the size tuning of firing rate responses. This effect is illustrated in Figure 6.8 where we plot predicted tuning index as function of c_{50} and w_S/w_I. We plot indices for the generator potentials (Figure 6.8A), for the firing rates of a low-threshold neuron (Figure 6.8B), and for the firing rates of a high-threshold neuron (Figure 6.8C). For given values of c_{50} and w_S/w_I (Figure 6.8A), a low threshold reduces the apparent strength of size tuning (Figure 6.8B) because it increases responses by a constant factor (the spontaneous firing rate). A high threshold, instead, increases the strength of tuning (Figure 6.8C) because it eliminates the responses with the lowest amplitudes.

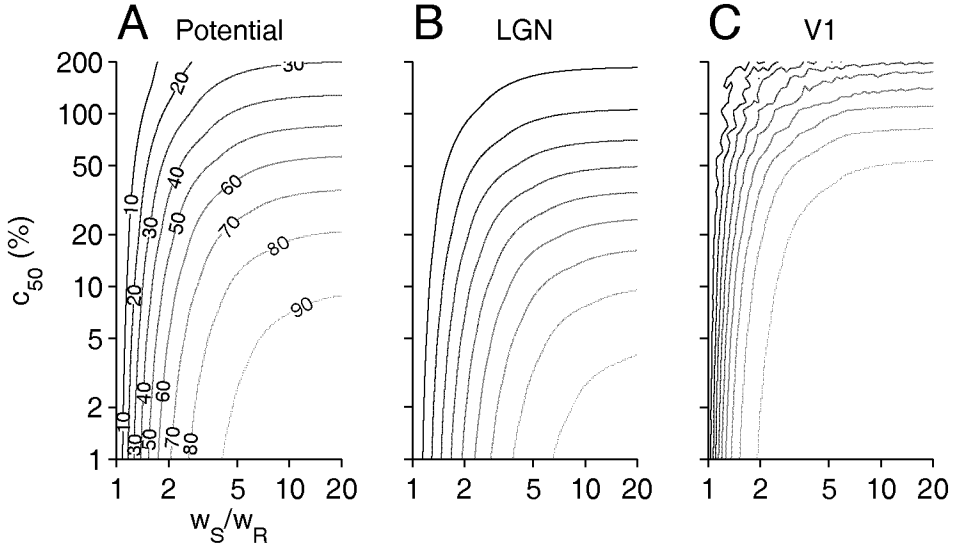

Figure 6.8. Effect of threshold on tuning of firing rates. **A**: *Contour plot* shows tuning of potentials as function of c_{50} and w_S/w_I. **B**: Tuning of firing rates for our median LGN neuron. V_T=-4.3 spikes/s. **C**: Tuning of firing rates for our median V1 neuron. V_T=3.9 spikes/s.

Indeed, the strength of size tuning is strongly correlated with spike threshold. Model fits provide us with an estimate of spike threshold, which we can relate to the strength of size tuning. We calculate the average tuning index obtained from the data at 100% contrast and condition the result on the estimated threshold (Figure 6.9). In both areas, increases in tuning strength coincide with increases in threshold (Figure 6.9A). This dependence is gentle in LGN neurons (Figure 6.9A, *closed symbols*) but strong in V1 neurons (Figure 6.9A, *open symbols*). Increasing threshold increases the strength of tuning because it abolishes

weak responses, a phenomena commonly referred to as the iceberg effect. Furthermore, for a given value of threshold, neurons from both areas have similar tuning indices.

But do model fits provide correct estimates of threshold? The spike threshold can also be estimated from the spontaneous firing rate of the neurons (Figure 6.9B). For low threshold values, spontaneous activity is high and inversely related to threshold. For high threshold values, the relationship between spontaneous activity and threshold falls apart, and spontaneous activity is no longer a good predictor of threshold. Interestingly, the size tuning index continues to increase for these high threshold values.

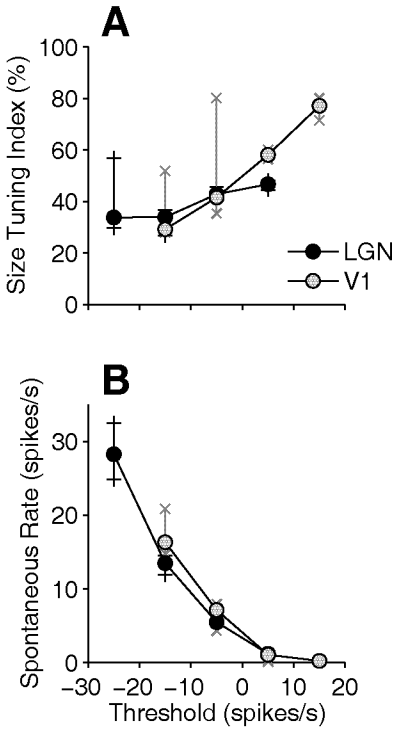

Figure 6.9. Tuning strength, threshold and spontaneous activity. **A**: *Symbols* indicate median tuning index as function of spike threshold for LGN neurons (*closed symbols*) and V1 neurons (*open symbols*). *Error bars* indicate 50 percentiles. **B**: Symbols indicate median spontaneous firing rate of the neurons as function of threshold.

The contribution of gain control to size tuning in individual neurons can be quantified from the generator potentials. We calculate tuning indices on generator potentials and look at their distributions (Figure 6.10). In LGN neurons, the distribution of tuning indices obtained for the generator potentials (Figure 6.10A) resemble that obtained for firing rates

(Figure 6.6B). On average, the strength of tuning is 47+/-21% in potentials and 44+/-18% in firing rates. In V1 neurons the tuning of potentials (Figure 6.10B) is much lower than that of firing rates (Figure 6.6D) averaging 47+/-24% in potentials and 60+/-25% in firing rates.

In fact, the differences between LGN and V1 neurons observed in the tuning of firing rates vanish when looking at the potentials. The distributions of tuning indices in LGN and V1 are statistically indistinguishable when calculated on potentials (two-sample Kolmogorov-Smirnov, p=0.8).

The contribution of threshold can be quantified by comparing the tuning of potentials to that of firing rates. We plot the tuning indices of firing rates as function those obtained for the potentials (Figure 6.10C). In LGN neurons firing rates and potentials have similar tuning indicating that threshold plays a minimal role. In these neurons the ratios of tuning of firing rates over those of potentials averages 101+/- 5% (mean±s.e.). In V1 neurons, firing rates are more tuned than potentials, with an average ratios of 129+/- 5%. This confirms that threshold plays an important role in determining the size tuning of V1 responses.

In summary, we have shown that LGN and V1 neurons exhibit different degrees of gain control but similar degrees of size tuning, once the effects of spike threshold are factored out. This finding holds on average, and from cell to cell.

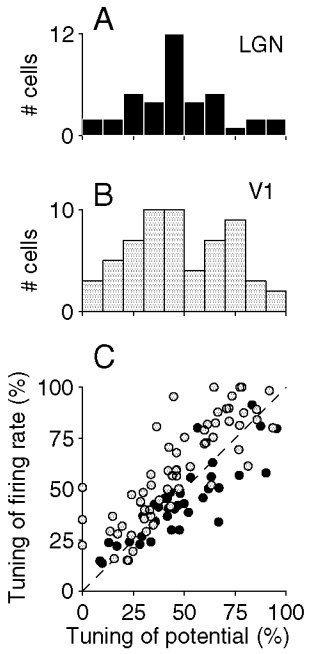

Figure 6.10. Size tuning indices of firing rates and model potentials. **A**: Distribution of tuning indices estimated from generator potential responses for LGN neurons. **B**: Same, for V1 neurons. **C**: *Scatters plots* show tuning indices of firing rates as function of indices for potentials.

6.2.3 Spatial extent of size tuning

We now consider the spatial extents of the integration field w_I and of the suppressive field w_S as obtained from model fits (Figure 6.11). These measures, together with the constant c_{50}, determine the shape of the size tuning curve. We compare the extents obtained for our populations of LGN and the V1 neurons.

The integration field is ~70% larger in V1 neurons than in LGN (Figure 6.11A). Its average width is 0.37±0.05 deg in LGN (N=39, mean±s.e.) and 0.53±0.03 in V1 (N=60). These measures, however, depend slightly on receptive field eccentricity (Figure 6.11A, *thick lines*). To test for possible effects of eccentricity, we calculate the average width for neurons with eccentricities in the range 5-20 deg (Figure 6.11A, *shaded area*), where our populations of LGN and V1 neurons overlap most. For this range, the width of the integration field averages 0.30±0.06 in LGN (N=17) and 0.52±0.08 in V1 (N=34).

The suppressive field is ~50% larger in V1 than in LGN (Figure 6.11B). Its average width is 1.09±0.25 deg in LGN (N=39, mean±s.e.) and 1.54±0.12 in V1 (N=60). These measures depend more strongly on eccentricity (Figure 6.11B, *thick lines*). To test for

possible effects of eccentricity, we calculate the average width for neurons with eccentricities in the (Figure 6.11A, *shaded area*), where our populations of LGN and V1 neurons overlap most. For neurons with eccentricities in the range 5-20 deg the width suppressive field averages 1.05±0.19 in LGN (N=17) and 1.65±0.19 in V1 (N=34).

In LGN and V1, suppressive field is about 3 times the size of the integration fields. In LGN, the suppressive field is on average 3.1+/-0.6 times the size of the integration field. In V1 this ratio is 2.9+/-0.6.

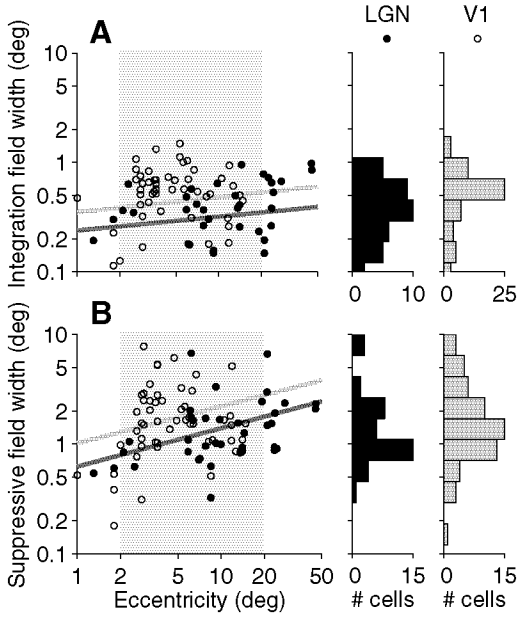

Figure 6.11. Spatial extent of the integration field and the suppressive field. **A**: *Scatter plot* shows width of integration field w_I as a function of receptive field eccentricity for LGN neurons (*closed symbols*) and V1 neurons (*open symbols*). *Histograms* depict distributions of field widths. *Thick lines* indicate results of linear regression (slopes = 0.13 and 0.14 on a double logarithmic axis). *Shaded area* indicate region where populations of LGN and V1 neurons mostly overlap. **B**: Same, for the suppressive field (regression slopes = 0.35 and 0.35).

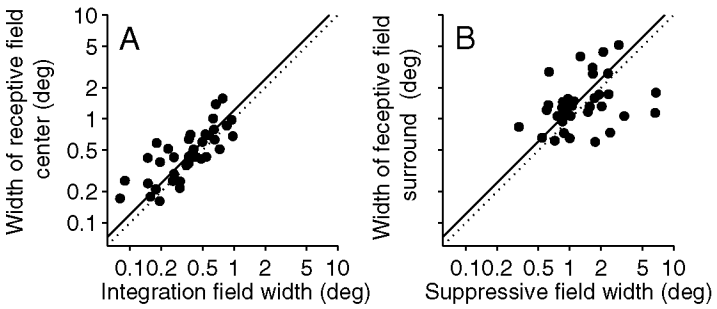

Figure 6.12. Relation to linear receptive field. A: Width of receptive field center vs. that of integration field. B: Width of receptive field surround vs. that of suppressive field.

How do the integration field and the suppressive field relate to the receptive field of the neurons? Recall that the integration field is the region of visual field providing the bulk of the drive to the neuron. Because stimuli had optimal spatial frequency, we would expect the integration field of LGN neurons to be about the same size as the receptive field center. By comparison, the suppressive field is the region over which gain control computes local contrast. We have shown in Chapter 3 the suppressive field to be about the size of the receptive field surround.

To confirm these suggestions, we estimate the receptive fields of the LGN neurons from responses to moving gratings of different spatial frequencies. We fit the responses to a difference-of-Gaussians receptive field and compare the results to our measures of integration field and suppressive field widths.

As expected we find that the integration field of LGN neurons is about the same size as the receptive field center (Figure 6.12A). We plot the width of the receptive field center versus that of the integration field. The data points fall near the diagonal line indicating that the two are similar. On average the receptive field center is 1.21+/-0.13 that of the integration field (median, bootstrap estimate, N=39).

Consistent with the results presented in Chapter 3, we find that the suppressive field of LGN neurons is about the same size as the receptive field surround (Figure 6.12B). We plot the width of the receptive field surround versus that of the integration field. Data points fall again near the diagonal line. On average the receptive field surround is 1.22+/-0.14 that of the suppressive field.

6.3 Discussion

We have studied the behavior of the contrast-response curves of LGN and V1 neurons. At low contrast, V1 responses grow on average more steeply with contrast than LGN responses (Figure 6.3). This effect is relatively independent of the diameter of the stimulus. At high contrast, V1 responses saturate more than LGN responses (Figure 6.2). This difference in saturation is more pronounced for small stimuli.

These differences in steepness and saturation reflect differences in gain control and spike threshold but not in the trial-to-trial variability of responses. For the range of parameters observed, only the differences in strength of gain control c_{50} and threshold V_T had notable effect of the shape of the average contrast-response curves (Figure 6.4). Surprisingly, the effects of spike threshold actually counteract the effects of gain control. For a fixed value of c_{50}, increasing threshold V_T made the contrast-response curves grow more linearly with contrast (Figure 6.4CD).

The differences in gain control observed between LGN and V1 may result from differences in the temporal frequency of the stimulus. We employed stimuli with optimal temporal frequencies to obtain robust responses from the neurons. V1 neurons, however, prefer lower temporal frequencies than LGN neurons (Ikeda and Wright, 1975; Movshon et al., 1978; Saul and Humphrey, 1990; DeAngelis et al., 1993b). In our population of V1 neurons, the median optimal temporal frequency is 3.1±0.4 Hz (bootstrap estimates). In our population of LGN neurons, the preferred temporal frequency is twice as high (7.4±0.9 Hz). As we explained in Chapter 2, the effects of contrast gain control are strongest for low temporal frequencies (Shapley and Victor, 1978b; Shapley and Victor, 1981). If indeed the effects observed in V1 responses reflect a retinal or thalamic gain control mechanism, then the differences in gain control might be explained by the differences in temporal frequencies of the stimulus.

We investigated the strength of size tuning and found LGN and V1 neurons to exhibit similar degrees of size tuning. Consistent with previous studies (Jones et al., 2000; Ozeki et al., 2004), we found that LGN and V1 neurons show distinct strength of size tuning when considering the tuning of firing rate responses (Figure 6.5AB and Figure 6.6). These differences, however, vanish when considering the tuning of generator potential responses (Figure 6.5AB and Figure 6.10AB). This suggested an important role for the spike

threshold. We could quantify this role using model fits because they provide us with estimates of the spike threshold. We found that, for a given threshold, LGN and V1 responses exhibited nearly identical degrees of tuning. We also found that threshold affects the tuning of V1 firing rates but not that of LGN responses (Figure 6.9A). This is because thresholds in V1 neurons generally have positive values (i.e. near zero spontaneous firing rate) whereas thresholds in LGN are negative (i.e. positive spontaneous firing rate). For positive values, increasing threshold increases the tuning of responses (Figure 6.10C, *open symbols*). For negative values, varying thresholds has little effect (Figure 6.10C, *closed symbols*).

An important feature of the model is its ability to infer information about subthreshold responses from the extracellular responses. Traditionally the spike threshold is estimated from the spontaneous firing rate of the neurons. For high threshold values, the relationship between spontaneous activity and threshold falls apart, and spontaneous activity is no longer a good predictor of threshold. That the size tuning index continues to increase for these high threshold values suggest that the model has successfully inferred threshold (Figure 6.9B).

We investigated the spatial extent of size tuning and found its scale to be constant relative to that of the integration field. At first glance, this result appears to support previous work suggesting the spatial extent of size tuning to be inconsistent with a thalamic origin (Angelucci et al., 2002). From a model in which V1 neurons linearly sum responses of multiple LGN neurons with spatially-displaced receptive fields, one expects the relative size of the suppressive field to be lower in V1 than in LGN. However, while in LGN the size of the integration field matches that of the receptive field (Figure 6.12), we cannot be sure this to be case in V1.

In conclusion, our findings indicate that contrast gain control is likely to contribute substantially to the suppressive phenomena of V1 responses. Quantifying this contribution, however, requires a more complete model of V1 responses that includes realistic inputs.

6.4 Methods

Methods for physiological recordings, stimulation and data analysis are described in the in General Methods. The experiments, the data set and the model are described in Chapter 5.

Chapter 7

Conclusions

In this thesis, we have investigated neural responses at two consecutive stages of the cat visual system, the lateral geniculate nucleus (LGN) and the primary visual cortex (V1). Our approach has been to devise simple models of the responses and to validate them on large sets of responses.

First, we have shown that three suppressive phenomena shaping LGN responses can be entirely accounted for by a mechanism of contrast gain control. We devised a model in which the output of the receptive field is divided by the output of a suppressive field, thereby controlling the gain of the neurons. The model quantitatively predicts (1) the phenomena of saturation of responses with contrast, (2) the masking of responses to a given stimulus by superposition of a second and (3) the size tuning of responses for stimulus size.

Second, we have studied the visual preferences of the suppressive field of LGN neurons and have shown that these properties are consistent with an origin in retina and thalamus but are inconsistent with an origin in feedback from cortex. The suppressive field is no larger than the receptive field surround. It is not selective for stimulus orientation, and it responds to a wide range of frequencies, including very low spatial frequencies and high temporal frequencies. Together these results provide a thorough characterization of contrast gain control in LGN.

Third, we have shown the suppressive field to compute the true standard deviation of the light intensities (i.e. root-mean-square contrast). We used white noise stimuli to measure the gain of LGN neurons and asked how gain is affected by the luminance statistics of the stimulus. We found that neural gain is strongly influenced by the standard

deviation but mostly invariant to skewness and kurtosis. These results validate the suppressive field model for a broad range of luminance statistics.

Fourth, we have devised and validated a probabilistic model of LGN and V1 responses. The model includes an integration field, a divisive suppressive field, a noise source and a rectification stage. The model accurately predicts the mean responses and also captures their trial-to-trial statistics. The model provides a tool to compare responses across brain areas.

Fifth, we have compared the responses of LGN and V1 neurons to stimuli of different contrasts and sizes, teasing apart the potential contributions of gain control, noise and spike threshold. Responses in V1 show more saturation with increasing contrast than responses in LGN; they exhibit distinct patterns of size tuning; they show more trial-to-trial variability. The contrast-response curves of V1 show indications of enhanced gain control whose effects are assuaged by the spike threshold. The size tuning curves in both areas are strikingly similar once the spatial scales are matched, and once the effects of spike threshold are accounted for. Overall, the differences in trial-to-trial variability appear too small to significantly affect the average response curves. These results indicate that the effects of gain control in the LGN input are likely to contribute substantially to the effects observed in the V1 responses.

7.1 Benefits of parsimonious modeling

The models have played a central role in this thesis, influencing the research at multiple levels. We have gone back and forth between levels so it is useful to make these explicit. The models have supported our research in at least three ways. First, they provided us with a robust and tractable description of large sets of responses, thereby capturing the computations performed by complex neural circuits. Second, they provided us with functional interpretation of the responses giving us some insights as to why some phenomena may arise. Finally, they revealed possible constraints about the underlying neural circuits.

The models introduced in this thesis provide parsimonious and accurate descriptions of large sets of responses. We took care in devising the simplest models that could explain the data. With only eleven free parameters, the suppressive field model introduced in Chapter 2

could accurately describe the responses of LGN neurons to hundreds of stimuli. With only seven free parameters, the compound model introduced in Chapter 5 could describe the responses of LGN and V1 neurons to >50 stimuli. In that case the model not only captured the average responses but also described their trial-to-trial statistics.

The main benefit of parsimonious modeling is tractability. By keeping the models simple we made it easier to estimate their parameters. We could estimate the parameters directly from the data and did not have to make educated guesses, as is often the case in the field (e.g. Heeger, 1992; Somers et al., 1998; Rao and Ballard, 1999; Schwartz and Simoncelli, 2001; Miller and Troyer, 2002). By keeping the models simple we allowed for the effects of individual components to be assessed. For example we could tease apart the contributions of gain control, noise in membrane potential and spike thresholds to the responses. By keeping the models simple we reduced (though not to zero) the likelihood that the effects of model parameters trade with each other. This allowed the use of model parameters for empirically describing the data.

Throughout this thesis we have sought to explain a variety of observations under the light of divisive normalization. A likely function of contrast gain control is to map the wide range of contrasts encountered in the natural environment onto the dynamic range of the neurons. Divisive normalization is a strategy to achieve this goal (Heeger, 1992). Several aspects of our results support this hypothesis. First and foremost, we have found the suppressive field to compute the true root-mean-square contrast of the stimulus. The suppressive field is largely insensitive to the skewness and the kurtosis of the stimulus. It is not selective for stimulus orientation. It responds to both spatial contrast and temporal contrast (i.e. low spatial frequencies). Second, the effects of gain control fall off at low contrast where the dynamic range of responses is no longer a limiting factor. Third, the suppressive field and the receptive field are coextensive. Had the suppressive field been much smaller or much bigger, it could not have adequately normalized the responses. Note that this framework has limitations. For example, contrast gain control also affects the dynamics of the responses (Shapley and Victor, 1978a; Shapley and Victor, 1981). These effects are more difficult to explain by the normalization framework.

Finally, parsimonious modeling has also revealed constraints about the physiological origins of the suppressive phenomena. We have argued in Chapter 3 that the visual preferences of the suppressive field are consistent with an origin in retina or thalamus but

not with an origin in feedback from cortex. The model suggested the masking experiments as a method to measure these properties. The model also yielded estimates of the properties of the receptive field. In Chapter 5 and 6 we have measured the contributions of gain control, noise in membrane potential and spike threshold to suppression in V1. Again, the model provided the framework to tell these apart.

7.2 Limitations of the approach

Despite its many benefits, there are also potential pitfalls to parsimonious modeling, the main of which is oversimplification. The models presented here also have significant limitations.

We have studied the effects of contrast gain control but have neglected other adaptation mechanisms. For example, we have not considered light adaptation (Shapley and Enroth-Cugell, 1984), which adapts responses to the mean light intensity of the stimulus. In a related project (Mante, 2005), we have shown the effects of light adaptation and contrast gain control to be functionally separable, indicating that they can safely be investigated separately. Another mechanism that has received considerable attention is slow contrast adaptation (Sanchez-Vives et al., 2000; Chander and Chichilnisky, 2001; Kim and Rieke, 2001; Baccus and Meister, 2002; Solomon et al., 2004b). This mechanism resembles the mechanism of contrast gain control but operates over long time scales (i.e. tens of seconds). By comparison, the effects of contrast gain control are nearly instantaneous and settle within tens of milliseconds after the onset of the responses (Victor, 1987; Albrecht et al., 2002; Baccus and Meister, 2002). Since the eyes typically fixate for only 200-300 ms, the role of slow contrast adaptation in natural vision is unclear. In most of this thesis, we have minimized the effects of slow adaptation by presenting stimuli of short durations. In Chapter 4, however, we used stimuli that lasted ~30s and therefore we could not distinguish the effects of gain control from those of slow adaptation. Finally, we have not considered the mechanism of cortical adaptation, whose properties have been thoroughly investigated elsewhere (Maffei et al., 1973; Albrecht et al., 1984; Ohzawa et al., 1985; Carandini and Ferster, 1997; Sanchez-Vives et al., 2000).

In this work we have studied the amplitudes of responses but have not considered the phase and the time course of the responses. In addition to affecting the amplitudes of responses, contrast gain control also affects their time courses. Increasing contrast reduces

the integration time of the neurons so that changes in response amplitudes are accompanied by changes in phase (Shapley and Victor, 1978b; Shapley and Victor, 1981). Moreover, the effects of gain control are also pronounced low temporal frequencies and weak at high temporal frequencies so that increasing contrast enhances responses to high temporal frequencies. At low temporal frequencies, the effects of varying stimulus contrast on response phase go hand in hand with those on response gain (Shapley and Victor, 1981; Victor, 1987; Baccus and Meister, 2002). In a related project (Mante, 2005), we have shown that increasing stimulus size results in changes in gain and phase similar to those of increasing contrast. This close correspondence at low frequencies allowed us to thoroughly characterize the properties of gain control from the response amplitudes only.

Another aspect of LGN responses we have not considered is the nonlinear responses of Y cells. In Chapters 2 and 3 we have studied the component of responses oscillating at the temporal frequency of the stimulus. This harmonic component accounts for most of the variance in the responses of X cells. The responses of Y cells, however, exhibit an additional nonlinear response component consisting in an elevated mean firing rate in response to moving gratings (Enroth-Cugell and Robson, 1966; Hochstein and Shapley, 1976; Derrington et al., 1979; Demb et al., 1999). We did not consider the nonlinear components of the responses and nevertheless found the model to accurately predict the harmonic responses of both X and Y cells. This suggests that the linear and nonlinear responses of Y cells could be modeled by separate modules. In fact, the subunits underlying the nonlinear responses of Y cells exhibit properties that closely resemble the properties of contrast gain control. They respond equally to light increments and decrements and are largely insensitive to stimulus phase. They are broadly selective for spatial frequency and respond to contrast falling over a large region of the visual field. These properties led Shapley and Victor to suggest that the nonlinear responses of Y cells and of contrast gain control might share a common origin. An intriguing possibility is whether a nonlinear pathway such as provided by the suppressive field could also be used to model the nonlinear responses of Y cells.

A limitation of our models relates to the encoding of firing rates. In this thesis we have assumed firing rates to be linearly related to the membrane potential of the neuron. Previous work has shown this to be a good approximation for responses in retina (Chichilnisky, 2001; Kim and Rieke, 2001; Zaghloul et al., 2003) and in primary visual

cortex (Carandini and Ferster, 2000). LGN neurons, however, exhibit distinct firing modes that are not captured by a linear encoder (Jahnsen and Llinas, 1984a, b; McCormick and Feeser, 1990; Sherman, 2001b): a tonic mode in which firing rate is linearly related to the synaptic input and a burst mode in which neurons respond in an all or nothing fashion. These modes have been associated with the state of wakefulness of the animal (Weyand et al., 2001). Tonic spikes are generally associated with the wakeful state whereas bursts are mostly observed during sleep and under anesthesia. However, we have reasons to think that bursts did not overly influence our results. First, when considering the harmonic response to moving gratings (Chapters 2-3), bursts mostly affect response phase but have little effects on response amplitude (Mukherjee and Kaplan, 1995; Mante, 2005). Second, when considering the mean firing rate responses to these stimuli (Chapters 5-6) bursts normally represent only a small fraction of the total number of spikes (Weyand et al., 2001). Finally, when stimulating with white noise (Chapter 4) the presence of bursts has little effect on the impulse response estimated with reverse correlation (Reinagel et al., 1999). Recent studies suggest that bursts may play a more important role with natural stimuli (Lesica and Stanley, 2004; Denning and Reinagel, 2005).

7.3 Future directions

An open question regards the role of thalamic inhibition in visual processing (Levick et al., 1972; Uhlrich et al., 1991; Blitz and Regehr, 2005). In our study of the suppressive phenomena, we could not distinguish the contributions of thalamic inhibition from those of the retinal input. Ganglion cells and LGN neurons have similar visual preferences so we could not use these properties to tell them apart.

One approach to assess the role of thalamic inhibition would be to look for s-potentials (So and Shapley, 1981; Cleland et al., 1983a; Lee et al., 1983; Kaplan and Shapley, 1984; Einevoll and Plesser, 2002). The s-potentials are low amplitude deflections some times observed in extracellular recordings of LGN neurons. These deflections are thought to reflect retinal spikes that have failed to elicit spikes in the LGN neuron. We could use the models devised in this thesis to quantitatively compare the responses of LGN neurons to their retinal input. We could ask whether thalamic inhibition indeed enhances gain control or whether its effects serve a different function.

This approach, however, would not distinguish between feedforward inhibition from interneurons and feedback inhibition from the perigeniculate nucleus (PGN). To test for a role of the PGN one may want to perform experiments with a test grating in one eye and a mask grating in the other eye. Indeed, PGN neurons can be driven from both eyes (Sanderson, 1971; Ahlsén et al., 1983; Xue et al., 1988; Uhlrich et al., 1991; Funke and Eysel, 1998), and have been linked to binocular suppressive effects in LGN (Sanderson et al., 1971; Pape and Eysel, 1986; Wang et al., 1994; Funke and Eysel, 1998). Again one could use the models developed to test the function of interocular suppression.

Another unresolved issue if the function of cortical feedback (Sillito and Jones, 2002; Worgotter et al., 2002; Alitto and Usrey, 2003). Several studies have sought to compare LGN responses with and without cortical feedback (Murphy and Sillito, 1987; Przybyszewski et al., 2000; Alitto et al., 2002; Sillito and Jones, 2002; Webb et al., 2002; Worgotter et al., 2002; Alitto and Usrey, 2003). These studies are difficult to interpret because of the lack of a null hypothesis as to what the behavior of an intact LGN neuron should be. Our models provide such a null hypothesis. They would help to distinguish the different effects of cortical feedback. They would distinguish overall excitatory effects from those of a thalamocortical gain control mechanism or from any other mechanism.

In conclusion, this research brings us closer to a central goal of visual neuroscience, the development of functional models of neural responses to visual stimuli. We have demonstrated that functional models can not only summarize the image processing performed by complex neural circuits but can also guide research in the underlying biophysics, and facilitate the study of subsequent visual stages. Functional models are particularly desirable for LGN, whose responses reflect the computations performed in retina, and constitute the principal input to the visual cortex. These models may also one day leave the laboratory and be implemented into what would constitute the pre-processing stage of thalamic visual prostheses (Margalit and Sadda, 2003; Pezaris and Reid, 2004).

Chapter 8

General Methods

8.1 Physiological Experiments

All procedures were approved by the Veterinary Office of Canton Zurich and by the Animal Care and Use Committee of the Smith-Kettlewell Eye Research Institute.

8.1.1 Pre-operative preparation

In the morning of the experiment the animal receives an intramuscular injection of a mixture of sedative and anesthetic (Ketamine 20-24 mg/kg i/m + Xylazine 1.0-1.2 mg/kg i/m) (Flecknell, 1996). This initial anesthesia is meant to last about 30-45 min, the time usually needed to perform the initial surgery and move on to the steady anesthesia regime; additional doses are given when necessary. An anticholinergic, atropine sulfate (0.04-0.06 mg/kg, i/m), is administered intramuscularly as salivary and bronchial secretion may interfere with respiration. This administration is repeated daily for the duration of the experiment.

8.1.2 Initial Surgery

(Flecknell, 1996)To be able to continuously and securely deliver anesthetics and other liquids intravenously, we place a catheter in the saphenous vein of each leg. We make a small incision in the skin over the vein, we isolate the vein trough gentle blunt dissection of the overlying fascia, we make a small cut in the vein using microscissors, we insert an appropriate cannula and we tie the cannula to the vein.

To be able to artificially respirate the animal and to monitor the end-tidal CO_2 output of the respiratory system we place a cannula in the trachea. We perform a standard tracheotomy, insert an appropriately sized cannula and we secure the cannula to the trachea with ligatures.

Once the first catheter is inserted into a vein, anesthesia is continued with a loading dose of Thiopental Sodium (Penthotal, 10-15 mg/kg i/v) (Flecknell, 1996).

8.1.3 Neurosurgery

After these simple surgical procedures the animal is moved a few feet to the experiment area. This area contains a stereotactic apparatus set on a vibration isolation table. These devices ensure maximal stability and allow precise placement of microelectrodes (for electrical recordings), and of a high-performance camera (for optical recordings). A respiration pump and injection pumps for syringes are located adjacent to this equipment.

In the experiment area we provide continuous infusion of barbiturate anesthetic, Thiopental Sodium (Pentothal, 0.5-4 mg/kg/hr, i/v) (Gillespie et al., 2001; Kara et al., 2002; Martinez et al., 2002; Thompson et al., 2003). We enhance the analgesic effect of the barbiturate with additional inhalation of nitrous oxide (N_2O) mixed with oxygen (O_2) in concentration of commonly 50% to 70% (Hammond, 1978; Hikasa et al., 1997).

Rate of barbiturate infusion needs to be adjusted for each animal, and greatly depends on the proportion of body fat. Fat tissues absorb barbiturates and release them slowly over an extended period. Infusion rate, thus, must be progressively decreased. Infusion rate is thus determined by monitoring the depth of anesthesia and by an estimate of lean body mass, i.e. the body weight minus the estimated proportion of fat tissues (Flecknell, 1996).

The pupils are dilated and accommodation paralyzed with topical atropine (drops). The nictitating membranes are pharmacologically retracted with phenylephrine (drops). The corneas are protected with contact lenses.

We employ standard neurosurgical techniques to expose a piece of the brain. As is common during neurosurgery, the head of the animal is fixed to the stereotactic apparatus. Contact points are first sprayed with local anesthetic (Lidocaine, 10 mg/spray). The cranium is then exposed and appropriate electrodes inserted to monitor the electro-encephalogram (EEG). From now on the EEG is continuously monitored to detect possible

signs of arousal (Rampil, 1998). A small craniotomy is performed over the area of the relevant brain structures, and the dura mater is opened to allow the insertion of microelectrodes or the imaging of brain activity. To improve stability and to prevent drying of the exposed regions, the craniotomy is covered with 3-4% agar in saline solution.

After neurosurgery, the animal is given an infusion of muscle relaxants to minimize eye movements (described in a later section).

The animal receives constant fluid replenishment. Fluid balance is maintained by continuous intravenous infusion of replacement-type solutions such as 0.9% NaCl or Normosol-R. Rate is 5-10 ml/kg/hr during the initial surgery and 2-5 ml/kg/hr later, to compensate for insensible loss (Muir, 2000). The bladder is periodically expressed to relieve urine accumulation.

The animal receives periodic administration of antibiotic to prevent possible infections (Cephazolin, 18-22 mg/kg IM, twice daily), and anti-edematic agents that prevent brain swelling (Dexamethasone, 0.3-0.5 mg/kg daily). The animal is periodically massaged to stimulate circulation.

8.1.4 Paralysis

After neurosurgery, the animal is given an infusion of muscle relaxants to minimize eye movements (pancuronium bromide, 0.20-0.40 mg/kg loading dose, 0.10-0.20 mg/kg/hr maintenance dose, in balanced solution). This infusion occurs through the second I/V line. The animal is artificially respirated with the mixture of Oxygen and Nitrous Oxide. End-tidal CO_2 is continuously monitored and maintained close to the physiological value of 32-34 mmHg by manipulation of the respiratory rate or volume.

During experiments, in the presence of muscle relaxants, the EEG constitutes the primary source of information on depth of anesthesia. The ideal depth of anesthesia is one in which activity is synchronized in the low-frequency range (2-10 Hz) and occasionally breaks into high-amplitude, low-frequency spindle episodes typical of sleep. Indeed, in human surgery the accepted standard for assessing the depth of anesthesia is now the Bispectral Index, which is entirely based on the EEG (Sigl and Chamoun, 1994).

8.1.5 Post-operative monitoring

The physiological state is continuously monitored and the depth of anesthesia is maintained in a surgical plane by adjusting the rate of delivery of the anesthetic agents. •

We measure temperature, which is kept at a comfortable 38.5 degrees Celsius through the thermostat.

- We measure heart rate, which is commonly between 120 and 190 beats per minute, but becomes higher under Penthotal anesthesia.

- We measure the electrocardiogram (ECG) waveform to detect possible problems with cardiac activity.

- We measure the electrical resistance of the chest to know the respiration rate. This rate is variable during anesthesia, and when we attach the respirator we set its rate to 20-40 breaths per minute, depending on body weight.

- We measure the saturation of oxygen in the blood (SPO2) using an optical sensor. When the signal is sufficiently strong to get a reading, we commonly observe normal values in the range of 97-99%.

- We measure the maximal end-tidal CO2 content of expired air, and adjust the respiration rate and volume to maintain it around the physiological value of 32-34 mmHg.

- We measure lung pressure to detect possible obstructions in the airway. Reference values are taken at the beginning of the experiment, and subsequent measurements are evaluated on the basis of the reference values.

- We measure the electro-encephalogram (EEG) to know the depth of anesthesia (see below).

During the brief surgical procedures at the beginning of the experiment, surgical plane of anesthesia is assessed with classic measures (Brown, 1994). We test for the absence of jaw tone and of reflex reactions to toe pinch or to outer eye stimulation. Corollary evidence on the depth of anesthesia is also provided by respiration rate and heart rate (Flecknell, 1996). As soon as it becomes available, we correlate these observations with the EEG..

8.1.6 Euthanasia

The duration of the experiments is determined by the general physiological state and in particular by the responsiveness of the brain, which we monitor continuously. When the physiological state deteriorates there is an overall decrease of responsiveness to visual stimulation. We end the experiment after 120 hours or earlier if the quality of the data deteriorates.

At the end of the experiment, while they are still under surgical anesthesia, the animals are euthanized through an overdose of anesthetic (Thiopental Sodium, >100 mg/kg).

8.2 Recordings

For recordings in LGN, a craniotomy was performed above the right LGN (Horsley-Clarke, coordinates ~9 mm lateral and ~6 mm anterior). Electrodes were lowered vertically until visual responses were observed. The location of LGN was determined from the sequence of ocular dominance changes during penetration. Nearly all cells were located either in the first contra-lateral layer (presumably lamina A) or in the first ipsi-lateral layer (presumably lamina A_1). Less than 10% of the cells were located in subsequent layers.

For recordings in V1, a craniotomy was performed above the right A17/18 (Horsley-Clarke, coordinates 0–3 mm lateral and 5–9 mm posterior).

Eyes were mapped on tangential screen by illuminating the back of the eye. Location of blind spot and area centralis were marked.

8.3 Visual stimulation

Visual stimuli were displayed using the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) and presented monocularly on a calibrated monitor with mean luminance of 32 cd/m^2 and refresh rate of 124 Hz. All stimuli were generated using lookup table animation.

Stimuli lasted 1–4 s and were presented in blocks of 3-6 repeats. Stimuli within blocks were presented in randomized order. Each block included one or more blank stimuli to estimate the spontaneous firing of the neuron.

8.4 Data acquisition

Extracellular signals were recorded with Quartz-coated Platinum/Tungsten Microelectrodes (Thomas Recordings), sampled at 12 kHz with a National Instruments data acquisition board, and stored for offline analysis using custom software.

8.5 Basic experiments

Once a neuron with well isolated spikes has been identified, we mapped its receptive field was mapped on a tangential semi-reflective screen.

All experiments included >1 blank stimuli to estimate the spontaneous firing rate. We then ran a series of experiments to estimate the basic properties of the neuron.

- Drifting gratings of 50% contrast, optimal temporal frequency, > 20 deg in diameter and >14 different logarithmically spaced spatial frequencies.
- Drifting gratings of 50% contrast, optimal spatial frequency, > 20 deg in diameter and >10 different logarithmically spaced temporal frequencies.
- Drifting gratings enclosed in small circular apertures (>2 deg) displayed at different spatial locations. The location eliciting maximal responses was selected as the center of the receptive field.

We ran the following additional experiment for each LGN neuron:

- We classified cells into X and Y types from responses to counterphase-modulated sinusoidal gratings of different spatial frequencies (Enroth-Cugell and Robson, 1966; Hochstein and Shapley, 1976). Grating phase was chosen to elicit the smallest first harmonic of responses. Cells are classified as Y type if the amplitude of the second harmonic to high spatial frequencies reliably exceeds that of the first harmonic. Most units encountered were of X type, consistent the known laminar distribution of LGN cells (Wilson et al., 1976). We did not attempt to identify the so-called W cells (Stone and Fukuda, 1974), as they likely include many different cell types (Rodieck et al., 1993; Wassle, 2004; Callaway, 2005) and they are largely confined to lamina C (Wilson et al., 1976).

We ran the following additional experiments for each V1 neuron:

- Drifting gratings of 50% contrast, optimal spatial and temporal frequencies, and 12 evenly spaced orientations.
- To classify V1 neurons into simple and complex, we calculated the ratio of the first harmonic of responses over the mean responses from responses to optimal moving gratings. Cells with ratios ≥ 1.0 were classified as simple while those with ratios <1.0 were classified as complex.

8.6 Data analysis

All data analysis was performed with MATLAB (The Mathworks, Cambridge, MA).

8.6.1 Measure of responses

For experiments with moving gratings, we calculated the kth harmonic components of the spike times relative to the temporal frequencies of the stimuli (Figure 8.1). Let $\{t_j\}$ denote the set of N spike times measured in response to a stimulus of temporal frequency f (Figure 8.1A), the harmonic component is obtained by calculating $r = \frac{1}{N}\sum_j \exp\left[-i \cdot 2\pi k f \cdot t_j\right]$. For k=0, evaluating r yields the mean firing rate (Figure 8.1B, *dashed curve*). For k=1, it yields a complex number, the first harmonic of response (Figure 8.1B, *solid curve*). Together the mean firing rate and first harmonic represent an accurate description of the response (Figure 8.1B, *shaded area*). In this work, we considered the magnitude of r and neglected its phase.

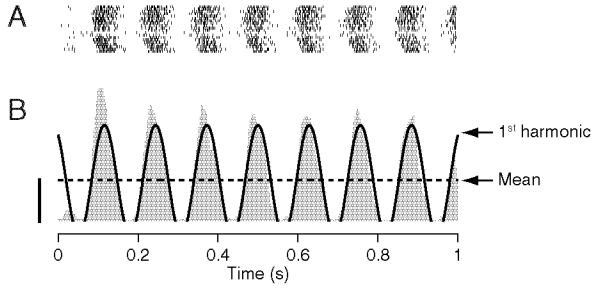

Figure 8.1. Measure of responses. **A**: Response of an LGN neuron to a moving grating of optimal spatial and temporal frequencies. Each dot indicates a spike. Each row corresponds to a single trial. **B**: *Shaded area* indicates corresponding firing rate response (calculated with 10 ms bins). *Dashed curve* indicates mean firing rate. *Solid curve* shows the time course of first harmonic of response. Scale bar indicates 100 spike/s.

8.6.2 Model fits

Fits of the deterministic model minimized square error between measured responses and model predictions. Square error is given by $\sum_{ij}\left(r_{ij}-m_j\right)^2$ where r_{ij} denote the observed response to trial i of stimulus j, and m_i represent the response predicted by the model.

Fits of probabilistic model maximize likelihood of observing the data.

8.6.3 Quality of fits

Percentage of explained variance

To measure the quality of model predictions across stimulus conditions we calculated the percentage of the variance in the averaged data that is explained by the model. Explained variance is given by $100\left(1-n^{-1}\sum_j\left(\bar{r}_j-m_j\right)^2\Big/\mathrm{var}\left[\bar{r}_j\right]\right)$ where $\bar{r}_j$ is the response to stimulus j averaged across trials and n is the number of stimuli.

Percentage of stimulus-driven variance

To quantify how well the model predictions r_t capture the measured responses s_t (both consisting of $t = 1, \ldots, M$ samples) we estimated the fraction of stimulus-driven variance in the responses accounted for by the model (Sahani and Linden, 2003; Machens et al., 2004):

$$\beta=\frac{\sigma_s^2-\sigma_e^2}{\sigma_s^2-\sigma_\eta^2}$$

where:

$$\sigma_s^2=\left\langle\frac{1}{M}\sum_t s_t^2\right\rangle$$

is the power in the response,

$$\sigma_e^2=\left\langle\frac{1}{M}\sum_t\left(r_t-s_t\right)^2\right\rangle$$

is the mean square distance between data and model and

$$\sigma_\eta^2=\frac{d}{d-1}\left[\left\langle\frac{1}{M}\sum_t s_t^2\right\rangle-\frac{1}{M}\sum_t\left\langle s_t^2\right\rangle\right]$$

is an estimate of the variance in s_t that is due to noise. Angular brackets indicate the average over d presentations of the same stimulus.

The quantity β is an intuitive measure of fit quality, similar but superior to the commonly used "percentage of the variance". A perfect model, i.e. a model that is a perfect description of a system, could never account for 100% of the variance, because some of that variance is due to noise. By estimating the fraction of the variance that is due to the noise, the quantity σ_η^2, and by subtracting it from the denominator in the definition of β, one eliminates this contribution. Rather reasonably, then, a perfect model is one that yields $\beta =$ 1, or 100%. Values $\beta > 1$ indicate overfitting, and $\beta = 0$ as usual indicates that a constant value would have been better than the proposed model. The assumptions behind this estimation are minimal (Sahani and Linden, 2003): the noise should have zero mean, should have a non-infinite variance, and should be independent between trials. The estimator σ_η^2 is unbiased, and holds even if the variance or other property of the noise depends on the signal strength (as is common with neural signals).

8.6.4 Model deviations

To measure the precision of model predictions in individual stimulus conditions we calculated z-scores of deviations between measured and predicted responses. The z-score of deviations is given by $z = (\bar{r} - m)/\hat{\nu}$ where $\hat{\nu}$ is a robust estimate of the standard deviation of responses calculated across trials. To obtain $\hat{\nu}$ we fitted a power-law relating standard deviations ν to the mean $\bar{r}$ and used the resulting predictions as estimate of standard deviation.

Appendix 1 - Divisive effect of gain control

Here we show that the model of retinal contrast gain control by Shapley and Victor (1981) predicts that contrast has a divisive effect on the responses. The model consists of a fixed low-pass filter followed by a high-pass stage whose time constant varies with contrast.

For simplicity, we consider the version of their model described by Victor (1987), where responses to a stimulus with contrast c and temporal frequency $\omega/(2\pi)$ are given by

$$R(c,\omega) = cL(\omega)\left(1 - \frac{h}{1 + i\omega\tau(c)}\right), \qquad 9.1$$

where $L(\omega)$ is a low-pass filter, and the expression in the large parenthesis is a high-pass filter with gain h and time constant $\tau(c)$. Following Victor (1987), we take the latter to depend on contrast c as

$$\tau(c) = \frac{\tau_0}{1 + (c/c_{50})^n}, \qquad 9.2$$

with n, τ_0, c_{50} constants.

Fixing $h = 1$ (Benardete and Kaplan, 1999), we substitute Equation 9.2 into Equation 9.1 and calculate the absolute value to obtain

$$|R(c,\omega)| = |L(\omega)| \frac{c}{\sqrt{\left(c_{50}^n + c^n\right)^2 + \omega^2 \tau_0^2 c_{50}^{2n}}} \qquad 9.3$$

where $|L(\omega)|$ is the amplitude of the linear response, and the right-hand term captures the effect of gain control. At low temporal frequencies (ω close to 0), the latter reduces to a hyperbolic ratio:

$$|R(c)| = |L(0)| \frac{c}{\left(c_{50}^n + c^n\right)}, \qquad 9.4$$

which demonstrates that contrast gain control in retina is divisive, as contrast appears not only in the numerator but also in the denominator.

Equation 2.1 follows from Equation 9.4 once one sets n=1. Indeed, this value is more consistent with the data (see discussion of Figure 3B) than the value n=2 used by Victor (1987).

Appendix 2 - Model implementation

The models in this thesis are implemented using a basic set of equations. We present here the main results.

Gaussian receptive field

Let F_0 be a 2d Gaussian receptive field of width σ,

$$F_0(\rho,\theta) = \frac{1}{2\pi\sigma^2}\exp\left[-\frac{\rho^2}{2\sigma^2}\right], \qquad 10.1$$

where ρ and θ denote the coordinates of visual space in the polar plane. The response V is obtained by multiplying $F(\rho,\theta)$ with the stimulus $S(\rho,\theta)$ and integrating over space:

$$V = \int_0^{2\pi}\int_0^{r} \rho f(\rho,\theta) s(\rho,\theta)\, d\rho d\theta\,. \qquad 10.2$$

Consider without loss of generality a uniform stimulus of amplitude c enclosed in a circular aperture of radius r:

$$\begin{aligned} S(\rho,\theta) &= c, \rho \le r \\ &= 0, \rho > r \end{aligned}. \qquad 10.3$$

It can be shown that the response to such stimulus follows:

$$\begin{aligned} V(r) &= c\int_0^{2\pi}\int_0^{r} \rho f(\rho)\, d\rho d\theta \\ &= c\left(1-\exp\left(-\frac{r^2}{2\sigma^2}\right)\right) \end{aligned}, \qquad 10.4$$

where r denotes stimulus radius. This expression holds only if the receptive field and the stimulus are concentric. We can generalize the result for a Gaussian receptive field whose center is off the origin by Δ deg:

$$F_\Delta(\rho,\theta) = \frac{1}{2\pi\sigma^2}\exp\left(-\frac{\rho^2+\Delta^2-2\rho\Delta\cos\theta}{2\sigma^2}\right). \qquad 10.5$$

In that case, the response is obtained by resolving the numerical integral

$$V(r) = \frac{c}{\sigma^2}\int_0^r \exp\left(-\frac{\rho^2+\Delta^2}{2\sigma^2}\right) I_0\left(\frac{\rho\Delta}{\sigma^2}\right) d\rho \qquad 10.6$$

where $I_0(z)$ denotes the modified Bessel function of the first kind. Next consider a slightly more complex stimulus, a grating stimulus of spatial frequency u and phase ϕ:

$$\begin{aligned} S(\rho,\theta) &= c\cdot\cos(2\pi u\rho\cos\theta+\phi), \rho \le r \\ &= 0, \rho > r \end{aligned} \qquad 10.7$$

Assuming F and S are centered on the origin, then the response is given by

$$V(r) = c\frac{\cos\phi}{\sigma^2}\int_0^r \exp\left(-\frac{\rho^2}{2\sigma^2}\right) J_0(2\pi u\rho)\, d\rho, \qquad 10.8$$

where $J_0(z)$ denotes the Bessel function of the first kind.

For the more general case of discentric receptive field and stimulus, the response follows

$$V(r) = c\frac{\cos\phi}{\sigma^2}\int_0^r \exp\left(-\frac{\rho^2+\Delta^2}{2\sigma^2}\right) J_0\left(\rho\left(2\pi u + \frac{i\Delta}{\sigma^2}\right)\right) d\rho, \qquad 10.9$$

where Δ denotes the distance between the receptive field and stimulus centers.

Difference-of-Gaussians

Consider a receptive field that consists of the difference of two Gaussians:

$$F(x,y) = \frac{1}{2\pi\sigma^2}\left(\exp\left[-\frac{x^2+y^2}{2\sigma_c^2}\right] - k_s\exp\left[-\frac{x^2+y^2}{2\sigma_s^2}\right]\right), \qquad 10.10$$

where σ_c and σ_s denote the widths of the center and the surround, and k_s denotes the relative strength of the surround. The selectivity of the receptive field for spatial frequency is obtained by computing the Fourier transform:

$$\hat{F}(\omega_x,\omega_y) = \frac{1}{2\pi}\left(\exp\left[-\frac{1}{2}(\omega_x^2+\omega_y^2)\sigma_c^2\right] - k_s\exp\left[-\frac{1}{2}(\omega_x^2+\omega_y^2)\sigma_s^2\right]\right). \qquad 10.11$$

Suppressive field

Here we present the primitives used to calculate the output of the suppressive field. The suppressive field computes the standard deviation of the stimulus *S(x,y,t)* falling over a Gaussian region of visual field *F(x,y)*:

$$c_{\text{local}} = \left(\iiint \left(S(x,y,t) \right)^2 F(x,y) \, dx \, dy \, dt \right)^{1/2} \qquad 10.12$$

Let the stimulus S be a grating of contrast c and radius r, then the output of the suppressive field is given by

$$c_{local} = \frac{c}{\sqrt{2}} \left(1 - e^{r^2/2\sigma} \right)^{1/2} . \qquad 10.13$$

This expression holds if both S and F are centered on the origin. For the model general case of S and F being Δ deg apart,

$$c_{local} = \frac{c}{\sigma^2 \sqrt{2}} \int_0^r \exp\left(-\frac{\rho^2 + \Delta^2}{2\sigma^2} \right) I_0\left(\frac{\rho \Delta}{\sigma^2} \right) d\rho . \qquad 10.14$$

Consider a stimulus consisting of the sum of two full-field gratings of contrasts c_1 and c_2, temporal frequency f, spatial frequency u, differing in orientation by φ deg.

$$c_{local} = \frac{1}{\sqrt{2}} \left(c_1^2 + c_2^2 + 2 c_1 c_2 \exp\left[4\pi^2 \sigma^2 u^2 \left(\cos\varphi - 1 \right) \right] \right)^{1/2} . \qquad 10.15$$

At low spatial frequency ($u \to 0$) and/or for small orientation differences ($\varphi \to 0$), the grating contrasts sum linearly:

$$\begin{aligned} c_{local} &= \frac{1}{\sqrt{2}} \left(c_1^2 + c_2^2 + 2 c_1 c_2 \right)^{1/2} \\ &= \frac{1}{\sqrt{2}} \left(c_1 + c_2 \right) \end{aligned} \qquad 10.16$$

If the gratings have different temporal frequencies, grating contrasts sum in root-mean-square fashion.

$$c_{local} = \frac{1}{\sqrt{2}} \sqrt{c_1^2 + c_2^2} . \qquad 10.17$$

Rectification

Let a response of the form $R(t) = \lfloor A\cos 2\pi ft + B \rfloor$ where $\lfloor \ \rfloor$ denotes the rectification. $V(t)$ is positive in the interval $\left[-t_0 \le t \le t_0\right]$ where $t_0 = \frac{1}{2\pi f}\arccos\left(-\frac{B}{A}\right)$. Applying Fourier decomposition to $R(t)$ yields the mean rate

$$R_0 = \frac{A}{2\pi}\sqrt{1-\frac{B^2}{A^2}} - B\arccos\left(-\frac{B}{A}\right), \qquad 10.18$$

the amplitude of the first harmonic

$$R_1 = \frac{B}{2\pi}\sqrt{1-\frac{B^2}{A^2}} - A\arccos\left(-\frac{B}{A}\right), \qquad 10.19$$

and that of the second harmonic

$$R_2 = \frac{\left(A^2 - B^2\right)}{3A\pi}\sqrt{1-\frac{B^2}{A^2}}\ . \qquad 10.20$$

Appendix 3 - Companion data CD

As required by the Swiss Federal Institute of Technology Zurich, we have compiled all original data and manipulations of data that led to this thesis onto a single compact disc (CD). Most data files and analysis scripts can be inspected and executed using the Release 14 of MATLAB (The Mathworks Inc, Natick, MA). Some files require additional applications such as Mathematical 4.1 (Wolfram Research Inc, Champaign, IL), Office 2002 (Microsoft Corp, Redmond, WA), or Acrobat 5.0 (Adobe Systems Inc, San Jose, CA) to be viewed.

The DVD is organized as follows. The directory `\Code` and the directories therein contain general-purpose MATLAB code common to all chapters (add these directories to your MATLAB path). The directory `\Projects\Registry` contains (1) a human-readable XML (Extensible Markup Language) database of all single units from which we recorded (accessible using routines from `\Code\CellRegistry`), (2) scripts used for the calculation of receptive field eccentricities. The remaining three subdirectories (`\Projects\*`) are named according to their corresponding chapters. Each project includes one or more of the following directories:

`.\maths` for the analytical derivation (to be viewed with Mathematica);

`.\tools` for more general-purpose MATLAB code;

`.\models` for the MATLAB implementations of the models;

`.\work` for the analysis scripts used to fit and to present the data;

`.\fits` or `.\analysis for` the original data and model fits (accessible using routines from `\Code\etc\shareit`);

`.\data` for additional derived data;

`.\diaries` for the text output of the fitting procedures.

The filenames sometimes end with a revision number (e.g. `*-01`). Files with the largest revision have been used for this work.

Bibliography

Ahlsén G, Lindström S (1983) Corticofugal projection to perigenculate neurones in the cat. Acta Physiol Scand 118:181-184.

Ahlsén G, Lindström S, Lo FS (1982) Functional distinction of perigeniculate and thalamic reticular neurons in the cat. Exp Brain Res 46:118-126.

Ahlsén G, Lindström S, Lo FS (1983) Maintained binocular connexions to perigeniculate neurones in visually deprived cats. Acta Physiol Scand 119:109-112.

Albrecht DG (1995) Visual cortex neurons in monkey and cat: effect of contrast on the spatial and temporal phase transfer functions. Vis Neurosci 12:1191-1210.

Albrecht DG, Hamilton DB (1982) Striate cortex of monkey and cat: contrast response function. J Neurophysiol 48:217-237.

Albrecht DG, Geisler WS (1991) Motion sensitivity and the contrast-response function of simple cells in the visual cortex. Vis Neurosci 7:531-546.

Albrecht DG, Farrar SB, Hamilton DB (1984) Spatial contrast adaptation characteristics of neurones recorded in the cat's visual cortex. J Physiol (London) 347:713-739.

Albrecht DG, Geisler WS, Frazor RA, Crane AM (2002) Visual cortex neurons of monkeys and cats: temporal dynamics of the contrast response function. J Neurophysiol 88:888-913.

Alitto HJ, Usrey WM (2003) Corticothalamic feedback and sensory processing. Curr Opin Neurobiol 13:440-445.

Alitto HJ, Usrey WM (2004) The Influence of Contrast on Orientation and Temporal Frequency Tuning in Ferret Primary Visual Cortex. J Neurophysiol.

Alitto HJ, Collins OA, Usrey WM (2002) The influence of corticothalamic feedback on the responses of neurons in lateral geniculate nucleus. Program No. 658.16. In: Abstract Viewer/Itinerary Planner. Washington, DC: Society for Neuroscience.

Alitto HJ, Weyand TG, Usrey WM (2005) Distinct properties of stimulus-evoked bursts in the lateral geniculate nucleus. J Neurosci 25:514-523.

Allison JD, Smith KR, Bonds AB (2001) Temporal-frequency tuning of cross-orientation suppression in the cat striate cortex. Vis Neurosci 18:941-948.

Anderson JS, Lampl I, Gillespie DC, Ferster D (2000) The contribution of noise to contrast invariance of orientation tuning in cat visual cortex. Science 290:1968-1972.

Angelucci A, Levitt JB, Walton EJ, Hupe JM, Bullier J, Lund JS (2002) Circuits for local and global signal integration in primary visual cortex. J Neurosci 22:8633-8646.

Baccus SA, Meister M (2002) Fast and slow contrast adaptation in retinal circuitry. Neuron 36:909-919.

Bauman LA, Bonds AB (1991) Inhibitory refinement of spatial frequency selectivity in single cells of the cat striate cortex. Vision Research 31:933-944.

Benardete EA, Kaplan E (1999) The dynamics of primate M retinal ganglion cells. Vis Neurosci 16:355-368.

Benardete EA, Kaplan E, Knight BW (1992) Contrast gain in the primate retina: P cells are not X-like, some M cells are. Vis Neurosci 8:483-486.

Benevento LA, Creutzfeldt O, Kuhnt U (1972) Significance of intracortical inhibition in the visual cortex. Nature 238:124--126.

Blakemore C, Tobin EA (1972) Lateral inhibition between orientation detectors in the cat's visual cortex. Exp Brain Res 15:439-440.

Blitz DM, Regehr WG (2005) Timing and specificity of feed-forward inhibition within the LGN. Neuron 45:917-928.

Bonds AB (1989) Role of inhibition in the specification of orientation selectivity of cells in the cat striate cortex. Vis Neurosci 2:41-55.

Bonin V, Mante V, Carandini M (2002) The contrast integration field of cat LGN neurons. Program No. 352.16. In: Abstract Viewer/Itinerary Planner, Online Edition. Washington, DC: Society for Neuroscience.

Bonin V, Mante V, Carandini M (2003a) Origins of size tuning in LGN neurons. Journal of Vision 3:19a.

Bonin V, Mante V, Carandini M (2003b) Predicting signal and noise in LGN and V1. Program No. 229.5. In: Abstract Viewer/Itinerary Planner. Washington, DC: Society for Neuroscience.

Bonin V, Mante V, Carandini M (2004a) Nonlinear processing in LGN neurons. In: Advances in Neural Information Processing Systems (Thrun S SL, Schölkopf B, ed), pp 1443-1450. Cambridge, MA: MIT Press.

Bonin V, Mante V, Carandini M (2004b) Integration of stimulus contrast in the early visual system. In: Computational and Systems Neuroscience. Cold Spring Habor, NY.

Bonin V, Mante V, Carandini M (2005) The early visual system only adapts to image mean and standard deviation. In: Computational and Systems Neuroscience. Salt Lake City, UT.

Bradley A, Scottun BC, Ohzawa I, Sclar G, Freeman RD (1987) Visual orientation and spatial frequency discrimination; a comparison of single cells and behavior. J Neurophys 57:755-772.

Brainard DH (1997) The Psychophysics Toolbox. Spatial Vision 10:433-436.

Brown MJ (1994) Principles of Anesthesia and Analgesia. In: Essentials for animal research: A primer for research personnel (Bennett BT, Brown MJ, Schofield JC, eds). Beltsville, Maryland: United States Depar tment of Agriculture - National Agricultural Library.

Budd JM (2004) How much feedback from visual cortex to lateral geniculate nucleus in cat: a perspective. Vis Neurosci 21:487-500.

Buracas GT, Zador AM, DeWeese MR, Albright TD (1998) Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. Neuron 20:959-969.

Cai D, DeAngelis GC, Freeman RD (1997) Spatiotemporal receptive field organization in the lateral geniculate nucleus of cats and kittens. J Neurophysiol 78:1045-1061.

Callaway EM (2005) Structure and function of parallel pathways in the primate early visual system. J Physiol 566:13-19.

Carandini M (2004) Amplification of Trial-to-Trial Response Variability by Neurons in Visual Cortex. PLoS Biol 2:E264.

Carandini M, Heeger DJ (1994) Summation and division by neurons in visual cortex. Science 264:1333-1336.

Carandini M, Ferster D (1997) A tonic hyperpolarization underlying contrast adaptation in cat visual cortex. Science 276:949-952.

Carandini M, Ferster D (2000) Membrane potential and firing rate in cat primary visual cortex. J Neurosci 20:470-484.

Carandini M, Heeger DJ, Movshon JA (1997) Linearity and normalization in simple cells of the macaque primary visual cortex. J Neurosci 17:8621-8644.

Carandini M, Heeger DJ, Senn W (2002) A synaptic explanation of suppression in visual cortex. J Neurosci 22:10053-10065.

Cavanaugh JR, Bair W, Movshon JA (2002a) Nature and interaction of signals from the receptive field center and surround in macaque v1 neurons. J Neurophysiol 88:2530-2546.

Cavanaugh JR, Bair W, Movshon JA (2002b) Selectivity and spatial distribution of signals from the receptive field surround in macaque v1 neurons. J Neurophysiol 88:2547-2556.

Chance FS, Abbott LF (2002) Gain Modulaton from Background Synaptic Input. Neuron 35:773-782.

Chander D, Chichilnisky EJ (2001) Adaptation to temporal contrast in primate and salamander retina. J Neurosci 21:9904-9916.

Cheng H, Chino YM, Smith EL, 3rd, Hamamoto J, Yoshida K (1995) Transfer characteristics of lateral geniculate nucleus X neurons in the cat: effects of spatial frequency and contrast. J Neurophysiol 74:2548-2557.

Chichilnisky EJ (2001) A simple white noise analysis of neuronal light responses. Network 12:199-213.

Chino YM, Kaplan E (1988) Abnormal orientation bias of LGN neurons in strabismic cats. Invest Ophthalmol Vis Sci 29:644-648.

Chubb C, Econopouly J, Landy MS (1994) Histogram contrast analysis and the visual segregation of IID textures. J Opt Soc Am A Opt Image Sci Vis 11:2350-2374.

Chubb C, Landy MS, Econopouly J (2004) A visual mechanism tuned to black. Vision Res 44:3223-3232.

Cleland BG, Lee BB (1985) A comparison of visual responses of cat lateral geniculate nucleus neurones with those of ganglion cells afferent to them. J Physiol (Lond) 369:249-268.

Cleland BG, Lee BB, Vidyasagar TR (1983a) Response of neurons in the cat's lateral geniculate nucleus to moving bars of different length. J Neurosci 3:108-116.

Cleland BG, Harding TH, Tulunay-Keesey U (1983b) Response to the length of moving visual stimuli of the brisk classes of ganglion cells in the cat retina. J Physiol 345:27-45.

Cucchiaro JB, Uhlrich DJ, Sherman SM (1988) Parabrachial innervation of the cat's dorsal lateral geniculate nucleus: an electron microscopic study using the tracer Phaseolus vulgaris leucoagglutinin (PHA-L). J Neurosci 8:4576-4588.

Cudeiro J, Sillito AM (1996) Spatial frequency tuning of orientation-discontinuity-sensitive corticofugal feedback to the cat lateral geniculate nucleus. J Physiol 490 (Pt 2):481-492.

Dan Y, Atick JJ, Reid RC (1996) Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. J Neurosci 16:3351-3362.

Daugman JG (1985) Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. J Opt Soc Am A 2:1160-1169.

Dawis S, Shapley R, Kaplan E (1982) Spatio temporal coupling and asymmetry in cat lateral geniculate nucleus x cells. Annual Spring Meeting Of The Association For Research In Vision And Ophthalmology Incorporated, Sarasota, Fla, Usa, May.

Dawis S, Shapley R, Kaplan E, Tranchina D (1984) The receptive field organization of x cells in the cat spatio temporal coupling and asymmetry. Vision Res 24:549-564.

Dean AF (1981) The relationship between response amplitude and contrast for cat striate cortical neurones. J Physiol (London) 318:413--427.

Dean AF, Tolhurst DJ (1986) Factors influencing the temporal phase of response to bar and grating stimuli for simple cells in the cat striate cortex. Exp Br Res 62:143--151.

DeAngelis GC, Ohzawa I, Freeman RD (1993a) Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. II. Linearity of temporal and spatial summation. J Neurophysiol 69:1118-1135.

DeAngelis GC, Ohzawa I, Freeman RD (1993b) Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. I. General characteristics and postnatal development. J Neurophysiol 69:1091-1117.

DeAngelis GC, Freeman RD, Ohzawa I (1994) Length and width tuning of neurons in the cat's primary visual cortex. J Neurophysiol 71:347-374.

DeAngelis GC, Robson JG, Ohzawa I, Freeman RD (1992) The organization of supression in receptive fields of neurons in cat visual cortex. J Neurophysiol 68:144-163.

Demb JB, Haarsma L, Freed MA, Sterling P (1999) Functional circuitry of the retinal ganglion cell's nonlinear receptive field. J Neurosci 19:9756-9767.

Denning KS, Reinagel P (2005) Visual control of burst priming in the anesthetized lateral geniculate nucleus. J Neurosci 25:3531-3538.

Derrington AM, Lennie P (1984) Spatial and temporal contrast sensitivities of neurons in lateral geniculate nucleus of macaque. J Physiol (Lond) 357:219-240.

Derrington AM, Lennie P, Wright MJ (1979) The mechanism of peripherally evoked responses in retinal ganglion cells. J Physiol 289:299-310.

Dubin MW, Cleland BG (1977) Organization of visual inputs to interneurons of lateral geniculate nucleus of the cat. J Neurophysiol 40:410-427.

Einevoll GT, Plesser HE (2002) Linear mechanistic models for the dorsal lateral geniculate nucleus of cat probed using drifting-grating stimuli. Network 13:503-530.

Enroth-Cugell C, Robson JG (1966) The contrast sensitivity of retinal ganglion cells of the cat. J Physiol (Lond) 187:517-552.

Felisberti F, Derrington AM (1999) Long-range interactions modulate the contrast gain in the lateral geniculate nucleus of cats. Vis Neurosci 16:943-956.

Felisberti F, Derrington AM (2001) Long-range interactions in the lateral geniculate nucleus of the New-World monkey, Callithrix jacchus. Vis Neurosci 18:209-218.

Flecknell PA (1996) Laboratory animal anesthesia, Second Edition. London: Academic Press.

Frazor RA, Geisler WS (2005) Local luminance and contrast in natural images. Vision Res *in press*.

Freeman T, Durand S, Kiper D, Carandini M (2002) Suppression without Inhibition in Visual Cortex. Neuron 35:759.

Funke K, Eysel UT (1998) Inverse correlation of firing patterns of single topographically matched perigeniculate neurons and cat dorsal lateral geniculate relay cells. Vis Neurosci 15:711-729.

Geisler WS, Albrecht DG (1997) Visual cortex neurons in monkeys and cats: detection, discrimination, and identification. Visual Neurosci 14:897-919.

Gillespie DC, Lampl I, Anderson JS, Ferster D (2001) Dynamics of the orientation-tuned membrane potential response in cat primary visual cortex. Nat Neurosci 4:1014-1019.

Girardin CC, Kiper DC, Martin KA (2002) The effect of moving textures on the responses of cells in the cat's dorsal lateral geniculate nucleus. Eur J Neurosci 16:2149-2156.

Gur M, Beylin A, Snodderly M (1997) Response variability of neurons in primary visual cortex (V1) of alert monkeys. J Neurosci 17:2914-2920.

Hammond P (1978) Inadequacy of nitrous oxide/oxygen mixtures for maintaining anaesthesia in cats: satisfactory alternatives. Pain 5:143-151.

Hansel D, van Vreeswijk C (2002) How noise contributes to contrast invariance of orientation tuning in cat visual cortex. J Neurosci 22:5118-5128.

Heeger DJ (1992) Normalization of cell responses in cat striate cortex. Vis Neurosci 9:181-197.

Hikasa Y, Yoshikai T, Takase K, Ogasawara S (1997) Comparisons of prolonged sevoflurane, isoflurane, and halothane anaesthesia combined with nitrous oxide in spontaneously breathing cats. Zentralbl Veterinarmed A 44:427-442.

Hochstein S, Shapley RM (1976) Linear and nonlinear spatial subunits in Y cat retinal ganglion cells. J Physiol 262:265-284.

Hubel D, Wiesel TN (1961) Integrative action in the cat's lateral geniculate body. J Physiol (Lond) 155:385-398.

Hubel DH, Wiesel TN (1977) Functional architecture of macaque visual cortex. Proc R Soc London Ser B 198:1-59.

Ikeda H, Wright MJ (1972) The periphery effect and its relation to the receptive field organization of 'transient' retinal ganglion cells. J Physiol 226:81P-82P.

Ikeda H, Wright MJ (1975) Spatial and temporal properties of `sustained' and `transient' neurones in area 17 of the cat's visual cortex. Exp Br Res 22:363--383.

Jahnsen H, Llinas R (1984a) Voltage-dependent burst-to-tonic switching of thalamic cell activity: an in vitro study. Arch Ital Biol 122:73-82.

Jahnsen H, Llinas R (1984b) Electrophysiological properties of guinea-pig thalamic neurones: an in vitro study. J Physiol 349:205-226.

Jones HE, Sillito AM (1991) The length-response properties of cells in the feline dorsal lateral geniculate nucleus. J Physiol (Lond) 444:329-348.

Jones HE, Sillito AM (1994) The length-response properties of cells in the feline perigeniculate nucleus. Eur J Neurosci 6:1199-1204.

Jones HE, Grieve KL, Wang W, Sillito AM (2001) Surround suppression in primate V1. J Neurophysiol 86:2011-2028.

Jones HE, Andolina IM, Oakely NM, Murphy PC, Sillito AM (2000) Spatial summation in lateral geniculate nucleus and visual cortex. Exp Brain Res 135:279-284.

Kaplan E, Shapley R (1984) The origin of the s slow potential in the mammalian lateral geniculate nucleus. Exp Brain Res 55:111-116.

Kaplan E, Benardete E (2001) The dynamics of primate retinal ganglion cells. Prog Brain Res 134:17-34.

Kaplan E, Purpura K, Shapley R (1987) Contrast affects the transmission of visual information through the mammalian lateral geniculate nucleus. J Physiol (London) 391:267--288.

Kara P, Reinagel P, Reid RC (2000) Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. Neuron 27:635-646.

Kara P, Pezaris JS, Yurgenson S, Reid RC (2002) The spatial receptive field of thalamic inputs to single cortical simple cells revealed by the interaction of visual and electrical stimulation. Proc Natl Acad Sci U S A 99:16261-16266.

Keat J, Reinagel P, Reid RC, Meister M (2001) Predicting every spike. a model for the responses of visual neurons. Neuron 30:803-817.

Kilavik BE, Silveira LC, Kremers J (2003) Centre and surround responses of marmoset lateral geniculate neurones at different temporal frequencies. J Physiol 546:903-919.

Kim KJ, Rieke F (2001) Temporal contrast adaptation in the input and output signals of salamander retinal ganglion cells. J Neurosci 21:287-299.

Kremers J, Silveira LC, Kilavik BE (2001) Influence of contrast on the responses of marmoset lateral geniculate cells to drifting gratings. J Neurophysiol 85:235-246.

Kremers J, Kozyrev V, Silveira LC, Kilavik BE (2004) Lateral interactions in the perception of flicker and in the physiology of the lateral geniculate nucleus. J Vis 4:643-663.

Kruger J, Fischer B (1973) Strong periphery effect in cat retinal ganglion cells. Excitatory responses in ON- and OFF- center neurones to single grid displacements. Exp Brain Res 18:316-318.

Lee BB, Virsu V, Creutzfeldt OD (1983) Linear signal transmission from prepotentials to cells in the macaque lateral geniculate nucleus. Exp Brain Res 52:50-56.

Lesica NA, Stanley GB (2004) Encoding of natural scene movies by tonic and burst spikes in the lateral geniculate nucleus. J Neurosci 24:10731-10740.

Levick WR, Oyster CW, Davis DL (1965) Evidence That Mcilwain's Periphery Effect Is Not a Stray Light Artifact. J Neurophysiol 28:555-559.

Levick WR, Cleland BG, Dubin MW (1972) Lateral geniculate neurons of cat: retinal inputs and physiology. Invest Ophthalmol 11:302-311.

Li C, Creutzfeldt O (1984) The representation of contrast and other stimulus parameters by single neurons in area 17 of the cat. Pflugers Archives 401:304--314.

Li C, Li W (1994) Extensive integration beyond the classical receptive field of cat's striate cortical neurons - classification and tuning properties. Vision Research 34:2337-2356.

Machens CK, Wehr MS, Zador AM (2004) Linearity of cortical receptive fields measured with natural sounds. J Neurosci 24:1089-1100.

Maffei L, Fiorentini A (1973) The visual cortex as a spatial frequency analyzer. Vis Res 13:1255--1267.

Maffei L, Fiorentini A, Bisti S (1973) Neural correlate of perceptual adaptation to gratings. Science **182**:1036-1038.

Mante V (2005) Luminance gain control and contrast gain control in the early visual system. In: Institute of Neuroinformatics. Zurich: Swiss Federal Institute of Technology.

Mante V, Bonin V, Carandini M (2002) Responses of cat LGN neurons to plaids and movies. Program No. 352.15. In: Abstract Viewer/Itinerary Planner. Washington, DC: Society for Neuroscience.

Mante V, Bonin V, Carandini M (2003) Light adaptation and contrast gain control as seen by LGN neurons. Program No. 229.6. In: Abstract Viewer/Itinerary Planner. Washington, DC: Society for Neuroscience.

Mante V, Bonin V, Carandini M (2004a) How local contrast determines gain and temporal integration in LGN. Program No. 409.10. In: Abstract Viewer/Itinerary Planner. Washington, DC: Society for Neuroscience.

Mante V, Bonin V, Carandini M (2004b) Interaction of gain control mechanisms in early visual system. Poster No. 133. In: Computational & Systems Neuroscience. Cold Spring Harbor, NY: Cold Spring Harbor Laboratory.

Mante V, Frazor RA, Bonin V, Geisler WS, Carandini M (2005a) Independence of gain control mechanisms in early visual system matches the statistics of natural images. In: Annual Meeting of the Vision Sciences Society.

Mante V, Frazor RA, Bonin V, Geisler WS, Carandini M (2005b) Independence of luminance and contrast in natural scenes and in the early visual system. Nat Neurosci 8:1690-1697.

Margalit E, Sadda SR (2003) Retinal and optic nerve diseases. Artif Organs 27:963-974.

Martinez LM, Alonso JM, Reid RC, Hirsch JA (2002) Laminar processing of stimulus orientation in cat visual cortex. J Physiol 540:321-333.

McCormick DA, Feeser HR (1990) Functional implications of burst firing and single spike activity in lateral geniculate relay neurons. Neuroscience 39:103-113.

McIlwain JT, Creutzfeldt OD (1964) Receptive fields of optic tract axons and lateral geniculate cells: peripheral extent and barbiturate sensitivity. J Neurophysiol 27:1154-1173.

McIlwain JT, Creutzfeldt OD (1967) Microelectrodes study of synaptic excitation and inhibition in the lateral geniculate nucleus of the cat. J Neurophysiol 30:1-21.

Mechler F, Reich DS, Victor JD (2002) Detection and discrimination of relative spatial phase by V1 neurons. J Neurosci 22:6129-6157.

Miller KD, Troyer TW (2002) Neural noise can explain expansive, power-law nonlinearities in neural response functions. J Neurophysiol 87:653-659.

Morrone MC, Burr DC, Maffei L (1982) Functional implications of cross-orientation inhibition of cortical visual cells. I. Neurophysiological evidence. Proc R Soc Lon B 216:335-354.

Movshon JA, Thompson ID, Tolhurst DJ (1978) Spatial and temporal contrast sensitivity of neurones in areas 17 and 18 of the cat's visual cortex. J Physiol (Lond) 283:101-120.

Muir W (2000) Handbook of veterinary anesthesia, 3rd Edition. St. Louis: Mosby.

Mukherjee P, Kaplan E (1995) Dynamics of neurons in the cat lateral geniculate nucleus: in vivo electrophysiology and computational modeling. J Neurophysiol 74:1222-1243.

Murphy PC, Sillito AM (1987) Corticofugal feedback influences the generation of length tuning in the visual pathway. Nature 329:727-729.

Murphy PC, Uhlrich DJ, Tamamaki N, Sherman SM (1994) Brain-stem modulation of the response properties of cells in the cat's perigeniculate nucleus. Vis Neurosci 11:781-791.

Nolt MJ, Kumbhani RD, Palmer LA (2004) Contrast Dependent Spatial Summation in the Lateral Geniculate Nucleus and Retina of the Cat. J Neurophysiol.

Ohzawa I, Sclar G, Freeman RD (1982) Contrast gain control in the cat visual cortex. Nature 298:266--268.

Ohzawa I, Sclar G, Freeman RD (1985) Contrast gain control in the cat visual system. J Neurophysiol 54:651-665.

O'Keefe LP, Levitt JB, Kiper DC, Shapley RM, Movshon JA (1998) Functional organization of owl monkey lateral geniculate nucleus and visual cortex. J Neurophysiol 80:594-609.

Ozeki H, Sadakane O, Akasaki T, Naito T, Shimegi S, Sato H (2004) Relationship between excitation and inhibition underlying size tuning and contextual response modulation in the cat primary visual cortex. J Neurosci 24:1428-1438.

Paninski L (2003) Convergence properties of three spike-triggered analysis techniques. Network 14:437-464.

Pape HC, Eysel UT (1986) Binocular interactions in the lateral geniculate nucleus of the cat: GABAergic inhibition reduced by dominant afferent activity. Exp Brain Res 61:265-271.

Passaglia CL, Troy JB (2004) Impact of noise on retinal coding of visual signals. J Neurophysiol 92:1023-1033.

Passaglia CL, Enroth-Cugell C, Troy JB (2001) Effects of remote stimulation on the mean firing rate of cat retinal ganglion cells. J Neurosci 21:5794-5803.

Peli E (1990) Contrast in complex images. J Opt Soc Am A 7:2032-2040.

Pelli DG (1997) The VideoToolbox software for visual psychophysics: transforming numbers into movies. Spatial Vision 10:437-442.

Pezaris JS, Reid RC (2004) Microstimulation in LGN produces focal visual percepts. Program No. 598.5. In: Abstract Viewer/Itinerary Planner. Washington, DC: Society for Neuroscience.

Pillow JW, Simoncelli EP (2003) Biases in white noise analysis due to non-Poisson spike generation. Neurocomputing 52-54:109-115.

Price DJ, Morgan JE (1987) Spatial properties of neurones in the lateral geniculate nucleus of the pigmented ferret. Exp Brain Res 68:28-36.

Przybyszewski AW, Gaska JP, Foote W, Pollen DA (2000) Striate cortex increases contrast gain of macaque LGN neurons. Vis Neurosci 17:485-494.

Rampil IJ (1998) A primer for EEG signal processing in anesthesia. Anesthesiology 89:980-1002.

Rao RP, Ballard DH (1999) Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. Nat Neurosci 2:79-87.

Reich DS, Victor JD, Knight BW, Ozaki T, Kaplan E (1997) Response variability and timing precision of neuronal spike trains in vivo. J Neurophysiol 77:2836-2841.

Reid RC, Victor JD, Shapley RM (1997) The use of m-sequences in the analysis of visual neurons: linear receptive field properties. Vis Neurosci 14:1015-1027.

Reinagel P, Godwin D, Sherman SM, Koch C (1999) Encoding of visual information by LGN bursts. J Neurophysiol 81:2558-2569.

Rodieck RW (1965) Quantitative analysis of cat retina ganglion cell response to visual stimuli. Vision Res 5:583-601.

Rodieck RW, Brening RK, Watanabe M (1993) The origin of parallel visual pathways. In: Contrast sensitivity (Shapley RM, Man-Kit Lam D, eds). Cambridge, Massachusetts: MIT Press.

Sahani M, Linden JF (2003) How Linear are Auditory Cortical Responses? In: Advances in Neural Information Processing Systems 15 (Becker S, Thrun S, Obermayer K, eds), pp 109--116. Cambridge, MA: MIT Press.

Sanchez-Vives MV, Nowak LG, McCormick DA (2000) Membrane mechanisms underlying contrast adaptation in cat area 17 *in vivo*. J Neurosci 20:4267-4285.

Sanderson KJ (1971) The projection of the visual field to the lateral geniculate and medial interlaminar nuclei in the cat. J Comp Neurol 143:101-108.

Sanderson KJ, Bishop PO, Darian-Smith I (1971) The properties of the binocular receptive fields of lateral geniculate neurons. Exp Brain Res 13:178-207.

Saul AB, Humphrey AL (1990) Spatial and temporal response properties of lagged and nonlagged cells in cat lateral geniculate nucleus. J Neurophysiol 64:206-224.

Saul AB, Humphrey AL (1992) Temporal-frequency tuning of direction selectivity in cat visual cortex. Vis Neurosci 8:365-372.

Sceniak MP, Hawken MJ, Shapley R (2001) Visual spatial characterization of macaque v1 neurons. J Neurophysiol 85:1873-1887.

Sceniak MP, Hawken MJ, Shapley R (2002) Contrast-dependent changes in spatial frequency tuning of macaque V1 neurons: effects of a changing receptive field size. J Neurophysiol 88:1363-1373.

Sceniak MP, Ringach DL, Hawken MJ, Shapley R (1999) Contrast's effect on spatial summation by macaque V1 neurons. Nat Neurosci 2:733-739.

Schwartz O, Simoncelli EP (2001) Natural signal statistics and sensory gain control. Nat Neurosci 4:819-825.

Sclar G (1987) Expression of "retinal" contrast gain control by neurons of the cat's lateral geniculate nucleus. Exp Brain Res 66:589-596.

Sclar G, Maunsell JHR, Lennie P (1990) Coding of image contrast in central visual pathways of the macaque monkey. Vision Res 30:1-10.

Sengpiel F, Blakemore C (1994) Interocular control of neuronal responsiveness in cat visual cortex. Nature 368:847-850.

Sengpiel F, Vorobyov V (2005) Intracortical origins of interocular suppression in the visual cortex. J Neurosci 25:6394-6400.

Shapley R, Victor JD (1978a) Contrast gain control in the cat retina. Invest Ophthalmol Visual Sci (Suppl).

Shapley R, Victor JD (1979) Nonlinear spatial summation and the contrast gain control of cat retinal ganglion cells. J Physiol (London) 290:141--161.

Shapley R, Kaplan E (1981) Spatial and temporal filtering in the cat lateral geniculate nucleus. Investigative Ophthalmology & Visual Science (Suppl).

Shapley R, Kaplan E, Soodak R (1981) Spatial summation and contrast sensitivity of X and Y cells in the lateral geniculate nucleus of the macaque macaca-fascicularis. Nature (Lond) 292:543-545.

Shapley RM, Victor JD (1978b) The effect of contrast on the transfer properties of cat retinal ganglion cells. J Physiol 285:275-298.

Shapley RM, Victor J (1981) How the contrast gain modifies the frequency responses of cat retinal ganglion cells. J Physiol (London) 318:161-179.

Shapley RM, Enroth-Cugell C (1984) Visual adaptation and retinal gain controls. Prog Retinal Res 3:263-346.

Sherman SM (2001a) Thalamic relay functions. Prog Brain Res 134:51-69.

Sherman SM (2001b) Tonic and burst firing: dual modes of thalamocortical relay. Trends Neurosci 24:122-126.

Sherman SM, Koch C (1986) The control of retinogeniculate transmission in the mammalian lateral geniculate nucleus. Exp Brain Res 63.

Sherman SM, Guillery RW (2004) Thalamus. In: The synaptic organization of the brain, 5th Edition (Shepherd GM, ed), pp xiv, 719 p. Oxford ; New York: Oxford University Press.

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. Visual Neurosci 13:605-613.

Shou TD, Leventhal AG (1989) Organized arrangement of orientation-sensitive relay cells in the cat's dorsal lateral geniculate nucleus. J Neurosci 9:4287-4302.

Sigl JC, Chamoun NG (1994) An introduction to bispectral analysis for the electroencephalogram. J Clin Monit 10:392-404.

Sillito AM, Jones HE (2002) Corticothalamic interactions in the transfer of visual information. Philos Trans R Soc Lond B Biol Sci 357:1739-1752.

Sillito AM, Cudeiro J, Murphy PC (1993) Orientation sensitive elements in the corticofugal influence on centre-surround interactions in the dorsal lateral geniculate nucleus. Exp Brain Res 93:6-16.

Simoncelli EP, Pillow JW, Paninski L, Schwartz O (2004) Characterization of neural responses with stochastic stimuli. In: The Cognitive Neurosciences, 3rd edition Edition (Gazzaniga M, ed): MIT Press.

Singer W, Creutzfeldt OD (1970) Reciprocal lateral inhibition of on- and off-center neurones in the lateral geniculate body of the cat. Exp Brain Res 10:311-330.

Singer W, Poppel E, Creutzfeldt O (1972) Inhibitory interaction in the cat's lateral geniculate nucleus. Exp Brain Res 14:210-226.

Smirnakis SM, Berry MJ, Warland DK, Bialek W, Meister M (1997) Adaptation of retinal processing to image contrast and spatial scale. Nature 386:69-73.

So YT, Shapley R (1981) Spatial tuning of cells in and around lateral geniculate nucleus of the cat x and y relay cells and peri geniculate inter neurons. J Neurophysiol (Bethesda) 45:107-120.

Solomon SG, White AJ, Martin PR (2002) Extraclassical receptive field properties of parvocellular, magnocellular, and koniocellular cells in the primate lateral geniculate nucleus. J Neurosci 22:338-349.

Solomon SG, Sun H, Lee BB (2004a) Surround suppression in magnocellular-pathway ganglion cells of the macaque retina. G88. In: Vision Science Society (Vision Jo, ed).

Solomon SG, Peirce JW, Dhruv NT, Lennie P (2004b) Profound contrast adaptation early in the visual pathway. Neuron 42:155-162.

Somers DC, Todorov EV, Siapas AG, Toth LJ, Kim DS, Sur M (1998) A local circuit approach to understanding integration of long-range inputs in primary visual cortex. Cereb Cortex 8:204-217.

Soodak RE, Shapley RM, Kaplan E (1987) Linear mechanism of orientation tuning in the retina and lateral geniculate nucleus of the cat. J Neurophysiol 58:267-275.

Stevens JK, Gerstein GL (1976) Spatiotemporal organization of cat lateral geniculate receptive fields. J Neurophysiol 39:213-238.

Stone J, Fukuda Y (1974) Properties of cat retinal ganglion cells: a comparison of W-cells with X- and Y-cells. J Neurophysiol 37:722-748.

Sun C, Chen X, Huang L, Shou T (2004) Orientation bias of the extraclassical receptive field of the relay cells in the cat's dorsal lateral geniculate nucleus. Neuroscience 125:495-505.

Tavazoie SF, Reid RC (2000) Diverse receptive fields in the lateral geniculate nucleus during thalamocortical development. Nat Neurosci 3:608-616.

Thompson JK, Peterson MR, Freeman RD (2003) Single-neuron activity and tissue oxygenation in the cerebral cortex. Science 299:1070-1072.

Tolhurst DJ, Movshon JA, Thompson ID (1981) The dependence of response amplitude and variance of cat visual cortical neurons on stimulus contrast. Exp Br Res 41:414--419.

Troy JB (1983) Spatial contrast sensitivities of X and Y type neurones in the cat's dorsal lateral geniculate nucleus. J Physiol (Lond) 344:399-417.

Uhlrich DJ, Cucchiaro JB, Humphrey AL, Sherman SM (1991) Morphology and axonal projection patterns of individual neurons in the cat perigeniculate nucleus. J Neurophysiol 65:1528-1541.

Usrey WM, Reid RC (2000) Visual physiology of the lateral geniculate nucleus in two species of new world monkey: Saimiri sciureus and Aotus trivirgatis. J Physiol 523 Pt 3:755-769.

Valerio R, Navarro R (2003) Input-output statistical independence in divisive normalization models of V1 neurons. Network 14:733-745.

Victor J (1987) The dynamics of the cat retinal X cell centre. J Physiol (London) 386:219-246.

Vidyasagar TR, Urbas JV (1982) Orientation sensitivity of cat LGN neurones with and without inputs from visual cortical areas 17 and 18. Exp Brain Res 46:157-169.

Vogels R, W. S, Orban GA (1989) The response variability of striate cortical neurons in the behaving monkey. Exp Brain Res 77:432-436.

Walker GA, Ohzawa I, Freeman RD (1999) Asymmetric suppression outside the classical receptive field of the visual cortex. J Neurosci 19:10536-10553.

Wang C, Dreher B, Burke W (1994) Non-dominant suppression in the dorsal lateral geniculate nucleus of the cat: laminar differences and class specificity. Exp Brain Res 97:451-465.

Wassle H (2004) Parallel processing in the mammalian retina. Nat Rev Neurosci 5:747-757.

Webb BS, Tinsley CJ, Barraclough NE, Easton A, Parker A, Derrington AM (2002) Feedback from V1 and inhibition from beyond the classical receptive field modulates the responses of neurons in the primate lateral geniculate nucleus. Vis Neurosci 19:583-592.

Weyand TG, Boudreaux M, Guido W (2001) Burst and tonic response modes in thalamic neurons during sleep and wakefulness. J Neurophysiol 85:1107-1118.

Wilson PD, Rowe MH, Stone J (1976) Properties of relay cells in the cat's lateral geniculate nucleus: A comparison of W-cells with X- and Y-cells. J Neurophysiol 39:1193-1209.

Worgotter F, Eyding D, Macklis JD, Funke K (2002) The influence of the corticothalamic projection on responses in thalamus and cortex. Philos Trans R Soc Lond B Biol Sci 357:1823-1834.

Wrobel A, Bekisz M (1994) Visual classification of X and Y perigeniculate neurons of the cat. Exp Brain Res 101:307-313.

Xu X, Bonds AB, Casagrande VA (2002) Modeling receptive-field structure of koniocellular, magnocellular, and parvocellular LGN cells in the owl monkey (Aotus trivigatus). Vis Neurosci 19:703-711.

Xue JT, Carney T, Ramoa AS, Freeman RD (1988) Binocular interaction in the perigeniculate nucleus of the cat. Exp Brain Res 69:497-508.

Yang Y, Jin J, Zhou Y, Shou T (2003) GABA(A) and GABA(B) receptors mediated inhibition affect the pattern adaptation of relay cells in the dorsal lateral geniculate nucleus (LGNd) of cats. Brain Res 959:295-303.

Zaghloul KA, Boahen K, Demb JB (2003) Different circuits for ON and OFF retinal ganglion cells cause different contrast sensitivities. J Neurosci 23:2645-2654.

Zaghloul KA, Boahen K, Demb JB (2005) Contrast adaptation in subthreshold and spiking responses of mammalian Y-type retinal ganglion cells. J Neurosci 25:860-868.

Zhou Y, Leventhal AG, Thompson KG (1995) Visual deprivation does not affect the orientation and direction sensitivity of relay cells in the lateral geniculate nucleus of the cat. J Neurosci 15:689-698.

VINCENT BONIN

The Smith-Kettlewell Eye Research Institute
2318 Fillmore Street
San Francisco, CA 94115, USA

+1 415 345 2006
Fax: +1 415 345 8455
vincent@ski.org

Education

SWISS FEDERAL INSTITUTE OF TECHNOLOGY — Zurich, Switzerland
PhD, Neuroscience (Dept. Physics), 2005

TECHNISCHE UNIVERSITÄT HAMBURG-HARBURG — Hamburg, Germany
M.S., Computer Science and Electrical Engineering, 2000

UNIVERSITÉ DE SHERBROOKE — Sherbrooke, Canada
B.Eng., Electrical Engineering, 1997

Employment

2002-current, Research Assistant, Smith-Kettlewell Eye Research Institute — San Francisco, USA

2000-2002, Research Assistant, Institute of Neuroinformatics, University of Zurich — Zurich, Switzerland

1999-2000, Research Assistant, Dept Computer Architecture, Univ. Politecnica de Catalunya — Barcelona, Spain

1998-1999, Research Assistant, Technische Universität Hamburg-Harburg — Hamburg, Germany

1996-1997, System Engineer, Mentor Engineering — Calgary, Canada

Awards and Honors

2004, Alumnus, "Computational Neuroscience: Vision". Cold Spring Harbor Laboratory — Cold Spring Harbor, USA

1999, Fellowship, German Academic Exchange Service — Bonn, Germany

1998, President's Award for Outstanding Performance, Tech. Univ. Hamburg-Harburg — Hamburg, Germany

Articles and Conference Papers

V Bonin, V Mante, M Carandini, "The suppressive field of LGN neurons", *Journal of Neuroscience*. In press.

V Mante, R Frazor, V Bonin, WS Geisler, M Carandini, "Independence of gain control mechanisms in early visual system matches the statistics of natural images", *Nature Neuroscience*. In press.

V Bonin, V Mante, M Carandini, "Nonlinear processing in LGN neurons", In: *Advances in Neural Information Processing Systems 16* (Thrun S, Saul L, Schölkopf B, eds). Cambridge, MA: MIT Press (2004).

P König, C Kayser, V Bonin, RP Würtz, "Efficient evaluation of serial sections by iterative Gabor matching" *J. Neurosci. Meth.* 111(2):141-50. (2001).

V Bonin, O Casals, B Van Houdt, C Blondia, "Performance Modeling of Differential Fair Buffer Allocation", 9th International Conference on Telecommunication Systems, Dallas, USA (2001).

V Bonin, F Cerdán, O Casals, "A Simulation Study of Differential Fair Buffer Allocation", *Joint IEEE* ATM Workshop 2000 and 3rd International Conference on ATM, Heidelberg, Germany (2000).

Abstracts

V Bonin, V Mante, M Carandini. "Response sparseness in the early visual system", Program 591.15. *Abstract Viewer/Itinerary Panner*. Washington, DC: Society for Neuroscience (2005).

V Bonin, V Mante, M Carandini. "The early visual system only adapts to image mean and standard deviation". Computational and Systems Neuroscience. Salt Lake City, UT (2005).

V Mante, RA Frazor, V Bonin, W. S. Geisler and M. Carandini. Independence of gain control mechanisms in early visual system matches the statistics of natural images. *Journal of Vision*, *5*(8), 371a, http://journalofvision.org/5/8/371/, doi:10.1167/5.8.371 (2005).

RA Frazor, V Mante, V Bonin, M Carandini. Dynamics of spatial frequency tuning in lateral geniculate nucleus [Abstract]. *Journal of Vision*, *5*(8), 428a, http://journalofvision.org/5/8/428/, doi:10.1167/5.8.428 (2005).

V Bonin, V Mante, M Carandini. "A fixed firing threshold and constant noise predict trial-to-trial variability in LGN and V1 neurons". Program No. 984.19. *Abstract/Itinerary Planner*. Washington, DC: Society for Neuroscience (2004).

V Bonin, V Mante, M Carandini. "Integration of stimulus contrast in the early visual system". Computational and Systems Neuroscience. Cold Spring Habor, NY (2004).

V Mante, V Bonin and M Carandini. How local contrast determines gain and temporal integration in LGN. Program No. 409.10. *Abstract Viewer/Itinerary Planner*. Washington, DC, Society for Neuroscience (2004).

V Mante, V Bonin and M Carandini. Interaction of gain control mechanisms in early visual system. Poster No. 133. Computational and Systems Neuroscience, Cold Spring Harbor, NY, Cold Spring Harbor Laboratory (2004).

V Bonin, V Mante, M Carandini. "Predicting signal and noise in LGN and V1". Program No. 229.5. *Abstract Viewer/Itinerary Planner*. Washington, DC: Society for Neuroscience (2003).

V Bonin, V Mante, M Carandini. "Origins of size tuning in LGN neurons". *Journal of Vision*, 3(9), 19a, http://journalofvision.org/3/9/19/, doi:10.1167/3.9.19. (2003).

V Mante, V Bonin and M Carandini. Light adaptation and contrast gain control as seen by LGN neurons. Program No. 229.6. Abstract Viewer/Itinerary Planner. Washington, DC, Society for Neuroscience (2003).

V Bonin, V Mante, M Carandini. "The contrast integration field of cat LGN neurons". Program No. 352.16. *Abstract Viewer/Itinerary Planner*. Washington, DC: Society for Neuroscience (2002).

V Mante, V Bonin and M Carandini. Responses of cat LGN neurons to plaids and movies. Program No. 352.15. Abstract Viewer/Itinerary Planner. Washington, DC, Society for Neuroscience (2002).

Languages

Fluent French, English, German and Spanish.