NORTHWESTERN UNIVERSITY

Mechanisms Underlying Response Suppression in Cat Visual Cortex

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Interdepartmental Neuroscience

By

Ian Micah Finn

EVANSTON, ILLINOIS

December 2007

Acknowledgements

Much of the work presented in this thesis was highly collaborative, and so I would first like to thank Nicholas Priebe and Hirofumi Ozeki for allowing me to work with them on some very important questions in visual neuroscience. Dr. Priebe in particular has been the source of much aid over the years, generously contributing everything from basic factual knowledge up to and including suggesting entirely new ideas. In addition, his insights about existing problems, and his willingness to share computer code, were both tremendously appreciated.

I thank Beth Boudreau, who is a terrific role model for scientific dedication and conscientiousness. Her technical skill and willingness to teach were invaluable during my experimental training.

I thank my advisor, David Ferster, for his flexible management, which allowed me to pursue projects that I found interesting and worthwhile. His facility with the process of writing scientific papers, and in doing so making a coherent whole out of a bunch of seemingly disconnected parts, is truly admirable. I would consider myself lucky if I have learned in any way to emulate his style.

I thank my family – My brothers Lucas and Adam and my parents Amy and Mark – for their enduring support and encouragement. And finally, I thank my wife Daphna, whose friendship and love I can always count on to see me through.

ABSTRACT

Mechanisms Underlying Response Suppression in Cat Visual Cortex

Ian Micah Finn

A comprehensive understanding of how image processing occurs in the primary visual cortex (V1) requires learning what aspects of neuronal responses are driven by strong feed-forward input from the lateral geniculate nucleus (LGN), and what aspects arise due to the densely recurrent network operating within the cortex itself. From an anatomical perspective, the preponderance of intracortical excitatory and inhibitory connections over feed-forward excitatory connections stands as strong evidence that intracortical input is critical for determining V1 response properties. A particularly appealing hypothesis is that intracortical inhibition operates to shape response selectivity in V1, as it does in the retina for stimuli of varying size.

A number of suppressive response patterns in V1 indirectly support the importance of intracortical inhibition, including: (1) contrast-invariant orientation tuning in simple cell membrane potential responses, (2) MAX-like responses in complex cells, and (3) surround suppression in both simple and complex cells. It is not clear, however, whether inhibition has a distinct role in determining these properties, or if they arise from more excitatory means. Recently, Hubel and Wiesel's hierarchical model, which invokes only excitatory feed-forward input, has been used successfully to explain the suppressive phenomena of sharp orientation

tuning and cross-orientation suppression. These results suggest that their model may be able to account for other suppressive behaviors in visual cortex.

In this thesis I present data from intracellular experiments designed to probe various mechanisms of suppression in V1. Based on my results, I conclude that an excitatory, feed-forward architecture can in fact explain the emergence of contrast-invariant orientation tuning in simple cells, and MAX behavior in complex cells, when modest extensions to the traditional model are considered. In addition, I conclude that surround suppression is mediated by cortical inhibition, but in a novel manner that reflects a critical and highly general contribution of inhibition to stabilizing the visual network.

Table of Contents

Figure List

Chapter I

Introduction

Learning how networks of living cells in the brain detect and segment sensory information is a formidable challenge. Employing mechanisms that as yet defy understanding, the brain processes stimuli with a fluidity, speed, and accuracy that has not been approached artificially, except in very well circumscribed domains. The rewards for discovering these mechanisms would likely be significant, as there are many real-world problems, such as prosthesis development, face recognition, auditory discrimination, and sensor design, that might find solutions, or inspiration from solutions, currently implemented by the brain.

To determine how cortical networks represent and carry out computations on sensory information, it is necessary to observe and interpret the behavior of many cells simultaneously, often a difficult experimental task. An alternative approach to decoding network behavior is in-vivo intracellular recording, which reveals the aggregate activity of all network elements that influence the computations of a single neuron. The question can then be asked: How do the computations of individual neurons reflect the organization of the cortical network they are embedded in?

Difficulties arise in studying this question owing to the extreme complexity of networks in most areas of neo-cortex, including sensory cortices. Neurons may receive hundreds to thousands of inputs, both excitatory and inhibitory, that potentially contribute to their computations (Ahmed et

al., 1994; Fitzpatrick, 1996; Callaway, 1998). While constraints based on experimental measurements can be placed on models of cortical operation, in general most observations of neuronal behavior are consistent with a wide variety of possible network architectures.

Feed-forward and recurrent neural networks

Primary visual cortex or V1 is a genuine area of neo-cortex, and as such exhibits the dense connectionist networks characteristic of other cortices. Primary visual cortex is unique, however, in having a well characterized input pathway from the thalamus, specifically the lateral geniculate nucleus or LGN (Hubel and Wiesel, 1961), which itself receives input directly from the retina. Several decades ago, Hubel and Wiesel took a conceptual leap into understanding V1 physiology by proposing that these feed-forward connections, and not recurrent intracortical connections, give rise to the fundamental visual receptive field property of orientation tuning (Hubel and Wiesel, 1962).

In feed-forward processing, successive stages of neuronal computation rely predominantly on the output of previous stages. In recurrent processing, neuronal computations at any stage depend largely on feedback from within that stage. By setting out their feed-forward model, Hubel and Wiesel benefited both experimentalists and modelers alike – because of its relative simplicity, the hypothesis of feed-forward processing generates conceptually and experimentally tractable questions in a way that is difficult to do for hypotheses involving recurrent processing. For example, it is more difficult to prove in the lab that simple cell orientation tuning is due to a winner-take-all type network operating in V1 than it is to demonstrate properties of orientation tuning that are consistent with input from spatially aligned, pre-synaptic LGN cells.

While it is likely that much of higher level cognition arises from neural networks that heavily employ recurrent connections, the inherent difficulty of unraveling recurrent contributions to sensory processing provides a clear impetus for investigating neuronal responses with an eye towards whether they can be accounted for fully or partially by feed-forward connections. This is particularly true in primary visual cortex, where recent work has revealed the critical role of sequential, feed-forward processing in encoding information about real-world stimuli (Anderson et al., 2000c; Lampl et al., 2001; Priebe and Ferster, 2006), a role which has often been misattributed to recurrent cortical connections. These misattributions have been most frequent for suppressive phenomena in V1, which, because LGN synapses are exclusively excitatory, are assumed to require cortically supplied inhibition.

Lateral inhibition

Proposals for the role of recurrent inhibitory connections in V1 processing began to appear in the late 1960s and early 1970s (Blakemore et al., 1970; Blakemore and Tobin, 1972). It had been observed that the responses of peripheral sensory cells, such as retinal ganglion cells and skin somatosensory neurons, were decreased when the stimuli that maximally excited them were enlarged beyond a certain spatial extent. This phenomenon was shown to be mediated by an increase in inhibition from receptors in the surround of a given central receptor (Wiesel, 1959), and was thus termed “lateral inhibition” (Fig. 2.01).

Figure 2.01. **Lateral inhibition in the retina.**

Reproduced from Wiesel, 1959.

A. Intracellular recordings from a retinal ganglion cell show a large depolarization and a high firing rate evoked in response to a small spot of light over the center of the receptive field.

B. An annulus of light surrounding the receptive field of the recorded ganglion cell produces a hyperpolarization and shuts off firing activity.

C. Illumination of both the center and surround of the ganglion cell's receptive field leads to a smaller depolarization and less firing than in **A**.

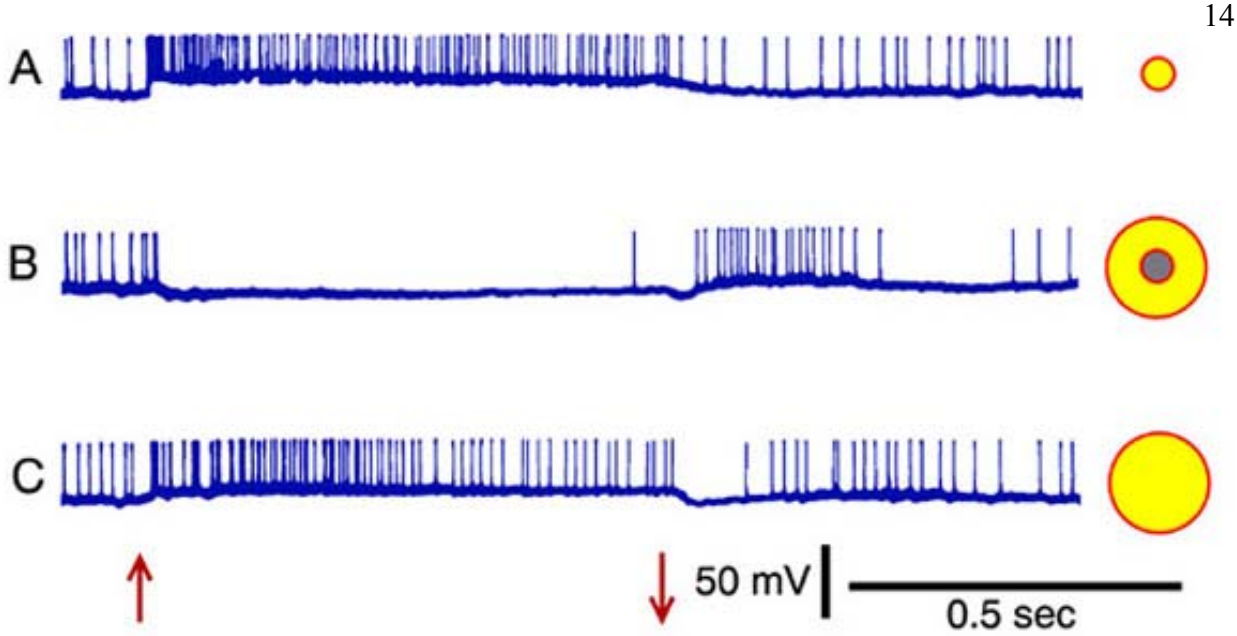
A
B
C
50 mV
0.5 sec

In both the retina and in the skin, lateral inhibition is critical for precisely localizing stimuli. Because it decreases the responses of neurons to non-optimally sized stimuli, lateral inhibition is also a powerful tool for sharpening feature selectivity. The importance of lateral inhibition in peripheral nervous tissues suggested that it might find implementations in neo-cortex, where it could underlie solutions to more general problems of discrimination. In particular, Blakemore et al. (1970) introduced the idea that inhibitory connections between neurons in V1 might give rise to the neurons' specificity when responding to oriented visual stimuli.

Two broad classes of mechanism were proposed to explain how lateral inhibition in the orientation domain might generate orientation specificity. One possibility was that neurons with orthogonal preferences for orientation might inhibit each other, thus blunting all responses except those to the preferred stimulus (Fig. 2.02A). An alternative possibility was that neurons with similar orientation preferences provide mutual inhibition to each other, but the tuning for inhibitory inputs is broader than the tuning for excitatory inputs. For a given neuron in this scenario, only stimuli of the preferred orientation would evoke a significant response (Fig. 2.02B).

Figure 2.02. **Models for the origin of orientation selectivity.**

Reproduced and modified from Ferster, 1986.

A. Large inhibitory inputs may be evoked by stimuli with orientation preferences that are very dissimilar to that of a recorded cell, sharpening orientation selectivity.

B. Conversely, large inhibitory inputs may be evoked by stimuli with similar, or even identical, orientation preferences to that of a recorded cell; however, the inhibitory tuning bandwidth would be larger than the excitatory tuning bandwidth.

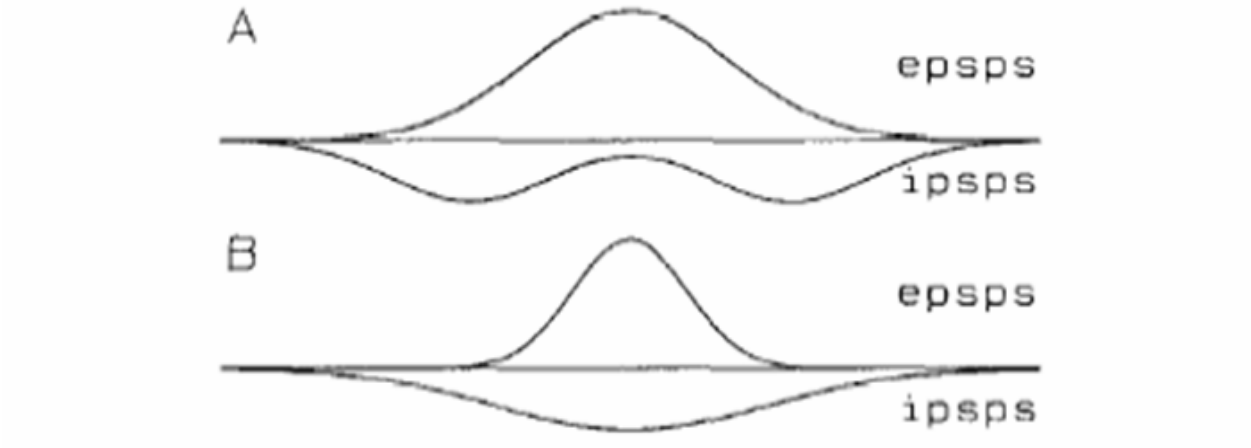
A
epsps
ipsps
B
epsps
ipsps

Sharpness of orientation tuning

A major limitation of studying V1 responses extracellularly is that the underlying process of synaptic integration, often of particular interest, is frequently obscured. Intracellular studies are thus vital to gain a full, and direct, understanding of the excitatory and inhibitory inputs evoked by visual stimuli that lead to the observation of a given spike train. Historically however, much evidence for lateral inhibition has been derived from indirect, extracellular measurements.

One observation that has been considered strong evidence for the importance of lateral inhibition is that when measured extracellularly using drifting gratings, tuning for orientation in simple cells appears much narrower than what would be expected based on predictions derived from receptive field maps (Gardner et al., 1999). The comparison of these measures seems reasonable, as an immediate consequence of Hubel and Wiesel's feed forward model is that, for simple cells, orientation tuning bandwidth should be derivable from a detailed two-dimensional map of the neuron's responses to small, unoriented patches of light. That is, if orientation tuning in simple cells is created by the convergence of spatially aligned LGN inputs, then mapping the extent of those inputs should reveal elongated subfields whose aspect ratio is directly related to the sharpness of orientation tuning exhibited by the neuron.

The mismatch between predictions of orientation selectivity made from simple cell receptive field maps and actual selectivity measured extracellularly with drifting gratings was taken to indicate that receptive field maps fail to reflect an intracortical contribution to orientation tuning that inhibits responses to non-preferred stimuli. In such a scenario, LGN relay cells would

establish the weak orientation tuning inherent in receptive field maps, but final tuning bandwidth would be determined by the magnitude of suppression ("sharpening") provided by the cortex.

What was overlooked in this scenario is that receptive field maps more closely reflect the pattern of synaptic input to simple cells than the pattern of output produced by simple cells. Thus the tuning bandwidth predicted by receptive field maps should be similar to the bandwidth measured intracellularly, and not extracellularly, in response to drifting gratings. Lampl et al. (2001) demonstrated very clearly that this was the case. Their results strongly suggested that orientation tuning bandwidth is narrowed extracellularly by the intrinsic non-linearity of neuronal spike threshold, which limits the responses that are observable extracellularly to those stimulus orientations that evoke supra-threshold intracellular responses.

The sharpness of orientation tuning story is paradigmatic of how recurrent, cortical contributions to sensory processing have traditionally, and sometimes mistakenly, been ascertained. Response suppression can be the result of one or more non-linear processes intrinsic to neurons, and does not necessarily imply the active involvement of cortex. Sharp orientation tuning arises primarily not from cortically supplied lateral inhibition, but from the well understood non-linearity of neuronal spike threshold.

Cross-orientation suppression

Another observation that has been considered evidence for cortical lateral inhibition is cross-orientation suppression, in which a simple cell's response is decreased when a non-preferred

(especially an orthogonally or "cross" oriented) stimulus is presented. Classic experiments examining the spike responses to orthogonally oriented stimuli employed a variety of techniques which were thought to enhance the visibility of inhibition when measured extracellularly. Background firing rates, normally very small or even zero, were elevated by applying drugs to increase cellular excitability, or by using constant visual stimulation (one-dimensional white noise, for example) superimposed on the stimulus of interest (Bishop et al., 1971; Bishop et al., 1973; Morrone et al., 1982; Ramoa et al., 1986). Firing activity was observed to decrease when stimuli orthogonal to the preferred were tested under these conditions, signaling to the authors that lateral inhibition, and in particular cross-orientation inhibition, was being supplied to the recorded neuron.

Additional evidence for cross-orientation suppression has been assumed from experiments in which both preferred and orthogonally oriented drifting gratings were presented simultaneously (plaid stimulus). Such plaid stimuli evoked much less firing than that measured in response to preferred stimuli alone (DeAngelis et al., 1992). The suppression in firing has been attributed to lateral inhibition from cortical neurons tuned to the orthogonal orientation.

Priebe and Ferster (2006) showed, however, that a majority of cross-orientation suppression could be explained by properties of the feed-forward geniculate input provided to simple cells. They observed that because LGN neurons exhibit gross non-linearities (such as saturation and rectification) in their firing rates, their responses to combinations of visual stimuli were non-linear. Ultimately, Priebe and Ferster demonstrated that a preferred and orthogonal stimulus, linearly combined, would yield sub-linear peak input to a model simple cell from a spatially

aligned population of LGN neurons. In this way they accounted for the experimentally observed suppression in simple cell responses without invoking intracortical connections.

The role of inhibition in primary visual cortex

The work of Lampl et al. (2001) and Priebe et al. (2006) strongly suggested that non-linear properties intrinsic to neurons could act on purely excitatory, feed-forward input and produce response patterns that otherwise would appear to demand cortical lateral inhibition. It is important to note, however, that neither study implies inhibition does not exist. Many direct measurements have in fact confirmed that inhibition is ubiquitous in V1 (Ferster, 1986; Anderson et al., 2000b; Anderson et al., 2001; Hirsch et al., 2003; Monier et al., 2003; Marino et al., 2005; Priebe and Ferster, 2005, 2006; Tucker and Fitzpatrick, 2006), and the majority of evidence suggests that it occurs between neurons sharing similar preferences for stimulus features. Functionally, however, inhibition does not seem to act the same way centrally as it does peripherally.

That cortical lateral inhibition may not be involved in determining critical features of V1 responses places renewed emphasis on questions regarding how inhibition contributes or does not contribute to sensory processing. Are there features of neuronal responses that arise solely or predominantly from lateral inhibition? Conversely, are there other suppressive phenomena that do not originate from the action of inhibition? What might the role of inhibition be in visual cortex? In this thesis I present three studies of V1 that address these questions. The first study on the contrast invariance of orientation tuning concerns another property of simple cell responses that has been thought of as strong evidence that lateral inhibition contributes to visual

computations. The second study on the MAX behavior of complex cells bears on whether the energy model, which is a particular instantiation of Hubel and Wiesel's simple cell to complex cell projection, can explain the responses of complex cells to pairs of visual stimuli without the need for inhibition. The third study about the nature of visual surround suppression has relevance to the question of how inhibition operates in the cortical network.

Contrast invariant orientation tuning

In their study of cross-orientation suppression, Sclar and Freeman (1982) demonstrated that, for a subset of cells with high spontaneous activity, the apparent inhibition evoked by non-preferred stimuli was contrast dependent. High contrast drifting gratings at non-preferred orientations suppressed firing activity more than identical gratings at lower contrasts. These observations were seen to have profound implications for how the orientation tuning of V1 neurons arises. In the same study, Sclar and Freeman demonstrated that the bandwidth for orientation tuning was independent of contrast – for the majority of cells, stimuli that evoked spikes at high contrast invariably evoked spikes at low contrast as well.

That orientation tuning bandwidth does not depend on contrast is paradoxical if, as laid out in Hubel and Wiesel's feed-forward model, geniculate relay cell input is responsible for establishing tuning. As the firing activity of relay cells is highly contrast dependent, there should exist a stimulus orientation that evokes suprathreshold input from the LGN at high contrast, but fails to do so at low contrast. Contrast dependent inhibition activated by non-preferred stimuli (cross-orientation inhibition) could potentially explain how cortical cells counteract strong LGN

input at high contrast (Troyer et al., 1998; Ferster and Miller, 2000), producing contrast-invariant orientation tuning.

Recent intracellular studies showing small or no depolarizations to orthogonally oriented stimuli (Anderson et al., 2000c; Lampl et al., 2001) would seem to argue in favor of the cross-orientation inhibition hypothesis, though the requisite inhibition does not appear to be present when measured directly (Anderson et al., 2000b; Priebe and Ferster, 2006), except perhaps in a small number of cells (Monier et al., 2003). The question of whether any simple cells depolarize in response to orthogonal stimuli, and if not what mechanism prevents them from doing so, thus remains open. In addition, extracellular contrast-dependent suppression in response to orthogonally oriented drifting gratings remains to be accounted for. The data in Chapter II of this thesis address these questions, and tie their answers into an overall explanation for the origin of contrast-invariant orientation tuning in simple cells that does not incorporate lateral inhibition.

MAX behavior in complex cells

The question of how complex cell receptive fields are created is fundamental to understanding how this broad class of neurons contributes to information processing in V1. Chapter II adds to the list of response properties in V1 simple cells that can be explained by a purely excitatory, feed-forward network architecture as described by Hubel and Wiesel. There remains a great deal of uncertainty, however, about the origin of many complex cell response properties.

Complex cells are highly non-linear (Hubel and Wiesel, 1962; Movshon et al., 1978b), and mapping the spatial extent of their receptive fields fails to reveal how their orientation or spatial

frequency selectivity arises; instead, a uniform, excitatory area is generally seen. Most theories for the formation of complex cells in the superficial layers (Layers II and III), however, assume input from pre-synaptic Layer IV simple cells, whose orientation and spatial frequency selectivities have much better understood origins. A simple to complex cell projection could thus account for complex cell tuning via inheritance. While evidence does support this idea (Alonso and Martinez, 1998; Martinez and Alonso, 2001, 2003), many questions remain. In particular, an intracellular study by Lampl et al. (2004) reported data that were not consistent with a commonly accepted model for the formation of complex cell receptive fields, the Adelson-Bergen energy model (Adelson and Bergen, 1985).

The energy model proposes an explicit arrangement of pre-synaptic simple cells and a linear mechanism for their integration by complex cells. Inputs from at least four simple cells are required. The simple cells are identical in almost all respects – for example, they have the same receptive field centers, orientation tuning, and spatial frequency tuning – except for their relative spatial phases, which are equally distributed between 0^o and 360^o. Each simple cell's output is half-wave rectified.

Lampl et al. (2004) recorded the responses of complex cells to pairs of small, rectangular bar stimuli of either light or dark polarity, separated by various distances from each other within the receptive field. The authors showed that the great majority of stimulus pairs evoked membrane potential responses that were significantly less than the sum of the complex cell's responses to the two stimuli individually. On closer analysis, they found that for a given stimulus pair, the

response to the pair was often similar in magnitude to the larger of the two responses evoked by each member of the pair. They referred to this behavior as MAX-like.

The consistent sublinear and MAX-like response patterns observed by Lampl et al. (2004) are not expected if paired stimuli interact in the receptive fields of pre-synaptic simple cells, as the energy model predicts they should. Previous extracellular studies have in fact uncovered such interactions, broadly confirming aspects of the energy model (Movshon et al., 1978b; Emerson et al., 1987; Szulborski and Palmer, 1990; Livingstone and Conway, 2003). As Lampl et al.'s intracellular results make a case for some form of inhibitory contribution to complex cell responses, while extracellular results can be accounted for by purely excitatory, feed-forward means, a reconciliation between the two types of study remains to be found. In addition, the exploration of MAX-like behavior in complex cells has particular interest because of a recent report from Serre et al. (2007). These authors showed that by making explicit use of MAX operations, a feed-forward model can equal human performance in a rapid categorization task.

Chapter III of this thesis thus explores complex cell responses to pairs of stimuli systematically and in great detail. Evidence is presented that both MAX-like and more Classical complex cells exist. A modified set of energy models is proposed to account for the variability in spatial integration observed over the population of complex cells.

Lateral inhibition in primary visual cortex

Considering all the suppressive response patterns observed in V1 cells, the phenomenon of surround suppression is arguably the most analogous to retinal lateral inhibition (Blakemore and

Tobin, 1972; DeAngelis et al., 1994; Li and Li, 1994). For some cells, a stimulus sized just to cover the receptive field (a "center" stimulus) produces more of a response than a larger stimulus that incorporates visual spatial locations outside of the receptive field (a "center + surround" stimulus). Lateral inhibition, if present in this scenario, would be considered spatial in origin, which is the form of lateral inhibition reported for the retina.

Unlike retinal lateral inhibition, however, cortical surround suppression is dependent on the form of the surround stimulus, and in particular on its orientation (Blakemore and Tobin, 1972; Cavanaugh et al., 2002b; Ozeki et al., 2004). Maximal surround suppression is evoked when an optimal center stimulus is combined with a surround stimulus of the same orientation. Shifting the surround orientation away from that of the center leads to weaker suppression. In some cases, suppression disappears entirely when the center and surround orientations are orthogonal to each other.

Drawing on the retinal analogy, increased intracortical inhibition has long been thought to mediate surround suppression, supplied by neurons with receptive fields in the surround of a given recorded neuron. Although possible, it is unlikely that surround suppression is inherited from the LGN, as was the case with cross-orientation suppression. The responses of LGN neurons do decrease when stimuli are enlarged beyond their receptive fields, but unlike in cortex the response decrease does not depend significantly on the orientation of the surround stimulus (Cleland et al., 1983b; Solomon et al., 2002; Nolt et al., 2004; Ozeki et al., 2004).

There are multiple possibilities for the manner in which intracortical connections could serve to produce surround suppression. The most direct possibility involves horizontal inhibitory connections within V1; however, these connections are generally restricted to fairly short distances (Buzas et al., 2001), making this possibility less likely. A second possibility is that surround suppression is accomplished via longer range, horizontal excitatory connections, which activate both inhibitory interneurons as well as other excitatory cells (McGuire et al., 1991; Angelucci et al., 2002; Stettler et al., 2002). In addition, these horizontal connections have been shown to connect cells with similar orientation preferences (Stettler et al., 2002). Finally, cortical areas outside of V1 could be the source of additional inhibition (Bair et al., 2003).

None of these possibilities directly challenge the essential architecture of Hubel and Wiesel's feed-forward model. However, finding that surround suppression is mediated by a sustained, orientation-dependent, increase in inhibitory conductance would be suggestive regarding how inhibition may contribute to other response properties of V1 neurons – for example, such inhibition could also underlie cross-orientation suppression. A single previous intracellular study looked at the effects of enlarging an optimal rectangular stimulus along an axis parallel to its orientation (Anderson et al., 2001). The authors reported that for some cells, relative to center stimulation alone, both inhibitory and excitatory conductances increased when the stimulus was enlarged. For other cells, the two conductances decreased. It is not clear what these results imply for classical surround suppression in which a stimulus is enlarged isotropically.

The origin of surround suppression thus remains an open question. Chapter IV of this thesis proposes a powerful model of cortical interconnectivity that gives rise to surround suppression as

a consequence of maintaining stability in the visual network. The model makes specific predictions for how inhibitory and excitatory conductances should change in response to surround stimulation. A number of these predictions were confirmed experimentally.

Chapter II

The emergence of contrast invariant orientation tuning in simple cells of cat primary visual cortex

Ian M. Finn[1,2], Nicholas J. Priebe[1,2], and David Ferster[1]

[1]Northwestern University Institute for Neuroscience
Department of Neurobiology and Physiology

[2]These authors contributed equally to this work

Northwestern University, Evanston IL 60208

Acknowledgements: We thank Kenneth. D. Miller for helpful discussions and Matteo Carandini, Stephen G. Lisberger, Larry Abbott and Wilson S. Geisler for comments on the manuscript. Supported by grants from the National Eye Institute (R01 EY04726) and National Institutes of Mental Health (P20 MH066239).

Abstract

Simple cells in primary visual cortex exhibit contrast-invariant orientation tuning, in seeming contradiction to feed-forward models relying on lateral geniculate nucleus (LGN) input alone. Contrast invariance has therefore been thought to depend on the presence of intracortical lateral inhibition. *In vivo* intracellular recordings instead suggest that contrast invariance can be explained by three properties of the excitatory pathway. 1) Depolarizations evoked by orthogonal stimuli are determined by the amount of excitation a cell receives from the LGN, relative to the excitation it receives from other cortical cells. 2) Depolarizations evoked by preferred stimuli saturate at lower contrasts than the spike output of LGN relay cells. 3) Visual stimuli evoke contrast-dependent changes in trial-to-trial variability, which lead to contrast-dependent changes in the relationship between membrane potential and spike rate. Thus, high-contrast, orthogonally-oriented stimuli that evoke significant depolarizations evoke few spikes. Together these mechanisms, without lateral inhibition, can account for contrast-invariant stimulus selectivity.

Introduction

In the classical view of sensory processing, generalized from Hartline's description of the limulus retina (Hartline, 1949), excitatory connections establish a bias in the selectivity of sensory neurons; lateral inhibition is then required to refine and sharpen this bias into the exquisitely selective responses sensory neurons often exhibit. According to this view, in the visual cortex excitatory, feed-forward connections from the lateral geniculate nucleus (LGN) establish the broad outlines of cortical receptive fields, including orientation bias and subfield organization (Hubel and Wiesel, 1962; Movshon et al., 1978a; DeAngelis et al., 1993a; Reid and Alonso, 1995), but these connections seem, on their own, to be unable to explain more subtle aspects of cortical responses, such as the sharpness of orientation tuning, cross-orientation suppression, and contrast invariance of orientation tuning. Lateral inhibition is thought to remedy the failures of the feed-forward model, either in the form of synaptic inhibition among neurons with different orientation tuning (Sompolinsky et al., 1990; Heeger, 1992; Somers et al., 1995; Troyer et al., 1998; Lauritzen and Miller, 2003; McLaughlin et al., 2003), or inhibition from neurons that are untuned for orientation (Hirsch et al., 2003; Lauritzen and Miller, 2003).

Despite the computational power of lateral inhibition, direct evidence that it shapes orientation selectivity in the cortex is equivocal (Ferster, 1986; Borg-Graham et al., 1998; Anderson et al., 2000b; Martinez et al., 2002). As an alternative to lateral inhibition, the failures of the feed-forward model can in part be accounted for by the inclusion of experimentally demonstrated nonlinear properties of the visual pathway, properties such as threshold, contrast saturation, synaptic depression and spike-rate rectification (Carandini and Ferster, 2000; Freeman et al.,

2002; Priebe and Ferster, 2006). Unlike orientation-specific inhibition, these nonlinearities (like untuned inhibition) are feature-blind: They operate independently of stimulus orientation, direction or size, but instead filter all signals as a function of stimulus strength (contrast) or response amplitude (spike rate).

This latter approach of incorporating non-linearities into the feed-forward model has been used to explain several fundamental aspects of simple-cell responses. The non-linearity of spike threshold can account for why simple cells' spike responses have sharper orientation tuning (Carandini and Ferster, 2000; Volgushev et al., 2000) and higher direction selectivity (Jagadeesh et al., 1997; Priebe and Ferster, 2005) than predictions derived from receptive field maps (DeAngelis et al., 1993a; Tolhurst and Heeger, 1997). Contrast saturation and spike-rate rectification of relay cells in the LGN can account for a large measure of cross-orientation suppression (Li et al., 2006; Priebe and Ferster, 2006).

One observation that remains difficult to reconcile with a purely feed-forward model is contrast invariance of orientation tuning (Skottun et al., 1987; Alitto and Usrey, 2004). As contrast increases, relay cell input to simple cells should increase at all orientations, including the orientation orthogonal to the preferred (Troyer et al., 1998; Ferster and Miller, 2000). Thus at higher contrasts, stimuli further and further from the preferred orientation, and ultimately at all orientations, should evoke suprathreshold depolarizations and elicit spikes, leading to a broadening of orientation tuning (the so-called iceberg effect). And yet, most simple cells respond with few or no spikes at the orthogonal orientation, and orientation tuning is largely

contrast invariant (Skottun et al., 1987; Anderson et al., 2000c; Ferster and Miller, 2000; Alitto and Usrey, 2004).

In models dependent on lateral inhibition, inhibitory input counteracts the excitatory relay-cell input that occurs at the null orientation, thus preventing a contrast-dependent broadening in orientation tuning. To ascertain whether tuned inhibition is required to refine cortical orientation tuning in this way, we recorded intracellularly from a large population of simple cells. We compared the contrast dependence of orientation tuning – both for membrane potential and spike rate – to the predictions of a feed-forward, excitation-only model based on the recorded behavior of geniculate relay cells. Consistent with the feed-forward model, many simple cells depolarized significantly in response to stimuli orthogonal to the preferred orientation. The amplitude of this depolarization was directly related to the fraction of direct synaptic input each cell received from the LGN. Thus, intracortical inhibition is not required to set the amplitude of the depolarization evoked by null stimuli. We also found that membrane potential responses to preferred stimuli saturated at lower contrasts than did spike responses of relay cells. These two properties had a significant effect on the contrast dependence of orientation tuning: tuning width did change with contrast, but less so than was expected from the feed-forward model.

Contrast-dependence of orientation tuning width was further reduced in the spike responses of simple cells by two features of the transformation between membrane potential and spike rate. The first is the expansive nonlinearity of threshold, previously described as a power-law (Hansel and van Vreeswijk, 2002; Miller and Troyer, 2002; Priebe et al., 2004), which amplifies small differences in membrane potential into large differences in spike rate. Second, we find that the

gain of the membrane-potential-to-spike-rate transformation is contrast dependent, falling with increasing contrast as a result of a concomitant fall in the trial-to-trial variability of responses. This change in gain helps prevent high-contrast stimuli of the non-preferred orientation from evoking spike responses, and consequently helps to generate contrast invariance in the spike responses of simple cells.

Our data thus support a model of orientation tuning in the visual cortex that operates without the need for lateral inhibition. Complex properties such as contrast-invariance can instead arise from the feed-forward pathway and its inherent nonlinearities.

Experimental Procedures

Animal preparation. Anesthesia was induced in young adult female cats with intramuscular Ketamine (10 mg/kg) and Acepromazine (0.7 mg/kg), and maintained with intravenous sodium thiopental (10 mg/kg initial, 2 mg/kg/hr maintenance). Paralysis was induced with intravenous vecuronium bromide (0.2 mg/kg/hr). Animals were artificially respirated through a tracheal cannula at a rate to maintain end tidal CO_2 at 4%. To stabilize the brain, the rib cage was suspended from a clamp on the cervical vertebrae and a bilateral pneumothoracotomy was performed. Rectal temperature was monitored and maintained at 38.3°C by a feedback controlled infrared lamp. EEG and EKG were monitored and rate of anesthesia was adjusted to maintain the regular occurrence of sleep spindles, and to prevent abrupt changes in heart rate. The pupils were dilated with atropine and nictitating membranes retracted with phenylephrine. The corneas were protected with contact lenses with 4-mm artificial pupils. Methods were approved by the Northwestern University Animal Care and Use Committee.

Recording. Patch recordings in current-clamp mode were obtained *in vivo* from area 17 of the visual cortex (within 5° of the representation of the area centralis). Electrodes were introduced through a craniotomy, which was protected by a solution of 3% agar in normal saline. Patch electrodes were filled with standard K^+-gluconate solution containing ATP and pH and calcium buffers as previously described. Signals were low-pass filtered, digitized at 4096 samples/sec and stored by computer using software written in LabVIEW (National Instruments, Austin, TX). Data were analyzed on-line to determine when enough trials had been performed to yield mean responses with low noise. Extracellular recordings were obtained from the cortex and the LGN

with lacquer-coated tungsten electrodes. Spikes were detected with a window discriminator and times of occurrence stored by computer (Bak DDS-2). Each neuron's receptive field was initially characterized by its tuning for location, size, orientation and spatial frequency using gratings. Using these preferences the response of the neurons was measured to a 2 or 4 second presentation of drifting gratings which varied in both contrast and orientation for cortical recordings or contrast alone for LGN recordings.

Visual stimulation. Drifting sinusoidal gratings of different orientation and contrast were displayed on a monitor using the Video Toolbox (Pelli, 1997) running in the Matlab environment on a Macintosh computer (Apple Computer, Cupertino, CA). Monitor mean luminance was 20 cd/m^2; refresh rate was 100 frames/sec and spatial resolution was 1024x768 pixels. The screen was placed at a distance of 48 cm from the cat's eyes and focused on the retina using auxiliary lenses and direct ophthalmoscopy.

Electrical stimulation. Electrical stimuli were delivered to the cortex through lacquer-coated tungsten electrodes with 200μ exposed tips. Electrodes were placed at a distance of 1 mm or less from the recording electrode and a depth of 400μ from the cortical surface. Stimuli were 200μs duration, electrode negative at an amplitude of 400μA or less. Previous experiments have shown that during the 50 ms following such stimuli, visual stimuli evoke no spikes in the surrounding cortical cells (Chung and Ferster, 1998). At the same time, visual responses of geniculate relay cells are unaffected.

Data analysis. All analyses were performed using custom software written in Matlab (The Mathworks, Natick, MA). For intracellular data, action potential and membrane potential responses were first segregated. Times of the occurrence of action potentials were determined using a simple threshold procedure. Action potentials were then removed from the membrane potential traces using a 5 ms median filter. The mean and standard deviation of the membrane potential were measured by aligning each response cycle, except the first cycle, binning the responses in 30 ms intervals, and computing the average and standard deviation at each time point.

Orientation Tuning. For almost any sensory stimulus in any modality, there are two independent aspects of tuning. The first is selectivity, the *amount* the response changes over the full range of stimuli, from best to worst. The second is width of tuning, or how rapidly the response falls off as the stimulus moves away from the preferred.

We most often use a Gaussian fit to the orientation tuning curve:

$$R(\theta) = B + A*\exp((\theta-\theta_p)^2/(2*\sigma^2))$$

where R is the response (membrane potential or spike rate), and is orientation. This formulation has the advantage that selectivity and tuning width are represented by two different parameters and are independent of one another: the selectivity is represented as null/preferred, or B/(B+A); tuning width is σ. Thus, cells with high selectivity can have either narrow or broad tuning (gray and black in the left graph below). Similarly, cells with low selectivity can have narrow or broad tuning (right graph).

We have shown previously that tuning width of the synaptic input to simple cells is largely dependent on the aspect ratio of the simple cells subfields (Lampl et al., 2001), and is then narrowed by threshold in the transformation to spike output, as had been postulated by many others (Palmer and Jones, 1984; DeAngelis et al., 1993b). In the feed-forward model (for gratings of infinite extent), stimulus selectivity in the synaptic input to simple cells depends only on the level of spontaneous activity in the presynaptic relay cells (Troyer et al., 1998). The lower the spontaneous activity, the lower the selectivity of the simple cell input, because of rectification of the relay cell responses.

Another traditional measure of orientation tuning Half-Width at Half Height (HWHH), where half-height is half the distance between the maximum response and 0 (not the offset, B). HWHH has the advantage of expressing orientation tuning as a single number, but has the disadvantage of depending both on selectivity and tuning width in a complex way. Circular variance, another commonly used measure of tuning, also depends on both selectivity and tuning width.

Contrast-response functions. Contrast response curves were fit using the Naka-Rushton curve ($R= A * C^n/(C^n + C_{50}^n)$). The C_{50} indicates the contrast at which a half-maximal response is generated. Fits were made on the peak response (F1+DC) to each cycle of the stimulus (except the first cycle, which was discarded) using the Gauss-Newton method. 95% confidence intervals for parameter estimates were computed from the Jacobian matrix and residuals using the Matlab function nlparci. Specific analyses are presented at the relevant places in the Results.

Power-law fits. Power law fits are based on the individual measurements of mean Vm and mean spike rate in each of the 30-ms epochs of the intracellular records. p and k in Equation 1 (or p, k_1 and k_2 in Equation 2) are adjusted until the summed least squared error between data and power law is minimized, using the Matlab function lsqcurvefit.

Results

Contrast dependence of orientation tuning in a simple feed-forward model

We begin by examining the properties of a purely linear feed-forward model in which we have expressly omitted inhibition of any type, even push-pull inhibition at the preferred orientation (Ferster, 1988; Hirsch et al., 1998). Our purpose here is to explore how well a purely excitatory feed-forward model can or cannot account for contrast invariance of orientation tuning in simple cells. The extent to which the model fails or succeeds would then lead to conclusions about how inhibition might or might not contribute to invariance. The comparison between model and data serves to highlight quantitatively where the recorded behavior of simple cells diverges from strict linearity, and what mechanisms might underlie this divergence.

Unlike in previous models (Somers et al., 1995; Troyer et al., 1998; Tao et al., 2004), we make no assumptions about the properties of geniculate relay cells, such as spontaneous activity, modulation amplitude, rectification, or contrast saturation. Instead, we constructed the model from the measured responses of geniculate X cells, recorded under the same conditions we used when recording from cortical simple cells. As a result, the model has only one free parameter -- the aspect ratio of the simple cell subfields -- which affects the width of orientation tuning, but has little effect on the response attribute we examine here: the change of tuning width with contrast.

To construct the model, we recorded extracellularly from 16 ON- and OFF-center geniculate X cells while presenting drifting gratings at 8 different contrasts (0, 4, 8, 12, 16, 20, 32, and 64%)

(Priebe and Ferster, 2006). We then averaged the responses of all 16 cells at each contrast (after shifting the response phases to be synchronous) and assigned these average responses to a template relay cell. The responses of 8 template ON-center relay cells with vertically offset receptive fields were summed to create the input to the simple cell from its ON subfield. Eight additional OFF-center relay cells were used to create the input from the OFF subfield (Fig. 2.01A, top; only 4 relay cells of each type are shown). The total relay-cell input was scaled so that an optimal, high-contrast stimulus evoked a 15-mV peak depolarization in the simple cell, similar to many recorded simple cells.

As in any feed-forward model, the alignment of the relay cell receptive fields makes the relative temporal phases of their responses dependent on the orientation of the stimulus: For the orientation orthogonal to the axis of displacement (null orientation), the relay cells respond asynchronously (Fig. 2.01B, 1st column, red and blue histograms). As a result, the null-oriented stimulus generates a rise in the mean potential, but no modulation component (Fig. 2.01B, 1st column, black histogram). For the preferred orientation, the ON and OFF relay cells respond synchronously (note that the red and blue ON- and OFF-center cell responses are therefore superimposed in Fig. 2.01B, 2nd column to make the violet histograms). As a result, the preferred stimulus generates a large modulation (F1 component) in the simple cell's membrane potential (Fig. 2.01B, 1st column, black histogram), which rides on top of a rise in mean potential (DC component). At lower contrasts, the responses are similarly shaped but with smaller amplitudes (Fig. 2.01A and B, 3rd and 4th columns).

At any given contrast, the modulation (F1) component of the relay cell input is tuned for orientation with a Gaussian-shaped tuning curve (Fig. 2.01C, left). The width of the tuning curve (σ = 32 degrees in Equation 3; see Methods) is dependent only on the aspect ratio of the subfields, with higher aspect ratios giving rise to narrower tuning (Palmer and Jones, 1984; DeAngelis et al., 1993b; Lampl et al., 2001). The aspect ratio we have chosen, 2.5:1, is the smallest for which the amplitude of the F1 component falls to 0 at an orientation of 90 degrees.

The mean (DC) component of the relay cell input – the total input to the simple cell averaged over one cycle of the grating – is independent of aspect ratio and is untuned for orientation (Fig. 2.01C, right) because the relay cells themselves are insensitive to orientation. Note that the DC component originates from rectification of the firing rate of relay cells. Visual stimuli modulate a relay cell's firing rate around its relatively low spontaneous rate, and so while the peak rate can increase more or less without bound, the trough is clipped at 0 spikes per second. As a result, for all but the lowest contrasts, the mean firing rate of relay cells increases with contrast (Troyer et al., 1998).

One measure we use below for quantifying the contrast invariance of orientation tuning is the peak of the simple cell response during a grating cycle, which is well approximated by summing the DC and F1 components. The tuning of the peak response therefore forms a roughly Gaussian shaped curve (the F1 component) riding on top of an offset from rest (the DC component), both of which increase with contrast (Fig. 2.01D).

To derive the spiking responses of the model simple cell, we first applied a threshold-linear transformation to the tuning curves of peak membrane potential (Fig. 2.01E and F). This transformation results in a strong dependence of tuning width on contrast, with a significant broadening as contrast increase. At the lowest contrasts the responses nearly disappear entirely; at the highest contrasts, spikes are evoked at all orientations. This broadening is the so-called iceberg effect of threshold and is in direct contradiction to the behavior of real simple cells: Few real simple cells respond with spikes to stimuli of the null orientation at any contrast, and the spike-rate responses of real simple cells show minimal contrast-dependent changes in orientation tuning width (Skottun et al., 1987; Anderson et al., 2000c; Alitto and Usrey, 2004).

A second, more realistic representation of the V_m-to-spike-rate transformation is a power law (Hansel and van Vreeswijk, 2002; Miller and Troyer, 2002; Priebe et al., 2004):

$$R(V_m) = k\lfloor \overline{V}_m - V_{rest} \rfloor_+^p \qquad (1)$$

where R is spike rate, $\overline{V}_m$ is trial-averaged membrane potential, V_{rest} is resting membrane potential, and the subscript, +, indicates rectification (R=0 for $V_m < V_{rest}$). The power law accounts for the effect of trial-to-trial variability by smoothing the threshold-linear relationship between mean membrane potential and mean spike rate (Anderson et al., 2000c). That is, even when a stimulus is weak and its mean response amplitude is far below physiological threshold, on a few trials the stimulus can carry the membrane potential above threshold and evoke spikes, leading to a small, but non-zero mean spike rate. In other words, variability smoothes (but does not completely linearize) the relationship between membrane potential and spike rate. We assign no theoretical significance to the power law but use it merely as a mathematical convenience to account for trial-to-trial variability. Other equivalent mathematical approaches are possible.

Smoothing the relationship between membrane potential and spike rate mitigates some of the effects of contrast on orientation selectivity, as shown in predictions of spike rate based on the power law (Fig. 2.01G-I). The widths of the resulting orientation tuning curves are much less dependent on contrast than those derived from the threshold-linear transformation. When normalized (Fig. 2.01H, right), the curves are closely superimposed. When displayed with an expanded vertical scale (Fig. 2.01I), however, the tuning curves reveal significant deviations from experimental results. First, a low-contrast stimulus of the preferred orientation evokes a smaller spike response than a high-contrast stimulus of the null orientation (Fig. 2.01I, left, red circles). Second, the curves still broaden visibly with increasing contrast (Fig. 2.01I, right).

Figure 2.01. **Contrast dependence of orientation tuning in a feed-forward model of simple cells**.

A. Receptive fields and responses (colored traces) for 8 of the 16 relay-cell inputs to the model simple cell.

B. Responses to both preferred and null-oriented gratings at high and low contrast are shown, as is the total input (black traces).

C. Orientation tuning curves for the F1 and DC components of the synaptic input to the simple cell.

D. Orientation tuning curve of the peak input to the simple cell (F1+DC).

E. A threshold-linear transformation between membrane potential and spike rate.

F. Orientation tuning curves (raw values and normalized) for peak spike rate as predicted by the threshold-linear transformation.

G. A power-law transformation between membrane potential and spike rate.

H. Same as E for the power-law transformation.

I. Same as G with amplified vertical scale.

J-L. Same as C, F and G with the DC component of the membrane potential response removed.

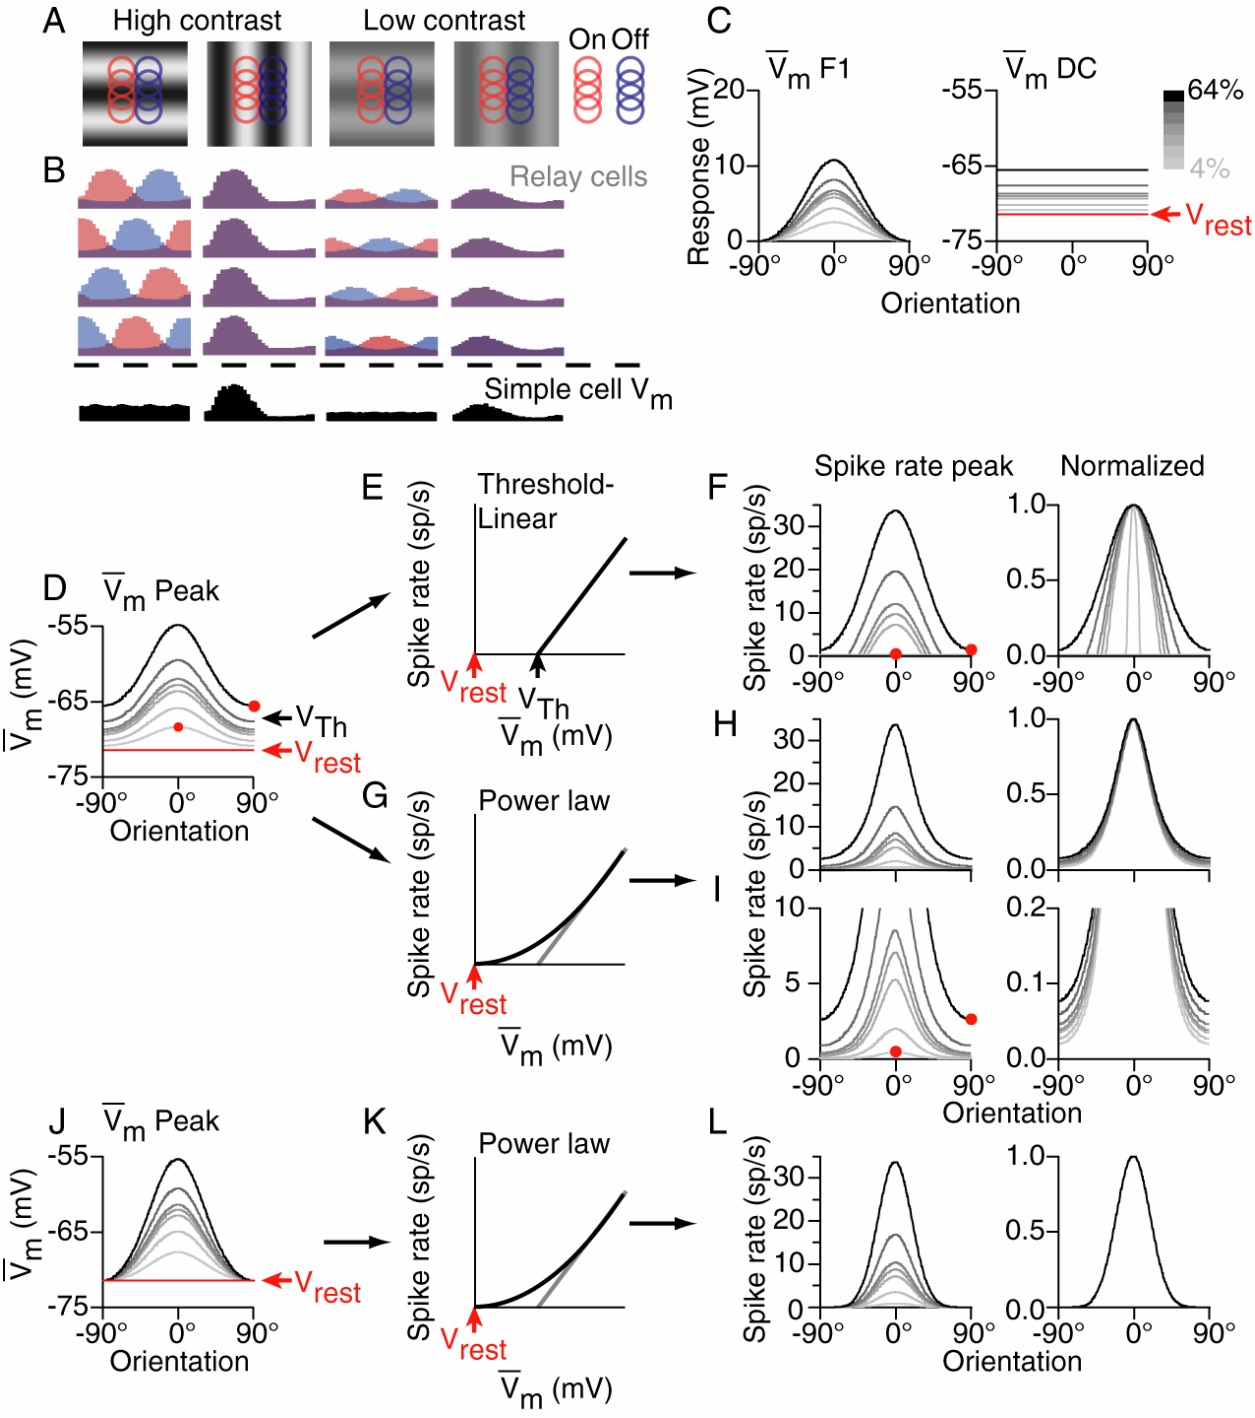
A
High contrast
Low contrast
On Off
B
Relay cells
Simple cell Vm
C
Vm F1
Vm DC
64%
4%
Vrest
Response (mV)
Orientation
D
Vm Peak
Vm (mV)
VTh
Vrest
Orientation
E
Threshold-Linear
Spike rate (sp/s)
Vrest
VTh
Vm (mV)
F
Spike rate peak
Normalized
G
Power law
H
I
J
Vm Peak
K
Power law
L

The response to high-contrast stimuli at the null orientation

In contrast to Figure 2.01D and I, for most simple cells we have studied intracellularly to date, high-contrast stimuli at the null-orientation evoke small membrane depolarizations, and very few (if any) spikes, relative to the preferred low-contrast response (Anderson et al., 2000c; Carandini and Ferster, 2000; Anderson et al., 2001). The orientation tuning curves for membrane potential in these cells look more like what is illustrated in Figure 2.01J-L. Here, the untuned DC components of the responses have been set to 0. Membrane potential responses are therefore perfectly contrast invariant (Fig. 2.01K), and the power law preserves invariance in the spike responses while narrowing the tuning width at all contrasts equally (Fig. 2.01L) (Miller, 1994; Anderson et al., 2000a; Hansel and van Vreeswijk, 2002). That the null response expected from the feed-forward model has apparently been suppressed in most cells has been attributed to cross-orientation inhibition or untuned inhibition (Sompolinsky and Shapley, 1997; Ferster and Miller, 2000).

To investigate whether the depolarization expected in response to null-oriented, high-contrast stimuli is consistently absent, we recorded intracellularly from 127 simple cells. Overall we found a wide range of behaviors, with some cells showing little depolarization at the null orientation (Fig. 2.02A), others showing moderate depolarization (B), and still others showing large depolarization (C). The model predicts that for any contrast the mean depolarization evoked by the null oriented grating (DC_N) should equal the mean depolarization evoked by the preferred orientation (DC_P). Within the recorded population, a significant number of cells echoed previous reports in showing a much smaller DC_N than DC_P (Fig. 2.02D, points below the unity line). Many points in Figure 2.02D did, however, fall on or near the unity line as predicted by the

feed-forward model of Figure 2.01. In only a small number of cells did the null-oriented stimulus cause a significant hyperpolarization of the membrane potential (Monier et al., 2003). The median DC_N/DC_P ratio for this population was 0.43 (Fig. 2.02E).

Figure 2.02. **Responses of simple cells to gratings of the preferred and orthogonal orientation**.

A-C. 8 cycles of response to a high-contrast drifting grating at the preferred (above) and orthogonal or null orientation (below) for three cells. Grating onset occurred after 250ms of blank stimulation.

D. The DC components of the responses to high-contrast gratings of the preferred and null orientation plotted against one another for 127 cells.

E. A histogram of ratios for the values in D.

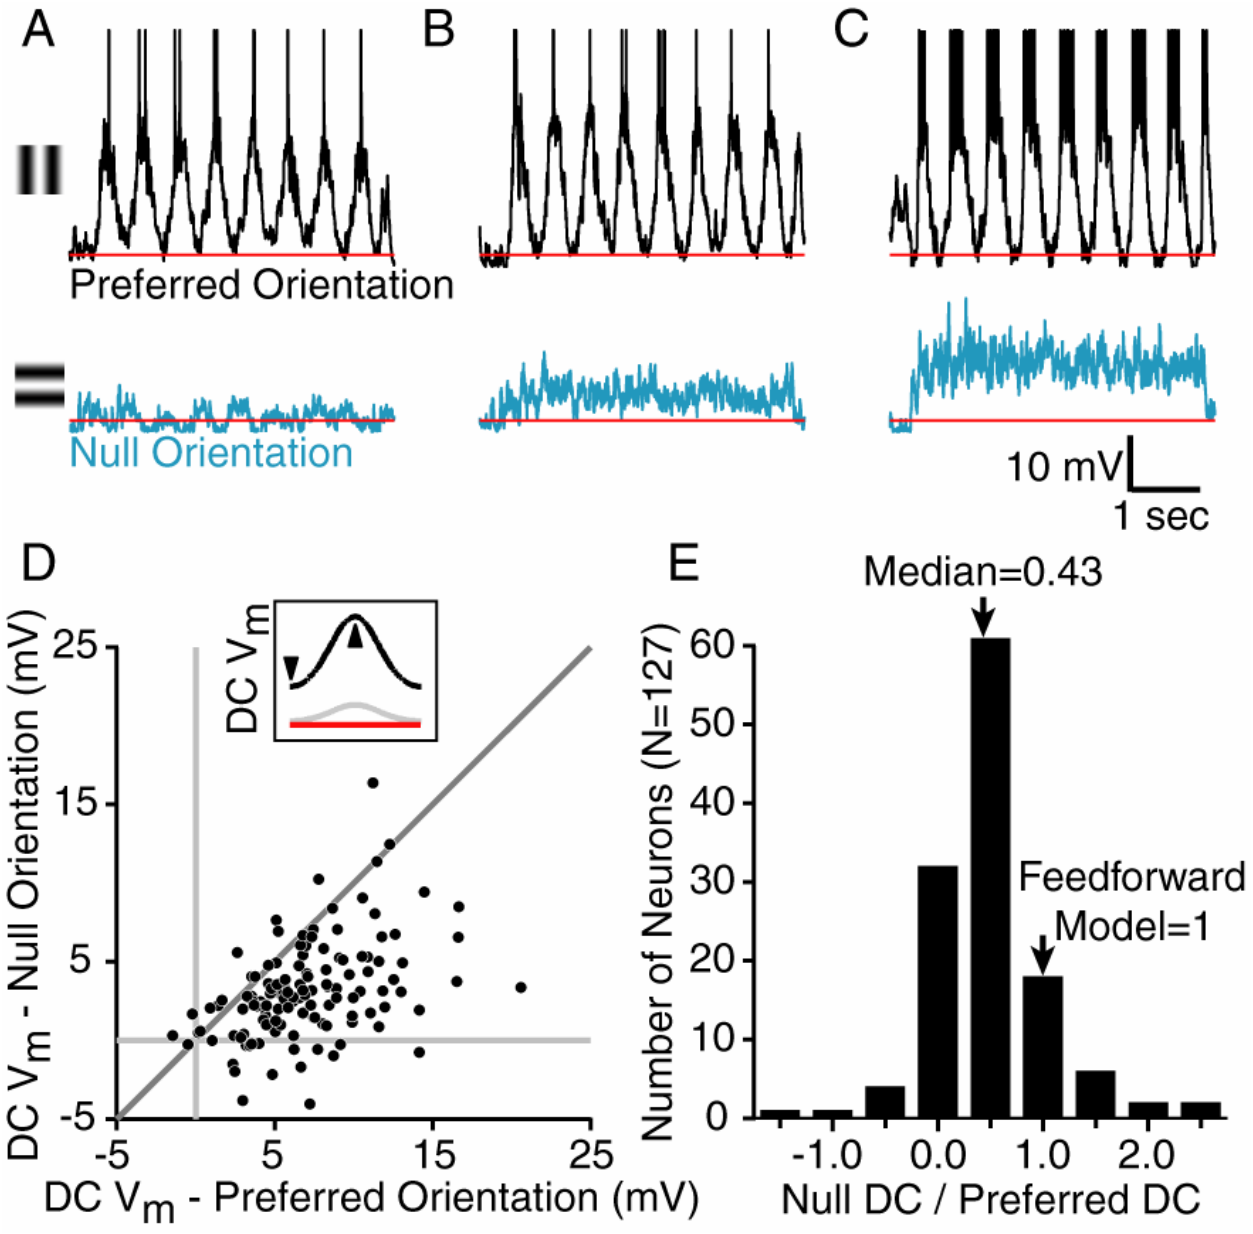
A
B
C
Preferred Orientation
Null Orientation
10 mV
1 sec
D
DC Vm
DC Vm - Null Orientation (mV)
DC Vm - Preferred Orientation (mV)
25
15
5
-5
E
Median=0.43
Feedforward
Model=1
Number of Neurons (N=127)
60
50
40
30
20
10
0
-1.0
0.0
1.0
2.0
Null DC / Preferred DC

The feed-forward model is based on the assumption that simple cells receive all of their excitatory input from geniculate relay cells. In reality, each simple cell receives a different proportion of its excitatory input from the LGN, with the remainder coming from other cortical cells (Chung and Ferster, 1998). If intracortical connections are formed among cells with similar preferred orientation, simple cells with a large fraction of cortical excitatory input should exhibit small DC_N/DC_P ratios because cortical cells respond little to null-oriented stimuli (Chung and Ferster, 1998). Conversely, simple cells that receive the bulk of their input from the LGN should exhibit large DC_N/DC_P ratios. We tested this expectation for 19 cells, measuring the relationship between DC_N/DC_P and the fraction of excitatory input the cell received from the LGN (%LGN Input).

The %LGN Input was measured for each simple cell by suppressing the responses of cortical neurons with electrical stimulation (Chung and Ferster, 1998). We presented a 20-ms flash of a high-contrast grating of optimal size, spatial frequency, spatial phase and orientation, with and without paired electrical stimulation of the nearby cortex. The electrical stimulus evokes a large IPSP in every nearby cortical cell and prevents spiking in response to the visual stimulus. To prevent the electrical stimulus from antidromically activating geniculate relay cells (Chung and Ferster, 1998), the stimulating electrode was inserted no deeper than 400 μm below the cortical surface, and stimulus amplitudes were kept in the range of 0.25-0.45mA (200μs duration, , electrode negative < 1mm distant from recording electrode). The response to paired electrical and visual stimulation (Fig. 2.03A-C, brown traces; the response to electrical stimulation alone has been subtracted off) is therefore dominated by direct, monosynaptic LGN input. The %LGN

Input is taken to be the amplitude of the paired response divided by the amplitude of the response to the flashed grating alone (Fig. 2.03A-C, top, black traces).

The DC_N/DC_P ratio was well correlated with the %LGN Input. The DC_N/DC_P ratios for the three cells of Figure 2.03A-C were 0, 0.45, and 0.85; their %LGN Input was 4%, 44%, and 86%. A scatterplot of DC_N/DC_P against %LGN Input for the 19 cells showed a strong correlation (Fig. 2.03D, $R^2 = 0.79$, slope = 0.76, Y-intercept = 0.13), with most of the points lying close to the unity line. A broad range of %LGN Input received by each cell can thus account for why the DC component of the grating response was often orientation tuned, i.e., why the average depolarization evoked by null-oriented stimuli in Figure 2.02D and E was often smaller than the average depolarization evoked by preferred stimuli. This result is diagramed in the cartoon of Figure 2.03E. The orientation tuning of relay cell responses, and the resulting input to a simple cell, is shown at the far left. In the center is depicted the input to a cell that receives 50% of its excitatory input from the LGN (top) and 50% from other cortical cells (bottom). These sum to produce the input pictured on the right. Thus, replacing some geniculate excitation with cortical excitation (from cells with similar preferred orientations) reduces the response at the null orientation while leaving the response at the preferred orientation unchanged.

An immediate question raised by the data in Figures 2.02 and 2.03 is whether simple cells with large a large DC_N/DC_P (large %LGN input) reside in cortical layer IV, the primary layer in which geniculate afferents terminate. An analysis of the data based on electrode depth measurements does show a significant negative correlation between electrode depth and DC_N/DC_P – superficial simple cells do not respond to orthogonally oriented gratings as much as

simple cells deeper in cortex (Fig. 2.04A). Figure 2.04B suggests that the average DC_N/DC_P is highest around 750 μm, which lies within upper layer IV according to slice studies (Toyama et al., 1981; Ferster and Lindström, 1983).

We note parenthetically that shunting inhibition evoked by the shock stimulus, in addition to inactivation of cortical inputs, could in theory reduce the size of the response to the flash, making the %LGN Input appear to be smaller than it actually was. This is likely not the case. 1) Shock-evoked conductance changes (Anderson et al., 2000b) are not likely to be that much larger than changes evoked by the flash alone (Hirsch et al., 1998). 2) %LGN Input near 100% would likely not be observed. 3) The average %LGN Input observed here is comparable to that measured in cortical cooling experiments (Ferster et al., 1996), which are not subject to shunting effects. We therefore take the shock-induced reduction in flash response as a reasonable approximation of geniculate input.

Figure 2.03. **The relationship between the response to null-oriented stimuli and the amount of input from the LGN**.

A-C. Top, responses to optimal flashed gratings with (brown) and without (black) paired electrical stimulation of nearby cortex for 3 cells. The response to electrical stimulation alone has been subtracted from the brown traces. The ratio of the amplitudes of the brown and black traces (F&S/F) is taken to be the proportion of synaptic input the cell receives directly from the LGN. The cell in A receives almost no direct input from the LGN; the cell in C receives almost exclusive input from the LGN. Middle and bottom, responses to high-contrast drifting gratings of the preferred and null orientation for the 3 cells. Inset in B shows 20 superimposed responses to electrical stimulation alone.

D. The ratio of responses to null and preferred stimuli (DC component) plotted against the proportion of input provided by the LGN (N=19).

E. Left, orientation tuning curves for the combined output from the relay cells exciting the model simple cell that receives input only from the LGN. Right, orientation tuning curves for a cell that receives half it's input from the LGN and half from other cortical cells with similar preferred orientation. The main effect is to reduce the response of the cell to stimuli of the null orientation.

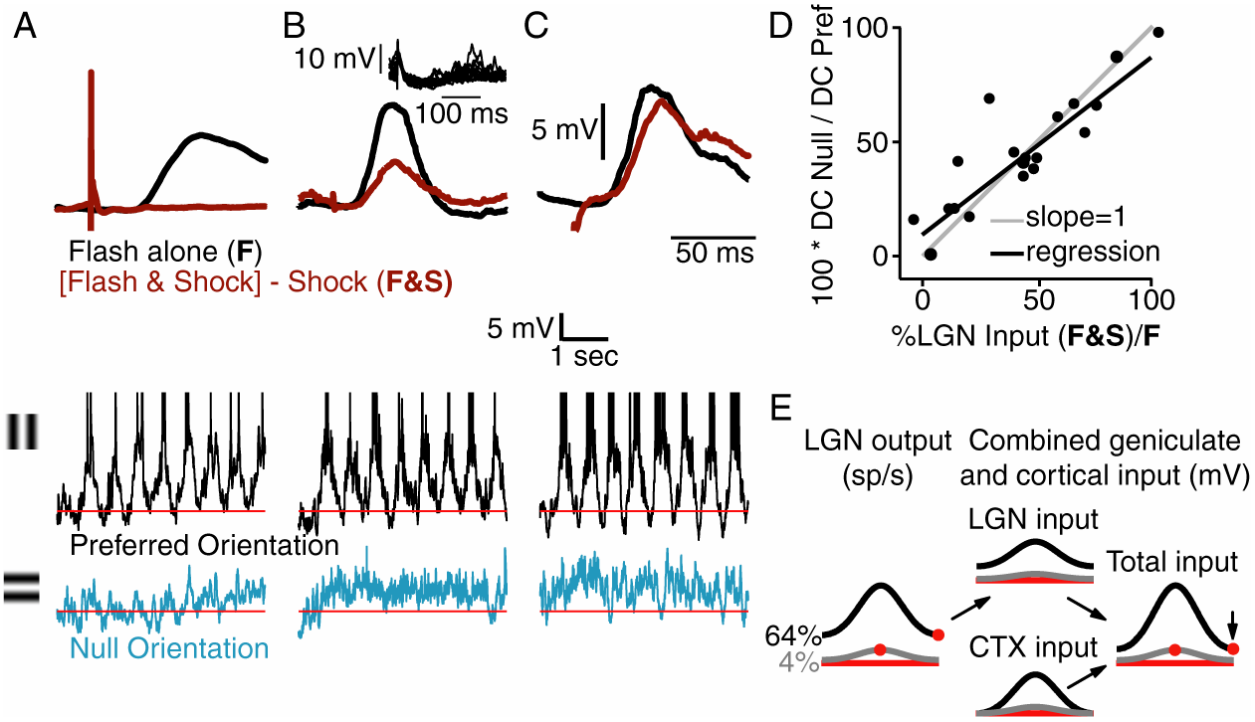
A
B
C
D
10 mV
100 ms
5 mV
50 ms
Flash alone (F)
[Flash & Shock] - Shock (F&S)
5 mV
1 sec
Preferred Orientation
Null Orientation
100 * DC Null / DC Pref
100
50
0
0
50
100
slope=1
regression
%LGN Input (F&S)/F
E
LGN output (sp/s)
Combined geniculate and cortical input (mV)
LGN input
Total input
64%
4%
CTX input

Figure 2.04. **Electrode depth and DCn / DCp**.

Left. Scatterplot of electrode depth versus DC_n / DC_p for a population of simple cells (N = 120). There is a significant negative correlation ($R = -.28$, $p < .005$).

Right. Data from A binned according to depth in 127.5 μm increments. Average DC_N/DC_P (dots) and S.E.M shown (blue lines) for each bin.

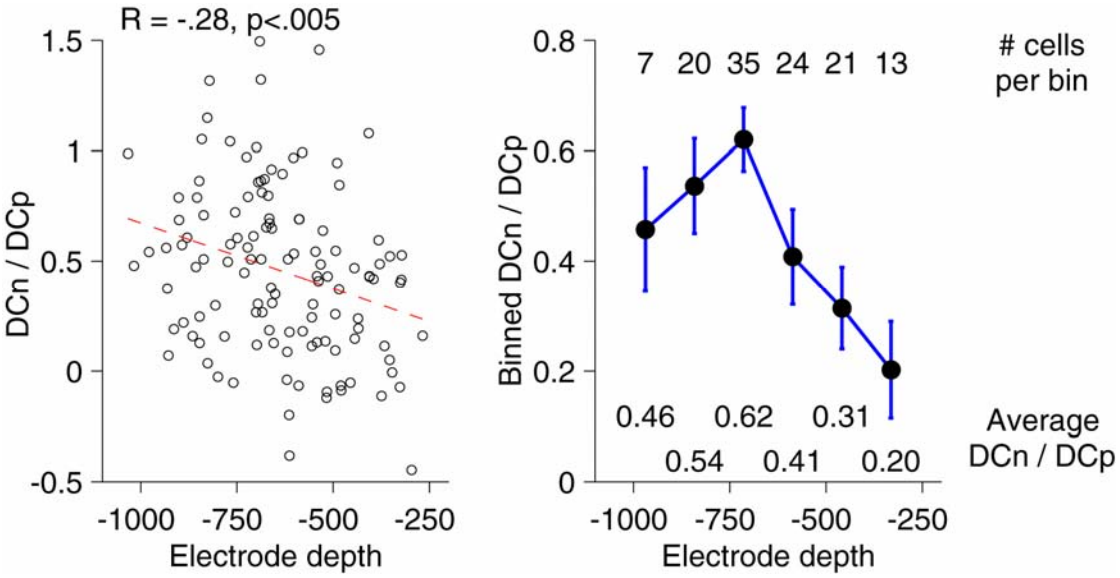
R = -.28, p<.005
1.5
1
0.5
0
-0.5
DCn / DCp
-1000
-750
-500
-250
Electrode depth
0.8
0.6
0.4
0.2
0
Binned DCn / DCp
7 20 35 24 21 13
cells per bin
0.46 0.62 0.31
0.54 0.41 0.20
Average DCn / DCp
-1000
-750
-500
-250
Electrode depth

Contrast saturation at the preferred orientation

The feed-forward model of Figured 2.01 makes specific predictions about the contrast dependence of membrane potential responses in simple cells. Specifically, membrane potential should saturate with contrast in the same way that geniculate relay cell spike responses do. We therefore compared the contrast saturation of membrane potential responses in 46 simple cells with the spike responses of 45 geniculate X cells. The X cell in Figure 2.05A and B reached half-maximal response amplitude (C_{50}) at 29% contrast. Overall, LGN cells had C_{50}'s between 4% and 35% contrast, with a median C_{50} of 15.9% (Fig. 2.05C). Simple cell membrane potential responses differed significantly from LGN cells. The membrane potential responses of the cell in Figure 2.05D and E reached nearly complete saturation at 16% contrast. The median C_{50} for simple cells was 7.6% (Fig. 2.05F).

Possible sources for the early simple cell contrast saturation include depolarization-induced reductions in driving force on synaptic currents, and activity-dependent synaptic depression (Kayser et al., 2001; Carandini et al., 2002), which can reach nearly 50% in geniculocortical synapses (Bannister et al., 2002; Boudreau and Ferster, 2005) and 80-90% in corticocortical synapses (Stratford et al., 1996; Abbott et al., 1997; Tsodyks and Markram, 1997). While not explicitly tested here, both of these potential mechanisms would affect off-orientation responses similarly to the preferred orientation response.

The early contrast saturation of membrane potential at the preferred orientation has important consequences for contrast invariance. The contrast can be lowered far more than would be expected from the feed-forward model before the synaptic input falls significantly. In other

words, real tuning curves – both for membrane potential and spike rate – change much less with contrast than is shown in Figure 2.01D (Fig. 2.05G).

Figure 2.05. **Contrast saturation in LGN and cortex**.

A. Spike responses of a geniculate relay cell to drifting gratings of different contrast.

B. Contrast response curve constructed from the peak (F1+DC) responses in A.

C. A histogram of C_{50}'s for 45 relay cells.

D-**F**. Same as A-C for the peak membrane potential responses of 46 cortical simple cells.

G. Same as Fig. 2.03E, but with the addition of early contrast saturation. The effect of early saturation is to raise the amplitude of responses to low-contrast stimuli of the preferred orientation.

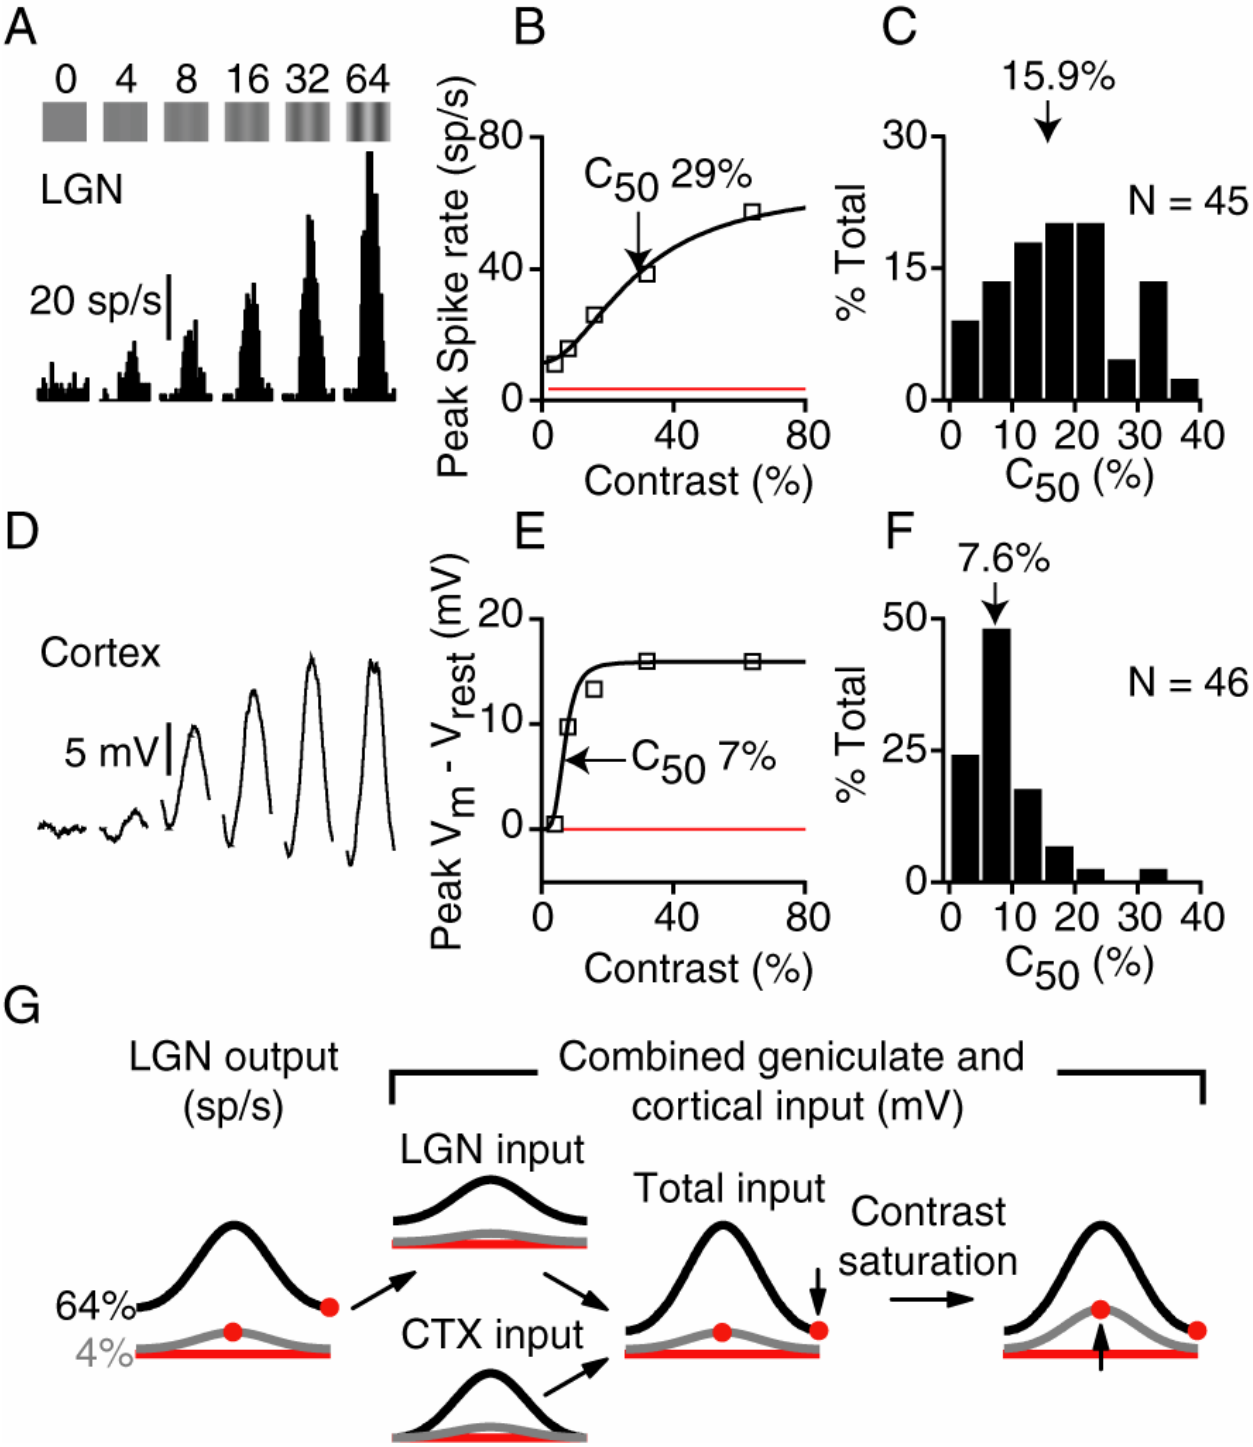
A
0 4 8 16 32 64
LGN
20 sp/s
B
Peak Spike rate (sp/s)
80
40
0
C50 29%
0 40 80
Contrast (%)
C
15.9%
30
15
0
% Total
N = 45
0 10 20 30 40
C50 (%)
D
Cortex
5 mV
E
Peak Vm - Vrest (mV)
20
10
0
C50 7%
0 40 80
Contrast (%)
F
7.6%
50
25
0
% Total
N = 46
0 10 20 30 40
C50 (%)
G
LGN output (sp/s)
Combined geniculate and cortical input (mV)
LGN input
Total input
Contrast saturation
64%
4%
CTX input

Comparison of responses to high-contrast null stimuli and low-contrast preferred stimuli

The principle difficulty for contrast invariance raised by the model of Figure 2.01 is the relationship between responses to high-contrast, null-oriented stimuli and low-contrast, preferred stimuli: The former are predicted to be larger than the latter (Fig. 2.01D, red dots), whereas in real simple cells – at least for spiking responses – the opposite is true (Fig. 2.06). In Figures 2.02-5, we illustrated two features of simple-cell input that tend to mitigate this problem – one, a smaller depolarization at the null orientation than predicted, and two, earlier than expected contrast saturation. As a result, for most cells, the ratio of the null response at high-contrast (64%) to the preferred response at low-contrast (4%, 8% or 12%; points in Fig. 2.06A) was lower than expected from the feed-forward model (lines of corresponding color). Nevertheless, the problem remains as to why the depolarizations, though of reduced amplitude, evoked almost no spikes. Low-contrast preferred stimuli evoked depolarizations on average only about twice the size of those evoked by the null orientation (Fig. 2.06A), and yet they evoked vastly more spiking (Fig. 2.06B). In several simple cells of Fig. 2.06A, the membrane potential responses to high-contrast null stimuli and low-contrast preferred stimuli were nearly equal, and yet there are no cells in Fig. 2.06B with comparable spiking responses to the two stimuli. A single expansive nonlinearity is insufficient to account for this differential amplification, thus suggesting that the transformation between membrane potential and spike rate depends on contrast.

Figure 2.06. **Lack of spiking responses to high-contrast stimuli of the null orientation**.

A and **B**. The response (A - membrane potential; B - spike rate) to high-contrast stimuli of the null orientation plotted against the response to low-contrast stimuli of the preferred orientation. Symbols of different shades of gray indicate the contrast of the low-contrast stimulus. Lines indicate the predictions of the feed-forward model in Fig. 2.01.

C. Same Figs. as 3E and 4G, with the addition of the V_m-to-spike-rate transformation, which differentially amplifies the responses to high-contrast preferred and low-contrast null stimuli while narrowing the tuning curves equally.

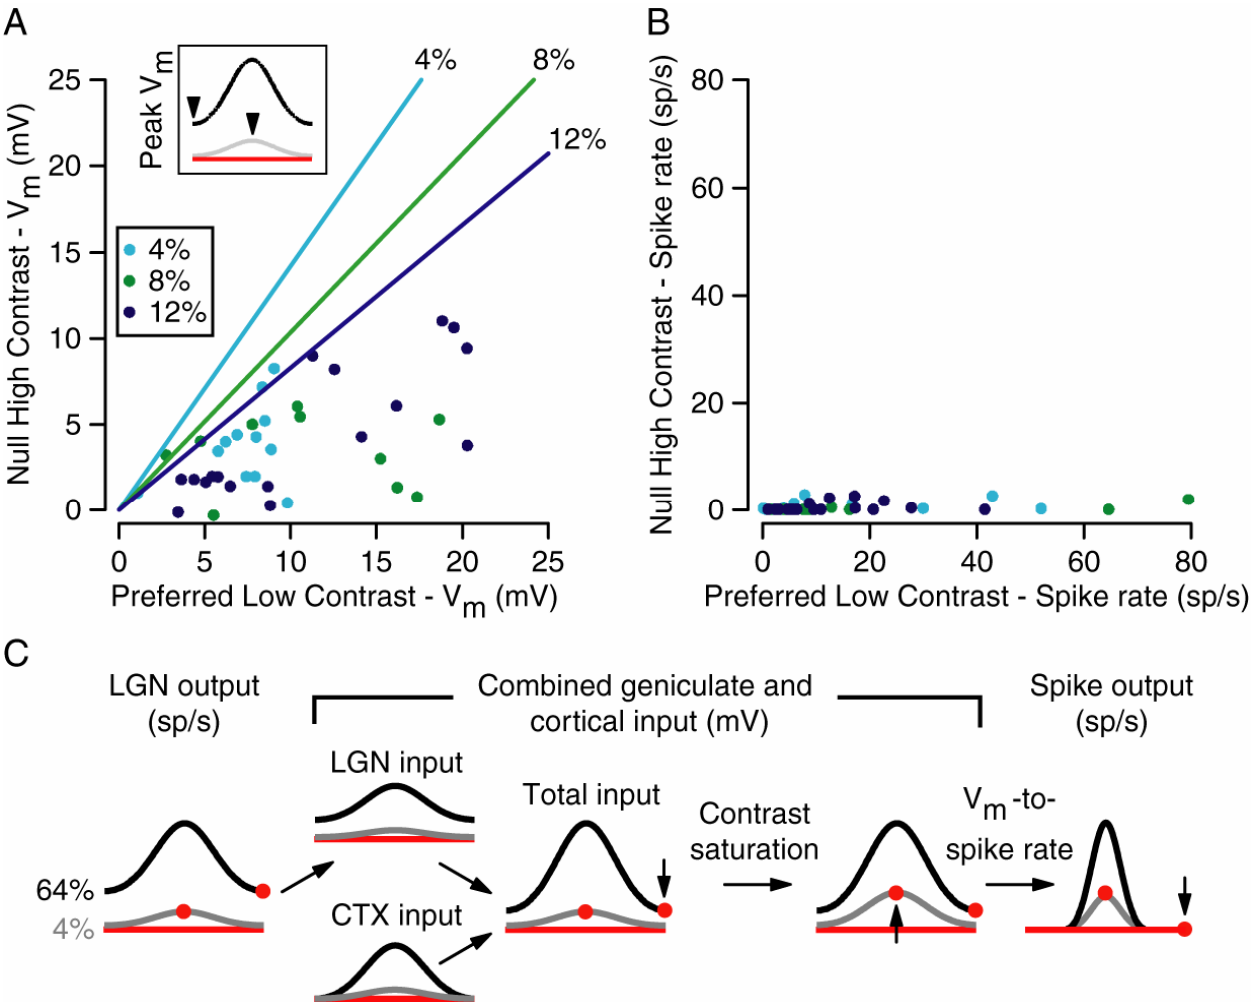
A
Peak V_m
4%
8%
12%
4%
8%
12%
Null High Contrast - V_m (mV)
Preferred Low Contrast - V_m (mV)
0
5
10
15
20
25
B
Null High Contrast - Spike rate (sp/s)
Preferred Low Contrast - Spike rate (sp/s)
0
20
40
60
80
C
LGN output (sp/s)
Combined geniculate and cortical input (mV)
Spike output (sp/s)
LGN input
Total input
Contrast saturation
V_m -to-spike rate
64%
4%
CTX input

The effects of contrast on the transformation between membrane potential and spike rate

Spike rate is plotted against membrane potential separately for high- and low-contrast stimuli in Figure 2.07A. Here, the cycle-averaged responses of all orientations were divided into 30-ms epochs, and mean spike rate was plotted against mean membrane potential for each epoch. As hinted at in Figure 2.06, a given mean depolarization at high contrast evoked fewer spikes than the same depolarization at low contrast. We quantified this trend by comparing spike rate responses for stimuli of high and low contrast that evoked the same mean membrane potential (Fig. 2.07B; 2-mV bins for membrane potential). Judging by the average slope of the plot in Figure 2.07B, on average a high-contrast stimulus evoked 62% of the spikes evoked by a low-contrast stimulus that gave rise to the same mean depolarization. Across 39 cells, this slope ranged from 0.01 to 1.2, with a median value of 0.49. Thus, stimulus contrast changed the gain of the membrane-potential-to-spike-rate transformation by a factor of about 2.

To understand the source of this change in gain between high and low contrast, it is important to note that the relationship plotted in Figure 2.07A is the *mean* membrane potential against *mean* spike rate, averaged across stimulus trials. As shown previously, this relationship approximates a power law (Eq. 1), or a threshold-linear curve smoothed by trial-to-trial variability (smooth curves in Fig. 2.07A). The parameters of the power law are mainly determined by (1) the resting membrane potential of the cell, (2) the cell's threshold, and (3) the amount of noise or trial-to-trial variability in the membrane potential responses (Chance et al., 2002; Hansel and van Vreeswijk, 2002; Miller and Troyer, 2002; Carandini, 2004). Of these three properties, biophysical threshold is unlikely to vary systematically with contrast. The resting potential is

also unlikely to change: The main stimulus-related influence on resting potential is contrast adaptation (Carandini and Ferster, 1997; Sanchez-Vives et al., 1997), which should not be a factor here because we randomly interleaved trials of different contrast. Since visual stimuli can have an effect on trial-to-trial variability (Monier et al., 2003), we speculated that contrast-dependent changes in gain might arise from contrast-dependent changes in trial-to-trial variability.

The effects of stimulus contrast on trial-to-trial variability are shown for one cell in Figure 2.07D-G. Here we refer to a stimulus trial as one complete cycle of the grating, six of which are shown in Figure 2.07D for three different stimuli. For the preferred orientation at high contrast (black), the cell responded consistently to each cycle of the grating with a 15-20 mV sinusoidal depolarization, giving rise to a cycle averaged response of similar peak amplitude (Fig. 2.07E black). To quantify the trial-by-trial variability, we median-filtered the membrane potential to remove spikes, smoothed the traces with a 30-ms sliding window, and then measured the trial-by-trial standard deviation at each point in time relative to the start of the cycle. The standard deviation is indicated in Figure 2.07E by gray shading surrounding the black trace, and in this case was relatively small compared to the size of the depolarization. For the preferred stimulus at low contrast (Fig. 2.07D, blue), the membrane potential response varied considerably from trial to trial: the response was almost as large as the high-contrast response on the 3rd trial, but only a small fraction of that on the 4th trial. As a result, the average peak membrane potential response was about half that evoked by the high-contrast grating (Fig. 2.07E, blue), yet the standard deviation of the membrane potential was far larger than its high-contrast counterpart (blue shading). This change in response variability is clearly a function of the stimulus and not of the

response amplitude: A high-contrast grating at the null orientation (Fig. 2.07D, green) evoked an average response that was comparable in peak amplitude to the low-contrast preferred response, yet had a standard deviation comparable to the high-contrast, preferred response (Fig. 2.07E, green trace and shading).

The effect of membrane potential trial-to-trial variability on spike rate can be seen in Figure 2.07F. At almost every point in time, the low-contrast preferred stimulus evoked a smaller average depolarization than the high-contrast null stimulus (Fig. 2.07E; green vs. blue), and yet it triggered more spikes because its higher variability more often carried the membrane potential above threshold (Fig. 2.07F; green vs. blue).

That trial-to-trial variability changes consistently with contrast is shown for this example cell in Figure 2.07G. When the peak amplitudes (F1+DC) for all of the individual stimulus trials were plotted against orientation for the cell in Figure 2.07D-F, the vertical spread of points was visibly greater at low contrast (compare Fig. 2.07G, left and right). To quantify the relationship between stimulus contrast and trial-to-trial variability across the population, for each cell we measured the trial-to-trial standard deviation of peak response amplitude at each contrast and orientation. Standard deviation at high contrast is plotted against standard deviation at low contrast for the preferred orientation in Figure 2.07H, and for the null orientation in Figure 2.07I. For both orientations, variability was, on average, higher at low contrast (51% more at the preferred and 30% more at the null). That trial-to-trial variability of the peak response amplitude is higher for low-contrast stimuli of other orientations is shown in Figure 2.08. The figure also illustrates, as

observed previously (Monier et al., 2003), that the trail-to-trail variability of responses to high-contrast stimuli was reduced relative to the blank stimulus.

A preliminary analysis looking at the relationship between trial-to-trial variability and amount of geniculate input received is shown in Figure 2.09. Fourteen simple cells recorded at both 4% and 64% contrast were considered. The top panels depict orientation tuning of the standard deviation of peak membrane potential, averaged across all 14 cells (top left), across 5 cells with small DC_N/DC_P ratios (< 0.3, top middle), and across 5 cells with large DC_N/DC_P ratios (> 0.75, top right). All cells were normalized to the value at their preferred orientation before averaging. Average membrane potential tuning curves for each population are depicted below. These data suggest that the difference in variability between low and high contrast is greater for cells receiving the majority of their excitation from the LGN.

Figure 2.07. **The contrast dependence of trial-to-trial variability and its effect on mean spike rate**.

A. The relationship between mean spike rate and mean membrane potential plotted separately for low-contrast and high-contrast stimuli in one simple cell. Each point is derived from one 30-ms epoch of a trial-averaged response (13 stimuli, 16 epochs each). Solid curves are power-law fits (Equation 1) to the data.

B. Average spike rate at high contrast plotted against spike rate at low contrast for each of 8 ranges of mean membrane potential in A. Solid line is a linear regression.

C. Slope of the regression (as in B) for 39 cells.

D. Six cycles of the responses of a simple cell to high- and low-contrast gratings of the preferred orientation (black and blue) and to a low-contrast grating of the null orientation (green).

E. Cycle-averages of the responses to the three stimuli, with standard deviation indicated by shading. The mean and standard deviation of the membrane potential were computed using a 30 ms sliding window.

F. Average spike responses for the three stimuli.

G. Orientation tuning curves for the peak (F1+DC) response of the cell at high and low contrast. Each point represents the peak response to a single cycle.

H. The trial-to-trial standard deviation of peak response amplitudes for low-contrast gratings plotted against the standard deviation for high-contrast gratings at the preferred and null orientations (52 cells).

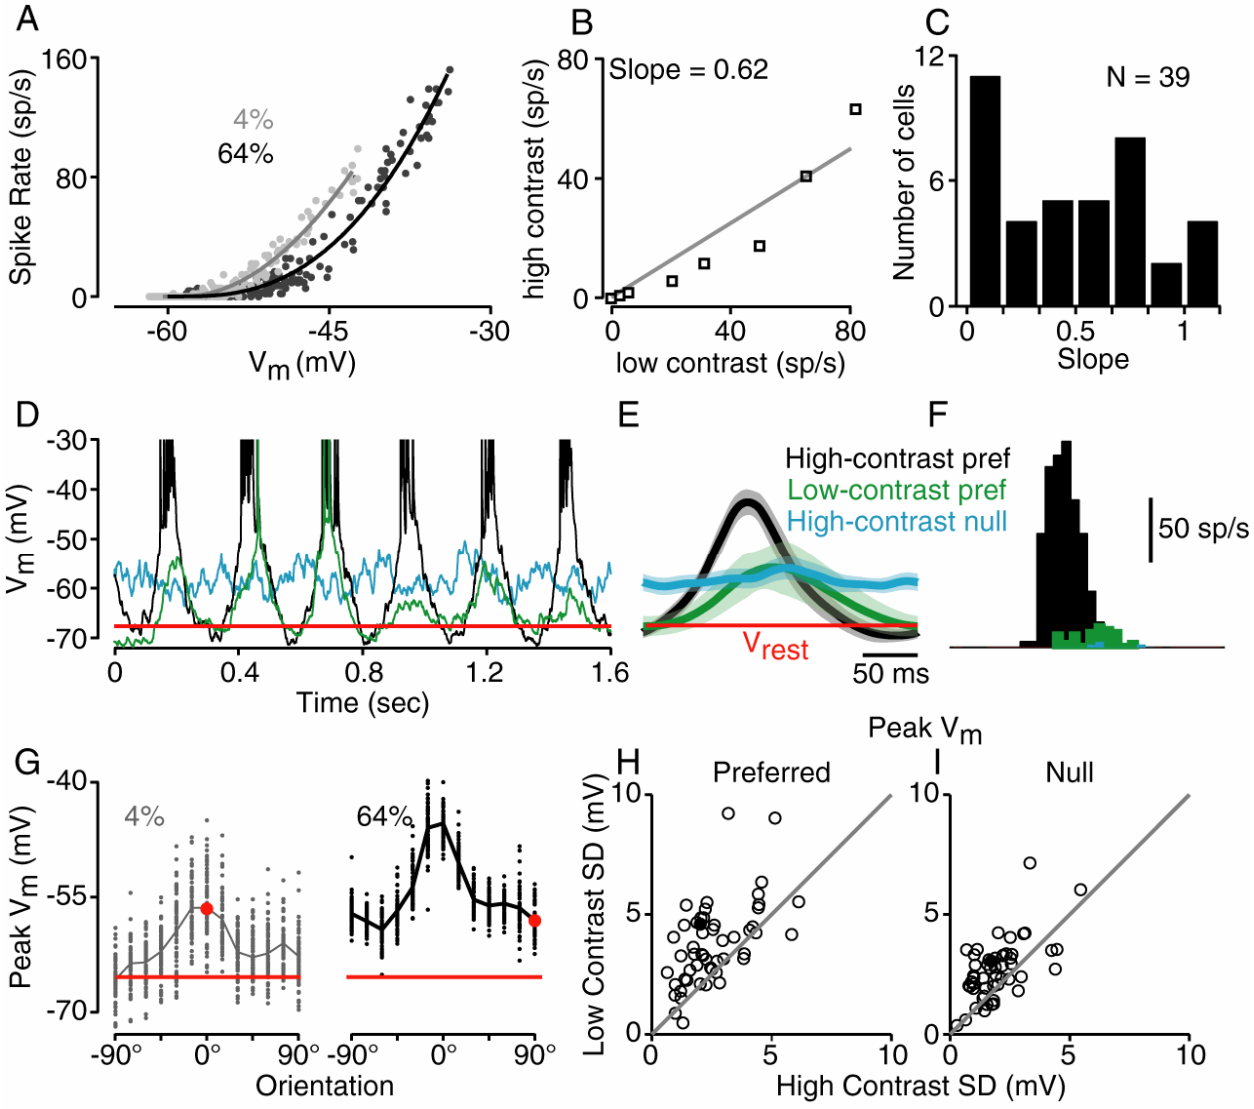
A
Spike Rate (sp/s)
160
80
0
4%
64%
-60
-45
-30
Vm (mV)
B
high contrast (sp/s)
Slope = 0.62
80
40
0
0
40
80
low contrast (sp/s)
C
Number of cells
12
6
0
N = 39
0
0.5
1
Slope
D
Vm (mV)
-30
-40
-50
-60
-70
0
0.4
0.8
1.2
1.6
Time (sec)
E
High-contrast pref
Low-contrast pref
High-contrast null
Vrest
50 ms
F
50 sp/s
G
Peak Vm (mV)
-40
-55
-70
4%
64%
-90°
0°
90°
Orientation
H
Peak Vm
Preferred
Low Contrast SD (mV)
10
5
0
High Contrast SD (mV)
I
Null

Figure 2.08. **Contrast and orientation dependence of trial-to-trial variability**.

A and **B**. Example membrane potential and standard deviation of membrane potential tuning curves at low and high contrast. Curves are Gaussian fits from which both peak and background parameters can be extracted.

C. Average SD at -90^{o}, -30^{o}, 0^{o}, 30^{o}, and 90^{o} over 52 cells. Each cell's SD was normalized to the maximum membrane potential response.

D. Plot of peak SD values extracted from fits like those in B.

E. Plot of background SD values.

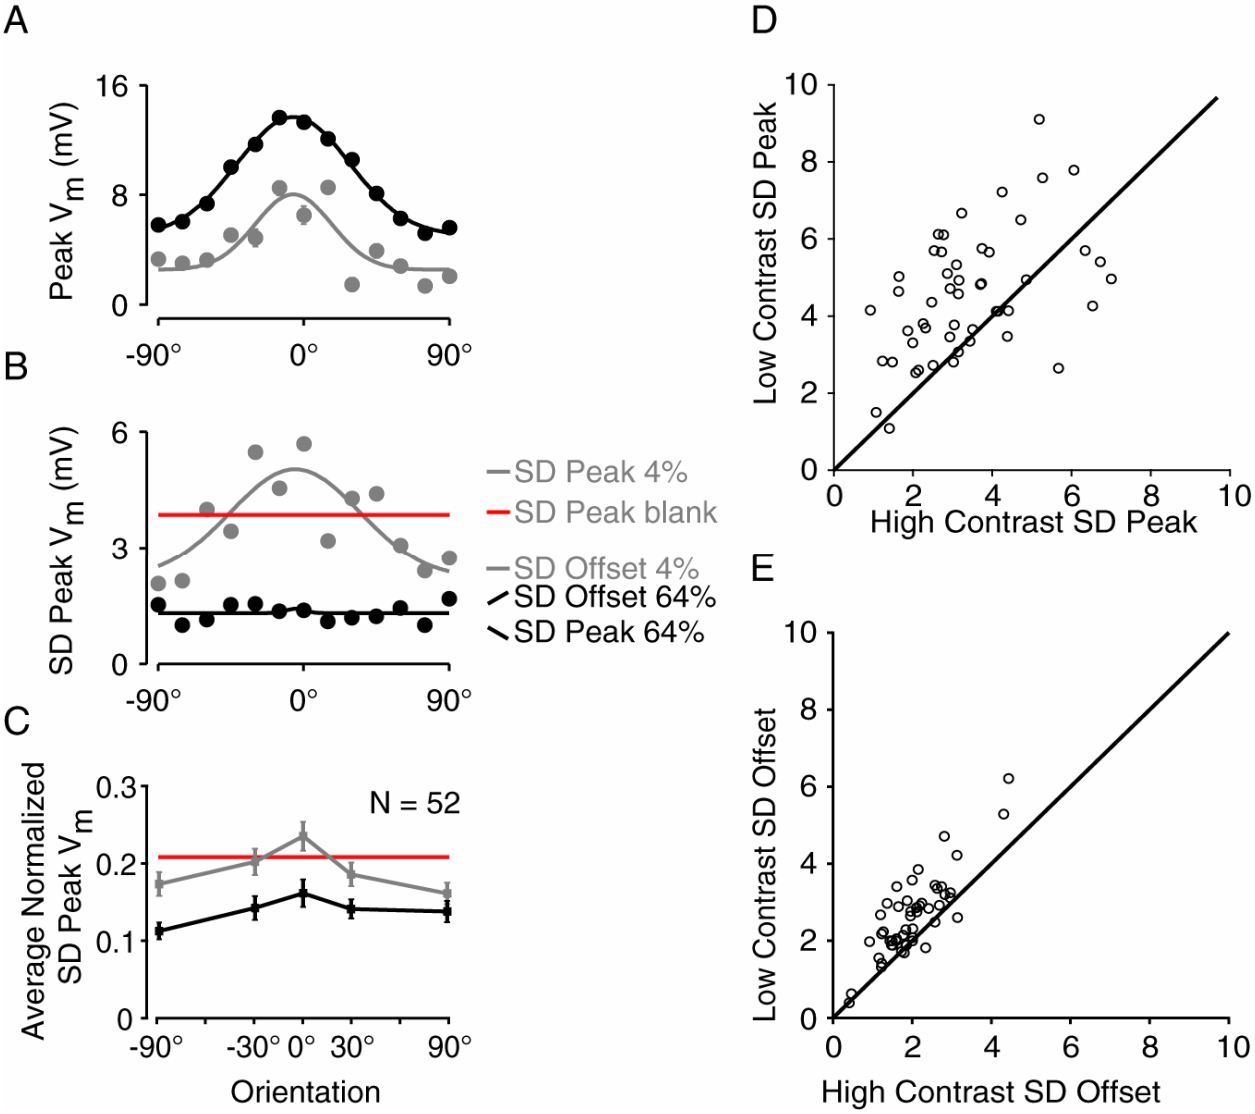
A
Peak Vm (mV)
16
8
0
-90°
0°
90°
B
SD Peak Vm (mV)
6
3
0
-90°
0°
90°
SD Peak 4%
SD Peak blank
SD Offset 4%
SD Offset 64%
SD Peak 64%
C
Average Normalized SD Peak Vm
0.3
0.2
0.1
0
N = 52
-90°
-30°
0°
30°
90°
Orientation
D
Low Contrast SD Peak
High Contrast SD Peak
E
Low Contrast SD Offset
High Contrast SD Offset

Figure 2.09. **Membrane potential variability and DC_N/DC_P**.

Top. Orientation tuning curves for the standard deviation of peak membrane potential (on a cycle by cycle basis) at low (4%) and high (64%) contrast, averaged across 14 cells (far left). Middle and right panels show the average of 5 cells each, all with small DC_N/DC_P ratios (< 0.3) or large DC_N/DC_P ratios (> 0.75). All tuning curves have been normalized to the standard deviation at the preferred orientation (0 deg). Blank mean (dashed line) and standard deviation (cyan) are also depicted for each population.

Bottom. Average membrane potential tuning curves for the same cells depicted above, again normalized to the preferred orientation.

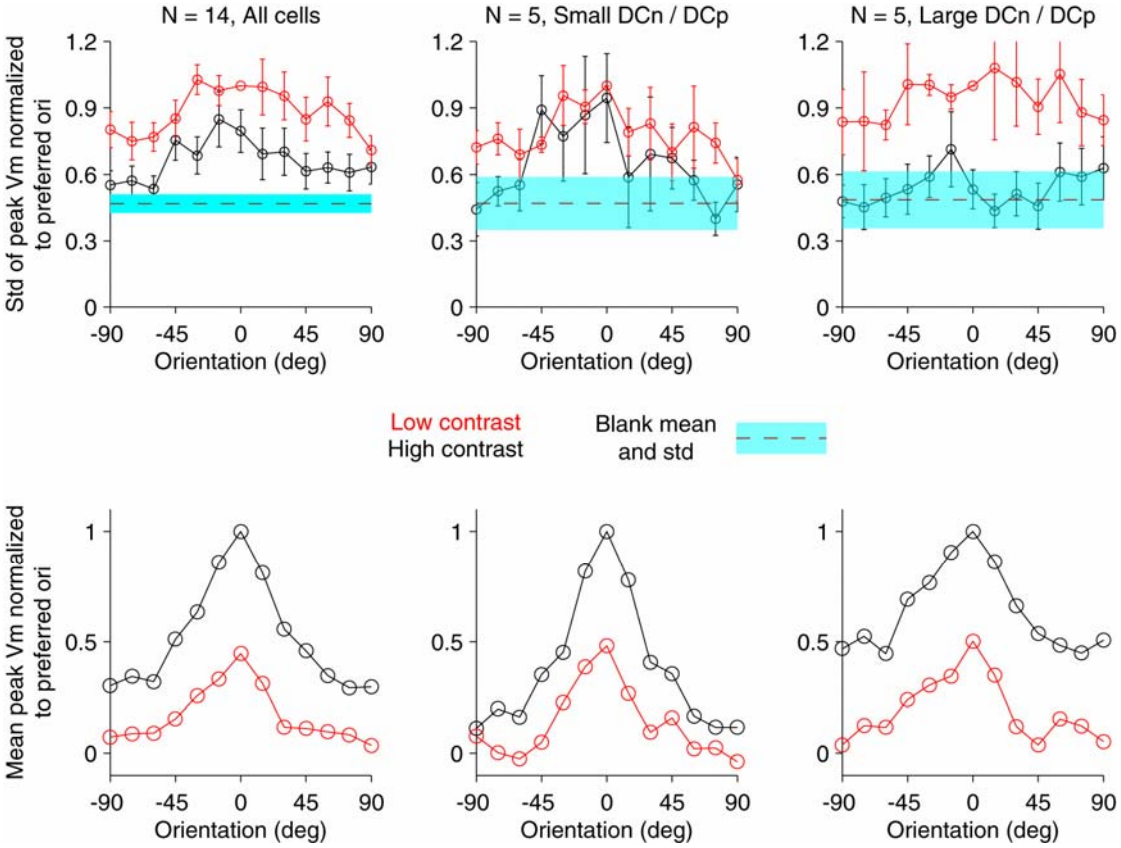
N = 14, All cells
N = 5, Small DCn / DCp
N = 5, Large DCn / DCp
Std of peak Vm normalized to preferred ori
Mean peak Vm normalized to preferred ori
Orientation (deg)
Low contrast
High contrast
Blank mean and std

The effect of trial-to-trial variability on firing rate

We have proposed that lower contrast leads to larger trial-to-trial variability and that larger variability in turn leads to higher spikes rate. The data in Fig. 2.07 show that trial-to-trial variability and the membrane-potential-to-spike-rate transformation both depend on contrast, but not that the spike rate transformation depends directly on variability. In order to make this connection, we plotted spike rate against mean membrane potential for different levels of variability (Fig. 2.10A and B). For each of 36 stimuli (3 contrasts and 13 orientations), responses were divided into 8 epochs of 30-ms duration. Mean potential, standard deviation of mean potential across cycles and mean spike rate were calculated for each epoch. The resulting 312 data points (8 epochs, 36 stimuli) were then grouped into bins of 2.25 mV in mean membrane potential and 0.625 mV of standard deviation, and the corresponding spike rates were averaged. Figure 2.10A (which illustrates a different cell than in Fig. 2.07A and B) indicates that increases in membrane potential variability lead to increases in spike-rate, much the same way that increases in mean membrane potential lead to larger spike-rates. To capture this trend, we applied an extension of the power law (Eq. 1) in which an increase in trial-to-trial variability (standard deviation) was essentially equivalent to an increase in mean potential:

$$R(\overline{V_m}, std_{V_m}) = k_1 \left\lfloor \left(\overline{V_m} - V_{rest}\right) + k_2 * SD_{V_m} \right\rfloor_+^p \qquad (2)$$

where $\overline{V_m}$ is the mean membrane potential averaged across trials, SD_{Vm} is the trial-to-trial standard deviation of the membrane potential, and k_1 and k_2 are constants. The fit of Equation 2 for the cell in Figure 2.10A is shown by the smooth curves. In Figure 2.10B the data points and fitted curves are replotted against the effective membrane potential, $\left(\overline{V_m} - V_{rest}\right) + 0.68 * SD_{V_m}$.

Here, the fitted curves, by construction, superimpose on one another, and the transposed points lie clustered along the fit.

In Figure 2.10C, the data from Figure 2.10A and B are replotted as a color-map of spike rate against mean and standard deviation of membrane potential. Colored lines show the trajectory of the membrane potential response in mean and in standard deviation over the course of 4 different stimuli (high-contrast preferred, black; high-contrast null, green; low-contrast preferred, blue; blank, red). A second example cell (same cell as in Fig. 2.07D-G) is shown in Figure 2.10D. The relationship between mean potential, standard deviation and spike rate is shown for the whole population in Figure 2.10E. For each cell the mean and standard deviation of the membrane potential were normalized to the largest stimulus-evoked depolarization; spike rates were normalized to the largest stimulus-evoked spike rate. The color maps for all 39 cells were then averaged together. The average image shows that the effect of increasing either mean membrane potential or membrane potential standard deviation is to increase spike rate.

To evaluate the effectiveness of Equation 2 in capturing the transformation between membrane potential and spike rate, we fit data from each cell to the equation and then made predictions of peak spike rate from mean membrane potential and membrane potential standard deviation. On average, a change in standard deviation was just over half as effective at increasing spike rate as a similar change in mean (k_2 had a roughly Gaussian distribution, with mean and sigma of 0.64 and 0.29). The predictions for all stimuli and all cells were then compared with the actual recorded spikes rates (Figure 2.10F). That the points cluster along the line of slope 1 indicates that Equation 2 captures the membrane potential to spike transformation well.

Figure 2.10. **The relationship between membrane potential mean, standard deviation, and spike rate**.

A. Mean and standard deviation of membrane potential and mean spike rate were measured in 30-ms epochs taken from the responses to gratings of different orientations and contrast. Data were binned into 2.25-mV intervals of mean potential and 0.625-mV intervals of standard deviation (SD) and then mean spike rate was plotted against mean membrane potential for 8 different SD intervals as indicated by the color legend. Curves are a fit to Equation 2.

B. Same data as in A, with spike rate plotted against mean membrane potential plus 0.68 times SD.

C. Same data as in A and B plotted as a color-map of spike rate against mean and standard deviation of membrane potential. Colored lines indicate the trajectory of mean and SD of membrane potential evoked by 4 different stimuli over the course of one grating cycle (high-contrast preferred, black; high-contrast null, blue; low-contrast preferred, green; blank, red). The mean and standard deviation of the membrane potential were computed using a 30 ms sliding window.

D. Same as C for the cell from Fig. 2.07D-G.

E. Same as C and D averaged over 39 cells. V_m and SD-V_m are normalized for each cell to the amplitude of the largest membrane potential response.

F. For 39 cells, the spike responses to stimuli of all orientations at high and low contrasts were calculated from Equation 2 using the corresponding membrane potential responses. The predicted spike rates are plotted against measured spike rates for each stimulus.

G. Data from Fig. 2.06B (spike-rate responses to high-contrast null and low-contrast preferred stimuli plotted against one another) replotted with a magnified y-scale.

H. Same as G, except that the spike rates plotted are predicted from membrane potential using Equation 2.

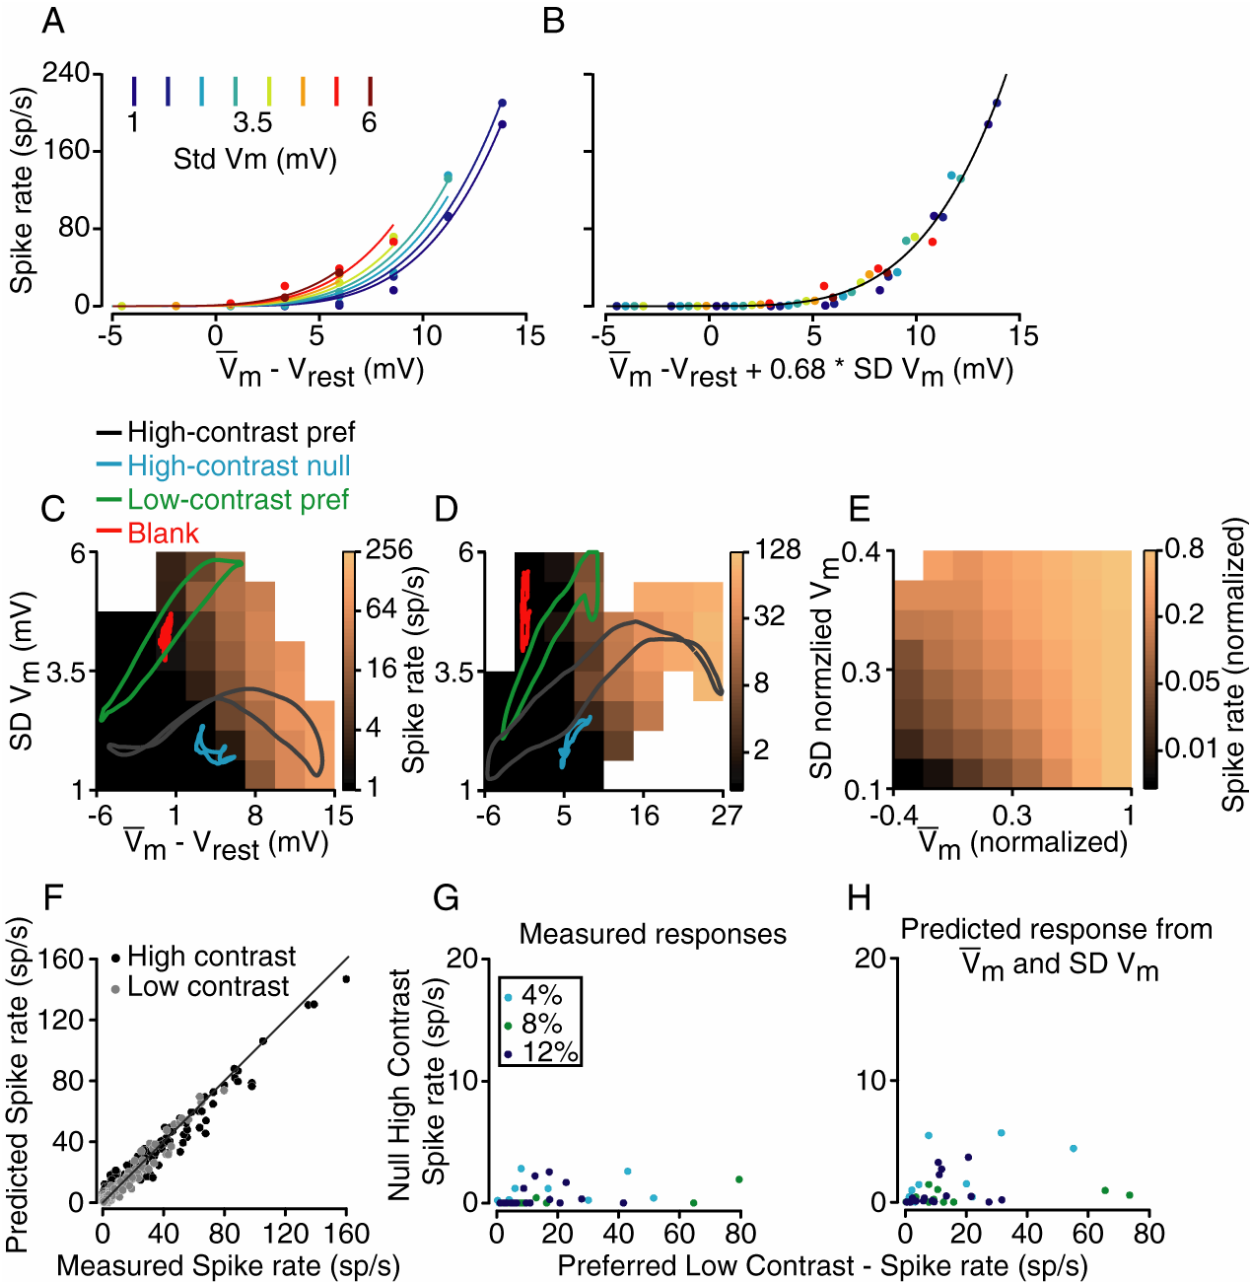
A
B
Spike rate (sp/s)
240
160
80
0
1
3.5
6
Std Vm (mV)
-5
0
5
10
15
$\overline{V}_m$ - V_{rest} (mV)
$\overline{V}_m$ -V_{rest} + 0.68 * SD V_m (mV)
High-contrast pref
High-contrast null
Low-contrast pref
Blank
C
D
E
SD V_m (mV)
6
3.5
1
-6
1
8
15
256
64
16
4
1
Spike rate (sp/s)
128
32
8
2
-6
5
16
27
SD normzlied V_m
0.4
0.3
0.1
-0.4
0.3
1
$\overline{V}_m$ (normalized)
0.8
0.2
0.05
0.01
Spike rate (normalized)
F
G
H
Predicted Spike rate (sp/s)
160
120
80
40
0
High contrast
Low contrast
0
40
80
120
160
Measured Spike rate (sp/s)
Measured responses
Null High Contrast Spike rate (sp/s)
20
10
0
4%
8%
12%
0
20
40
60
80
Preferred Low Contrast - Spike rate (sp/s)
Predicted response from $\overline{V}_m$ and SD V_m
20
10
0
0
20
40
60
80

Accounting for contrast invariance in spike-rate responses of simple cells

The amended power law accounts for why the spiking responses to high-contrast null stimuli are so much lower than the spiking responses to low-contrast preferred stimuli that evoke similar mean depolarizations (Fig. 2.06A and B). Figure 2.10G contains the same data as Figure 2.06B (spike response to high-contrast null stimuli plotted against response to low-contrast preferred stimuli) with an expanded y-axis; Figure 2.10H shows the predictions of spike rate derived from Equation 2. These predictions accurately reconstructed the differential amplification of spike rate responses to stimuli at high and low contrast.

All of the effects described so far – the mixing of cortical and geniculate excitatory input, early contrast saturation in membrane potential responses, and contrast-dependent trial-to-trial variability – should serve to make orientation tuning of simple cells' spike responses relatively invariant to changes in stimulus contrast. As discussed above, these effects are more important for cells that receive the majority of their excitatory input from the LGN, and thus respond to null oriented stimuli with a significant depolarization. The cell in Figure 2.11A-D, for example depolarized approximately 3 mV in response to low-contrast stimuli and 6 mV to high-contrast stimuli at the null orientation (Fig. 2.11A, -90° and 90°). This cell also exhibited early contrast saturation: By 4% contrast, the membrane potential responses were at least half the size of the 64% responses at all orientations.

The membrane-potential-to-spike-rate transformation for the cell was, as expected, highly non-linear: Even though the membrane potential response at the null orientation was almost 45% of the size of the preferred response (Fig. 2.11A and C), the null spike response was zero (B and D).

Finally, the cell showed contrast- and orientation-dependent (Monier et al., 2003) changes in the trial-to-trial standard deviation of the membrane potential (Fig. 2.11E). Spike-rate predictions derived from mean membrane potential and standard deviation of membrane potential produced orientation tuning curves (Fig. 2.11F) that were very similar to the ones derived from the cell's actual responses (Fig. 2.11D).

A second example cell with no depolarizing response to null-oriented stimuli is shown in Figure 2.11G-M). The cell did show considerable contrast saturation in that the 4% responses were over half as large as the 64% responses. Because of the lack of depolarization evoked by null stimuli, the orientation tuning of the membrane potential was largely contrast invariant, and by virtue of the power law (Equation 2), the spike responses were also invariant but with narrower tuning widths (half-width at half height: 26° for membrane potential vs. 13° for spike rate).

Figure 2.11. **Contrast invariance of orientation tuning in two simple cells**.

A. Cycle-averaged membrane potential responses to gratings of high and low contrast and different orientations.

B. Corresponding spike responses.

C-G. Orientation tuning curves at high and low contrast for mean membrane potential, spike rate, standard deviation of membrane potential and predicted spike rate (from Equation 2).

G-M. Same as A-F for a second cell.

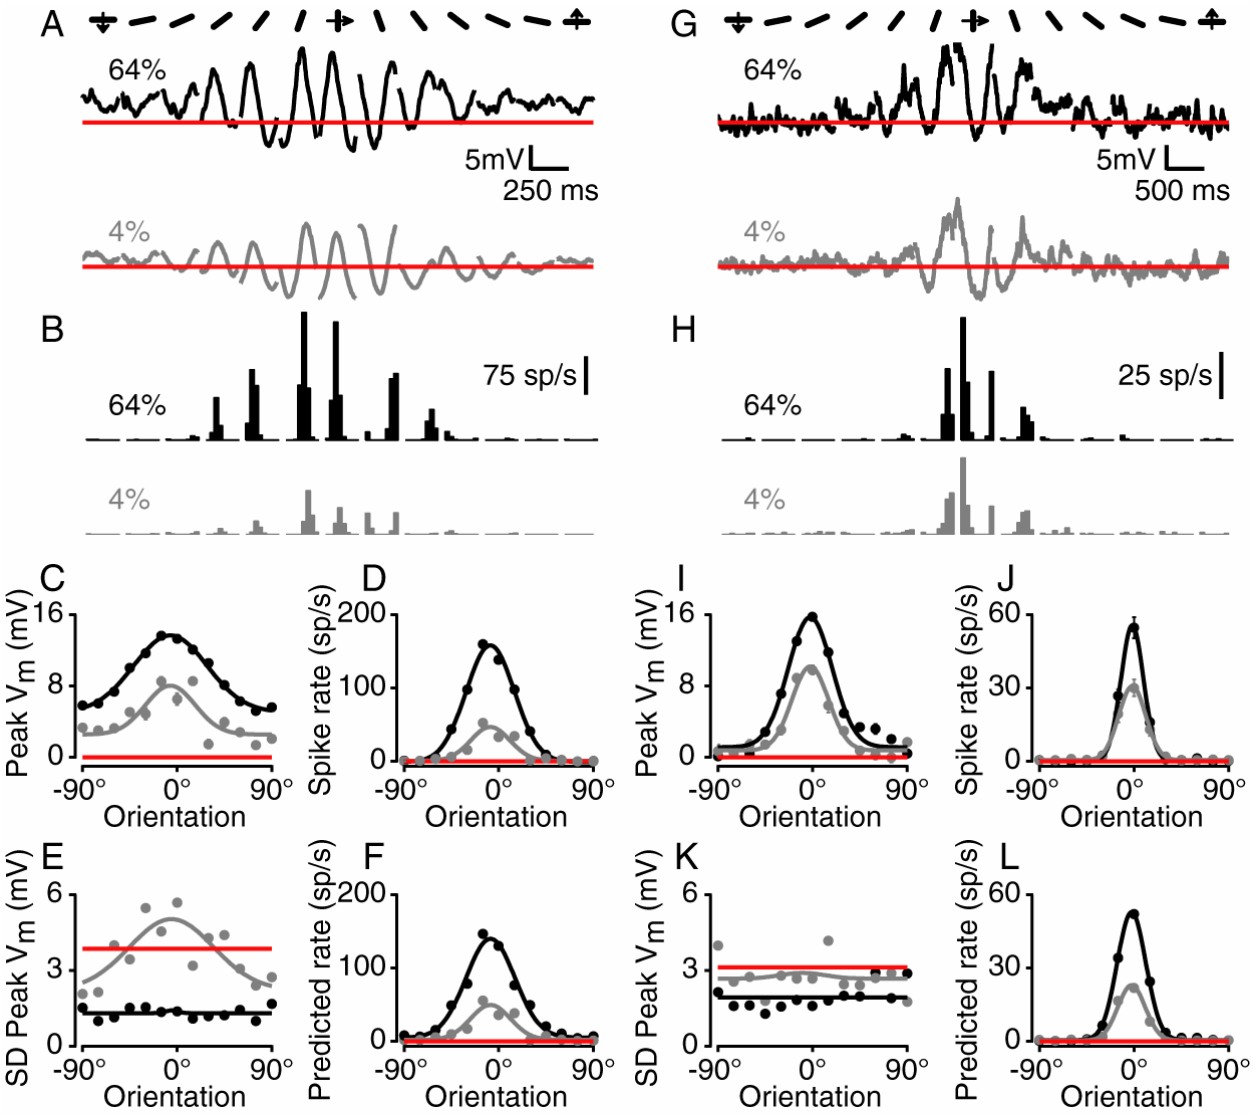
A
64%
5mV
250 ms
4%
B
64%
75 sp/s
4%
G
64%
5mV
500 ms
4%
H
64%
25 sp/s
4%
C
Peak Vm (mV)
16
8
0
-90°
0°
90°
Orientation
D
Spike rate (sp/s)
200
100
0
-90°
0°
90°
Orientation
I
Peak Vm (mV)
16
8
0
-90°
0°
90°
Orientation
J
Spike rate (sp/s)
60
30
0
-90°
0°
90°
Orientation
E
SD Peak Vm (mV)
6
3
0
-90°
0°
90°
Orientation
F
Predicted rate (sp/s)
200
100
0
-90°
0°
90°
Orientation
K
SD Peak Vm (mV)
6
3
0
-90°
0°
90°
Orientation
L
Predicted rate (sp/s)
60
30
0
-90°
0°
90°
Orientation

The relationship between tuning width and contrast is summarized for the population in Figure 2.12. Here we compare half width at half height (HWHH) of the orientation tuning curves, plotting the width at the lowest contrast tested against the width at 64% contrast (For a discussion of the relationship between HWHH, tuning and tuning curve offset, see Fig. 2.13). For the membrane potential responses, tuning width in many cells narrowed at low contrast relative to high contrast (Fig. 2.12A). The narrowing was predicted by the feed-forward model as a result of the untuned DC component of the geniculate input (Fig. 2.01B, right). The mean difference in HWHH between low and high contrast was 7.3^{o} (Fig. 2.12D). Tuning widths for the spike responses were far narrower than for membrane potential (median 32^{o} vs. 14.5^{o}; note change in scale between Fig. 2.12A and 2.12B-C), and much less dependent on contrast, with a 0.3° narrowing on average between high and low contrast (Fig. 2.12B and E). There were, however, a small number of cells that did not demonstrate contrast invariance in orientation tuning (Fig. 2.14).

Finally, we plot the widths of orientation tuning curves derived from predicted spike rates (Fig. 2.12C and F). Mean narrowing between predicted high and low-contrast widths (0.78°) was comparable to that seen in recorded spike rate; the distribution of contrast-dependent changes in tuning widths was broader, however, than that observed for the data (compare Fig. 2.12B and E). Not only were predicted and measured changes in tuning width similar over the population, they were also similar on a cell-by-cell basis, as shown in Figure 2.15.

Figure 2.12. **Contrast dependence of orientation tuning width**.

A-**C**. Half-width at half height (HWHH) of the orientation tuning curves at high and low contrast compared for mean membrane potential, measured spike rate, and spike rate predicted from Equation 2.

D-**F**. Histograms of low-contrast HWHH minus high-contrast HWHH.

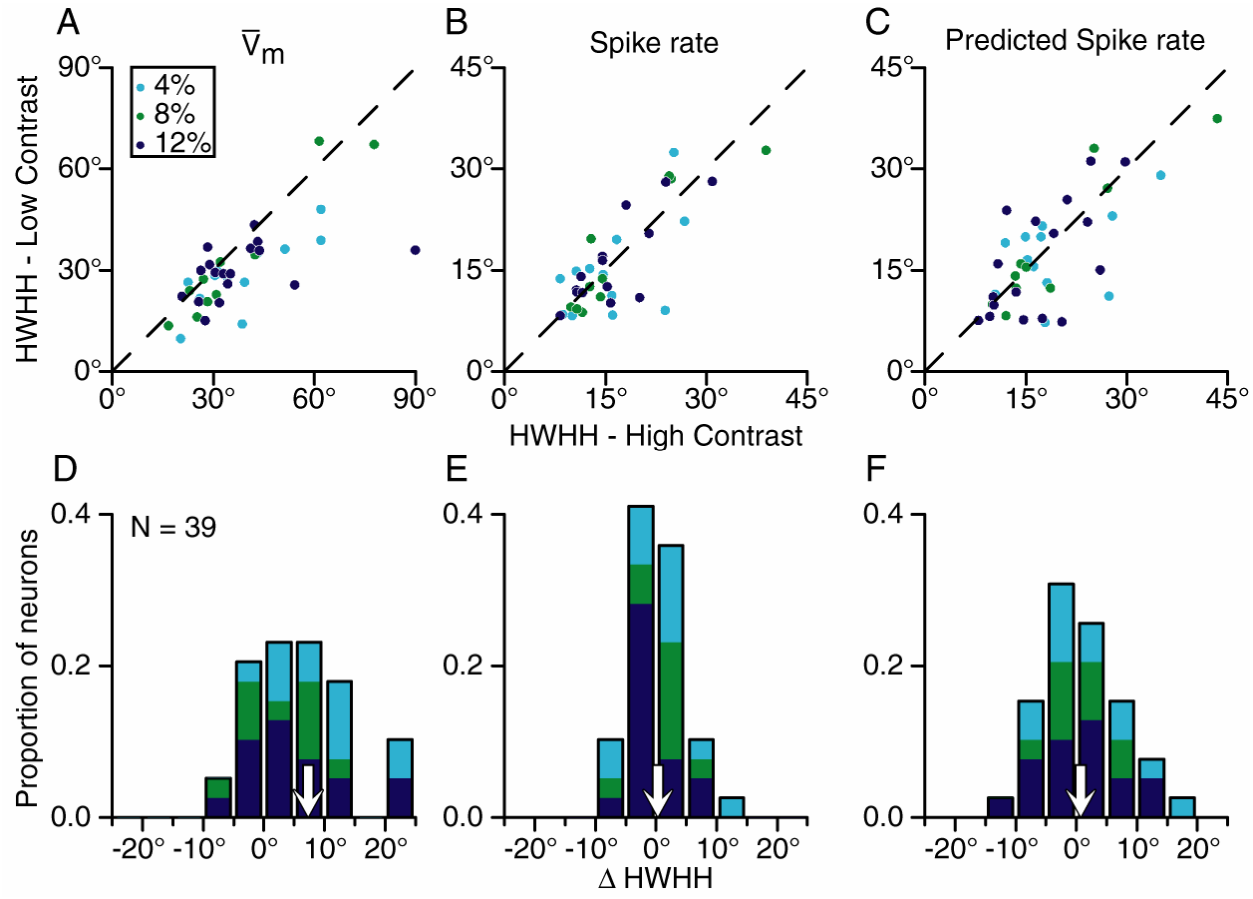
A
$\bar{V}_m$
4%
8%
12%
HWHH - Low Contrast
90°
60°
30°
0°
0°
30°
60°
90°
B
Spike rate
45°
30°
15°
0°
0°
15°
30°
45°
HWHH - High Contrast
C
Predicted Spike rate
45°
30°
15°
0°
0°
15°
30°
45°
D
N = 39
Proportion of neurons
0.4
0.2
0.0
-20° -10° 0° 10° 20°
E
0.4
0.2
0.0
-20° -10° 0° 10° 20°
Δ HWHH
F
0.4
0.2
0.0
-20° -10° 0° 10° 20°

Figure 2.13. **Relationship between half-width at height (HWHH), width and the offset of the tuning curve**.

A-C. Example tuning curves for high and low contrast (black and gray curves respectively). Each tuning curve has the same tuning width (□= 32 degrees, Equation 3), but the high contrast tuning curves have different background offsets (A:0%, B:20% and C:40%). The HWHH for high and low contrast are indicated in the top left corner of each panel. The ratio of the peak responses to null-oriented, high-contrast stimuli and to preferred low-contrast stimuli are indicated in the top right corner of each panel. As this ratio increases, the difference between low and high contrast HWHH also increases.

D. The contour map of HWHH for different combinations of tuning width and offset (□ and b, Equation 3).

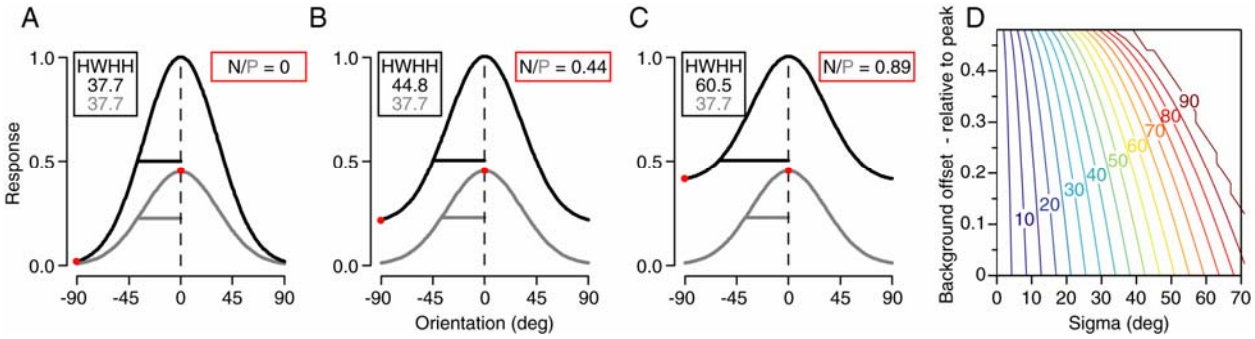
A
B
C
D
HWHH
37.7
37.7
N/P = 0
HWHH
44.8
37.7
N/P = 0.44
HWHH
60.5
37.7
N/P = 0.89
Response
Orientation (deg)
Background offset - relative to peak
Sigma (deg)
10
20
30
40
50
60
70
80
90

Figure 2.14. **Contrast dependence of orientation tuning measured from spike rate**.

A. Orientation tuning curves for 3 different contrasts in 4 cells (low contrast, 4%, 4%, 4% and 12%; medium contrast 8%, 8%, 16%, and 20%; high contrast 64%). Tuning width has an increasing dependence on contrast going from left to right. The data include the cells in Fig. 9B plus 26 cells that were recorded extracellularly.

B. Tuning width (half width at half height) for low contrast plotted against high contrast. Cells in which the tuning width at high and low contrast were statistically different are outlined in red. Statistical significance was taken to be non-overlapping 95% confidence intervals computed using the Gauss-Newton method.

C. A histogram of the ratio of tuning width at high and low contrasts. Red bars indicate statistical significance of the difference in tuning widths, as in B.

D. Tuning width as a function of contrast for the 4 example cells in A. Lines indicates a linear fit between orientation tuning width and contrast.

E. A histogram of the slopes of the linear fits between orientation tuning width and contrast for all cells. Bars in read indicate slopes that are significantly different from 0.

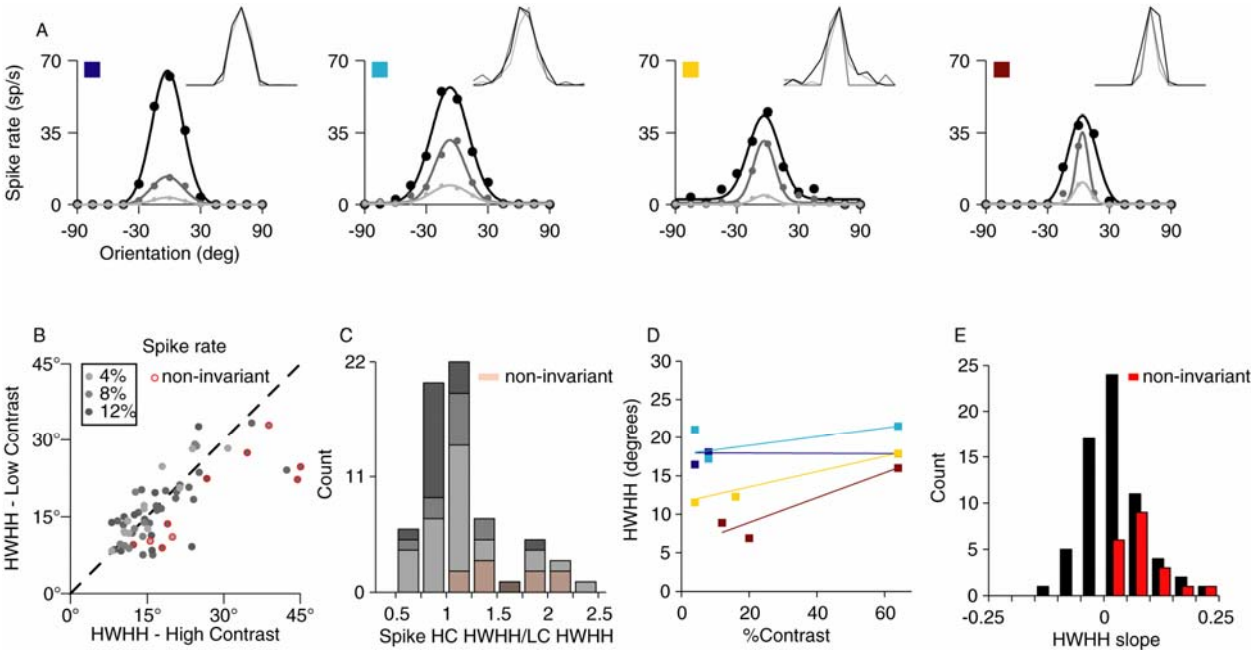
A
Spike rate (sp/s)
70
35
0
-90
-30
30
90
Orientation (deg)
B
Spike rate
4%
8%
12%
non-invariant
HWHH - Low Ccntrast
HWHH - High Contrast
45°
30°
15°
0°
C
non-invariant
Count
22
11
0
0.5
1
1.5
2
2.5
Spike HC HWHH/LC HWHH
D
HWHH (degrees)
30
25
20
15
10
5
0
20
40
60
%Contrast
E
non-invariant
Count
25
20
15
10
5
0
-0.25
0
0.25
HWHH slope

Figure 2.15. **Prediction of the difference in HWHH at low and high contrast compared to actual difference**. N=39.

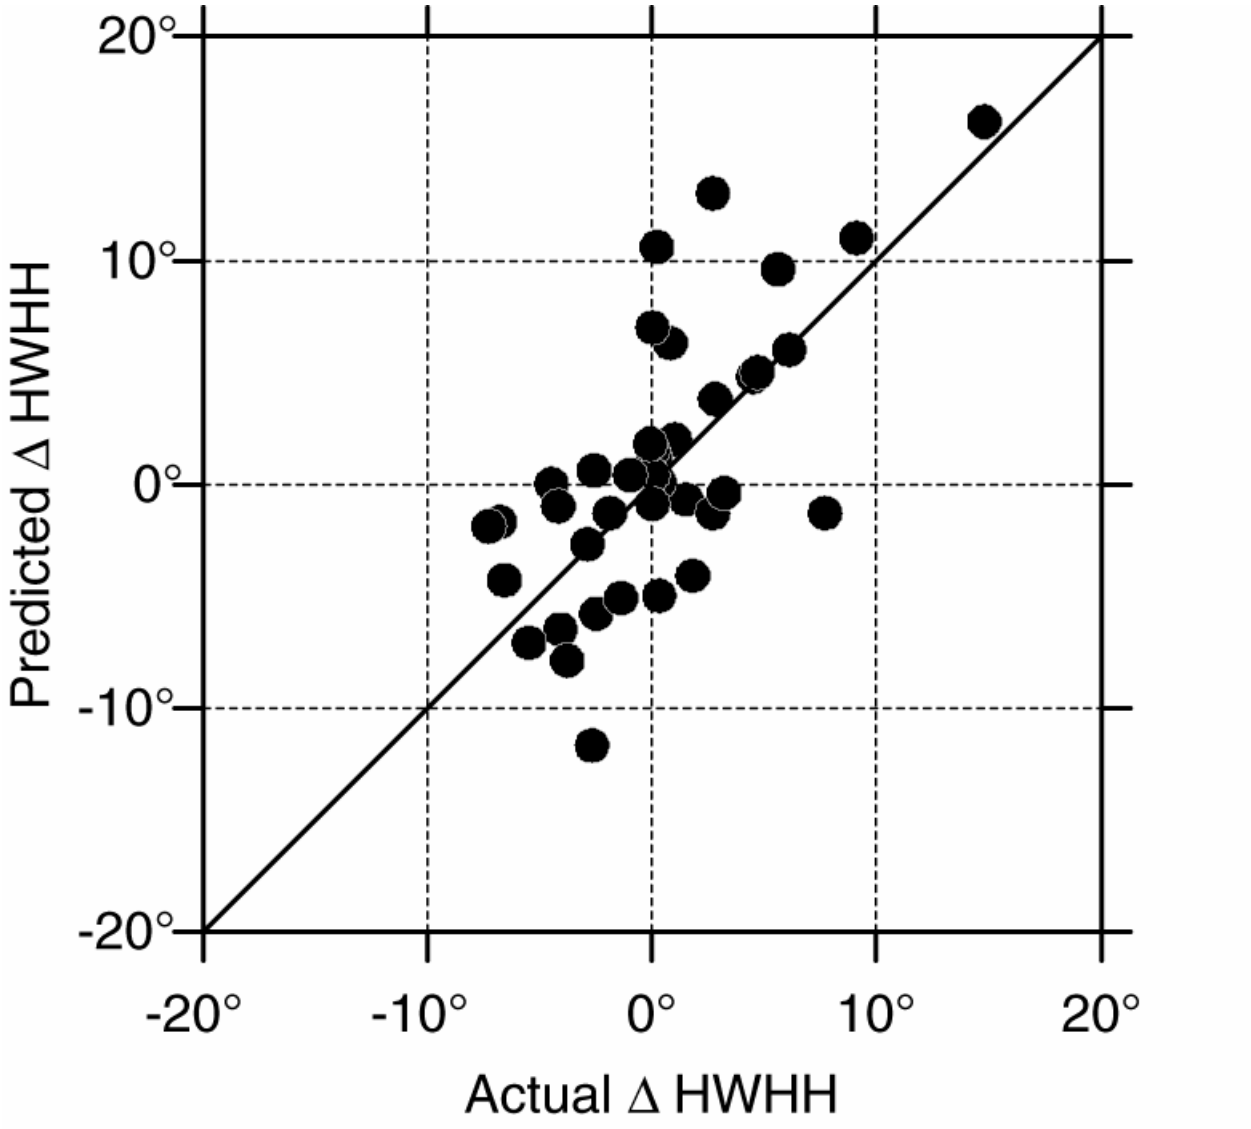
20°
10°
0°
-10°
-20°
-20°
-10°
0°
10°
20°
Predicted Δ HWHH
Actual Δ HWHH

Discussion

Two views of cortical computation have been proposed to account for the selectivity of sensory neurons. In one view, excitatory afferent input provides a rough sketch of the world, which is then refined and sharpened by lateral or feedback inhibition. In the alternative view, excitatory afferent input is sufficient, on its own, to account for sensory selectivity. We have studied which of these two viewpoints is most appropriate to describe one feature of cortical simple cells, namely, contrast invariant orientation tuning. A purely linear feed-forward model, incorporating only excitatory input from the LGN, predicts that the width of orientation tuning in simple cells broadens with contrast, breaking contrast invariance. Lateral inhibition, in the form of cross-orientation inhibition, is one mechanism that could restore contrast invariance by antagonizing feed-forward excitation at non-preferred orientations. We find instead that the predicted broadening is suppressed by three independent mechanisms, none of which requires inhibition. First, many simple cells receive only some of their excitatory input from geniculate relay cells (Fig. 2.03), with the remaining excitatory input originating from other cortical neurons with similar preferred orientations (Ferster et al., 1996; Chung and Ferster, 1998). Second, contrast-dependent changes in the trial-to-trial variability of responses lead to contrast-dependent changes in the transformation between membrane potential and spike rate. Third, membrane potential responses of simple cells saturate at lower contrasts than are predicted by a feed-forward model. We consider each of these mechanisms in turn.

Convergence of cortical and thalamic input

Because relay cells of the LGN are insensitive to orientation and because their firing rate responses are rectified, any feed-forward model based solely on input from relay cells predicts that drifting gratings of the non-preferred (orthogonal) orientation evoke a significant unmodulated rise in membrane potential, equal in size to the mean of the preferred response. Contrary to this prediction, few simple cells had been reported to depolarize or spike in response to orthogonal (null) stimuli, prompting the suggestion that null-evoked excitation from the LGN is suppressed by lateral inhibition in the orientation domain, i.e. cross-orientation inhibition or untuned inhibition (Sompolinsky and Shapley, 1997; Ferster and Miller, 2000). We find here that null-oriented stimuli do evoke a significant depolarization in many simple cells, the magnitude of which is directly proportional to the amount of excitatory input each cell receives from the LGN. In simple cells that receive the majority of their input from the LGN, the depolarization matches the prediction of the feed-forward model; in cells that receive only half of their excitation from the LGN and the rest from other cortical cells, the null-evoked depolarization is half as large as expected from the model. This match is inconsistent with the presence of cross-orientation inhibition: If cross-orientation inhibition were active in cortex, the points in Figure 3D would have consistently fallen below the unity line. In the extreme case of perfect null suppression, for example, all of the points would have fallen along the line Y=0. Our results do not exclude the possibility that simple cells receive untuned inhibition (Heeger, 1992; Troyer et al., 1998; Tao et al., 2004), which would decrease the DC component of the response at all orientations nearly equally. Our results are also consistent with previous evidence that inhibition in simple cells is often tuned to the preferred orientation (Ferster, 1986; Anderson et al., 2000b; Martinez et al.,

2005) (though see (Borg-Graham et al., 1998)). The push-pull organization of inhibition at the preferred orientation would have little effect on the peak amplitude of the response.

Contrast-dependent changes in trial-to-trial variability

Although the depolarization evoked by stimuli of the null orientation is on average smaller than what is predicted by the feed-forward model, it is often non-zero, and in many cells can be quite large (Figs 2 and 3). Why, then, do so few simple cells respond with spikes at the null orientation? This failure to respond seems especially paradoxical when low-contrast stimuli at the preferred orientation evoke comparable changes in membrane potential and yet do produce spike responses. The difference lies not in the average response (the membrane potential averaged over multiple trials), but in the trial-to-trial variability of the responses: Because the low-contrast preferred responses vary significantly from trial to trial, the membrane potential regularly rises above threshold. High contrast null responses, however, vary less from trial to trial and thus cross threshold less often. Note that the trial-to-trial variability for high-contrast, null-oriented stimuli can be even lower than the variability in the absence of a stimulus (Fig. 7C and D). Thus a reduction in variability could be the source of the null-evoked suppression of spontaneous firing that has been observed in some simple cells (Alitto and Usrey, 2004).

In trying to predict the mean spike rate that is evoked by a given stimulus, then, one must consider not only the mean depolarization evoked by the stimulus, but also the response variability (Azouz and Gray, 2000). A rise in either will tend to trigger spikes. Here, we have expressed this relationship as an extension of the power law relationship between mean potential

and mean spike rate (Hansel and van Vreeswijk, 2002; Miller and Troyer, 2002), where the mean potential is replaced by the sum of mean and membrane potential variability (Equation 2). We note, however, that this formulation was not derived analytically. It constitutes a convenient extension of the power law that fits the data in a simple way, but the exact form of the equation carries no theoretical significance.

The contrast dependence of trial-to-trial variability that we report here is somewhat larger than was reported previously (Anderson et al., 2000c). We attribute the difference between the two studies to a difference in the way trial-to-trial variability was measured. Anderson et al. averaged variability over the course of an entire grating cycle, whereas we analyzed variability at the peak of the response, where the cells most often fire action potentials. Because variability changes over the course of a cycle (see voltage trajectories in Fig. 7C and D), the two measures are not equivalent.

Saturation

In addition to contrast-dependent changes in trial-to-trial variability, a second aspect of cortical responses that contributes to contrast invariant orientation tuning is the early saturation of simple cell membrane potential responses relative to what is predicted by the feed-forward model. The spiking responses of geniculate relay cells are, on average, only half saturated at 16% or 20% contrast. In the majority of simple cells, however, membrane potential responses at these contrasts are at or near complete saturation. As a result, membrane potential responses at high (64%) and relatively low contrasts (~8%) are more similar to one another than expected. It is important to note, then, that previous studies of invariance were made using contrasts that are

generally higher than the half-saturation point (C_{50}) for simple cell membrane potential responses (Skottun et al., 1987; Anderson et al., 2000c; Alitto and Usrey, 2004). Only at contrasts below 8% does contrast invariance in spike-rate responses occasionally break down, in qualitative agreement with feed-forward predictions (Supplementary Figure 3).

There are several mechanisms that could account for early saturation in the membrane potential responses of simple cells. Thalamocortical and cortico-cortical depression (Abbott et al., 1997; Tsodyks and Markram, 1997; Kayser et al., 2001; Bannister et al., 2002; Carandini et al., 2002; Boudreau and Ferster, 2005) could disproportionately reduce the synaptic efficacy of synaptic input from higher contrast stimuli. So would the reductions in driving force on synaptic currents that occur as a result of strong depolarizations. If changes in depression or driving force at the thalamocortical synapse contributed to early saturation, we would expect the C_{50} of simple cell input to be similar at all orientations. If depression at corticocortical synapses contributed, then we would expect C_{50} to change with orientation, and possibly to vary with the %LGN input a cell received.

Conclusions

Contrast-invariance of orientation tuning is one instance of contrast-gain control (Heeger, 1992). Contrast-invariance requires the preservation of orientation tuning width for different contrasts; contrast-gain control requires the preservation of contrast response functions for different orientations. These requirements are equivalent: if a response to a stimulus of the non-preferred orientation saturated earlier or later in contrast than the response to the preferred orientation, the orientation tuning curves at different contrasts, when normalized, would not superimpose and

would not have the same widths. Our results show, then, that gain control is a property intrinsic to the feed-forward inputs to simple cells.

Gain control constitutes a general problem for sensory systems: How to distinguish changes in stimulus attributes from changes in stimulus strength. Both a drop in stimulus strength and a change in stimulus attribute away from the preferred will cause a reduction in the amplitude of a depolarizing response. Changes in trial-to-trial variability provide a mechanism by which to disambiguate these two events, such that a change in stimulus attribute causes a complete loss of spiking responses, whereas a drop in strength does not. In other words, changes in variability allow a cell to modulate the strength of its response without changing its selectivity. This mechanism requires only a spike threshold and stimulus-dependent changes in response variability, and so may be generally applicable to other parts of the visual system and to other sensory modalities. We note, however, that we have measured this effect by comparing different trials in a single cell. If the brain is to rely on such a mechanism in real time, there must be similar (and uncorrelated) variability among the responses of different cells within a given trial.

If stimulus-dependent changes in trial-to-trial variability really do contribute to preserving contrast invariance, where do they arise? One possible source is the changes in the excitability of the cortical circuit that underlie transitions between cortical UP and DOWN states (Anderson et al., 2000a; Shu et al., 2003; MacLean et al., 2005). Another possible source, at least at the null orientation, is shunting inhibition (Borg-Graham et al., 1998; Monier et al., 2003). We find no evidence, however, for either cortically-mediated inhibition or excitation at the null orientation, in that the amplitude of the null-evoked depolarization exactly matches the prediction of the

feed-forward model when the %LGN Input each cell receives is taken into account (Fig. 2.03). This result suggests that the trial-to-trial variability of responses evoked by the null orientation may also originate in the LGN. In support of this suggestion, Kara et al. observed significant contrast-dependent changes in trial-to-trial variability (Fano factor) in individual geniculate relay cells (Kara et al., 2000).

Since Hartline described it in the late 1940's (Hartline, 1949), lateral inhibition has been assumed to shape receptive field selectivity in many sensory domains. In visual cortex, however, much of the detailed behavior of simple cells can be captured by a simple feed-forward model lacking lateral inhibition. Many phenomena assumed to arise from visually selective intracortical inhibition can be accounted for by non-specific nonlinearities that occur at several different stages of processing. The sharpness of orientation tuning (Carandini and Ferster, 2000; Volgushev et al., 2000) and direction selectivity (Jagadeesh et al., 1997; Priebe and Ferster, 2005) arise from spike threshold and the so-called iceberg effect. Cross-orientation suppression can be explained by rectification and contrast saturation in geniculate relay cells (Priebe and Ferster, 2006). Here we show that one of the hallmarks of visual cortex, the contrast invariance of orientation tuning, may arise from contrast saturation, spike threshold, and trial-to-trial variability, all of which operate within the local cortical circuit and its thalamic inputs.

Chapter III

Computational diversity in complex cells of cat primary visual cortex

Ian M. Finn[1] and David Ferster[1]

[1]Northwestern University Institute for Neuroscience
Department of Neurobiology and Physiology

Northwestern University, Evanston IL 60208

Acknowledgements: We thank Tomaso Poggio, Ulf Knoblich and Nicholas Priebe for helpful discussions, and Ilan Lampl, Jose-Manuel Alonso and Maximilian Reisenhuber for comments on the manuscript. Supported by grants from the National Eye Institute (R01 EY04726) and National Institutes of Mental Health (P20 MH066239).

Abstract

A recent study has suggested that complex cells perform a MAX-like operation on their inputs. When two bar stimuli are presented within the receptive field, regardless of their relative separation, the cell's response is similar in amplitude to the larger of the responses elicited by the individual stimuli. This description of complex cells seems at odds with the classical energy model in which complex cells receive input from multiple simple cells with overlapping receptive fields. The energy model predicts – and experiments have confirmed – that bar stimuli should facilitate or suppress one another depending on their relative separation. We have recorded intracellularly from a population of complex cells and studied their responses to paired bar stimuli in detail. A wide range of behavior was observed, from the more classical separation-dependent interactions to purely MAX-like responses. We also found that the more MAX-like a cell was, the broader its spatial frequency tuning as measured with drifting gratings. These observations are consistent with energy models in which classical complex cells receive input from simple cells with similar preferred spatial frequencies, and MAX-like complex cells receive input from simple cells with disparate preferred spatial frequencies. Generalized energy models, then, can account for diverse modes of computation in cortical complex cells.

Introduction

Simple cells in primary visual cortex are well described by feed-forward models in which their basic response properties are derived from the lateral geniculate nucleus (LGN) (Hubel and Wiesel, 1962). The nature of the circuitry that gives rise to cortical complex cells is, by comparison, much less clear. Similar to simple cells, complex cells are selective for orientation and spatial frequency (Hubel and Wiesel, 1962; Movshon et al., 1978b). Unlike simple cells, complex cells lack obvious substructure in their receptive fields (Hubel and Wiesel, 1962; Movshon et al., 1978b; Szulborski and Palmer, 1990), and as a group appear to be more heterogeneous.

How can we best account for the aspects of complex cell responses that are shared with simple cells as well as those that are disparate? Hubel and Wiesel (1962) proposed that complex cell tuning is inherited from simple cell progenitors. In support of their hierarchical model, extracellular experiments designed to detect second-order structure in complex cell receptive fields have revealed simple-cell like patterns (Movshon et al., 1978b; Emerson et al., 1987; Szulborski and Palmer, 1990; Livingstone and Conway, 2003). In particular it has been reported that the spike rate response to an oriented bar flashed in the center of a complex cell's receptive field was modulated by the presence of a second simultaneously presented bar in a manner that depended on the separation between the two bars (Movshon et al., 1978b). The dependence of this interaction effect on bar separation was reminiscent of the subfield structure of simple cell receptive fields, and it predicted for individual complex cells the shape of their spatial frequency tuning curves.

In a recent model it was proposed that complex cells might perform a very different computation on their inputs, one resembling a MAX-like operation (Riesenhuber and Poggio, 1999, 2002; Serre et al., 2007). That is, when presented with pairs of stimuli, the response of a complex cell would resemble the larger of the responses to the two stimuli alone. MAX-like behavior has been observed in extracellular recordings from primate areas V4 (Gawne and Martin, 2002) and IT (Sato, 1989) and in intracellular recordings from complex cells in cat area V1 (Lampl et al., 2004).

The MAX-like computation reported by Lampl et al. (2004) is distinct from that measured in previous experiments (Movshon et al., 1978b; Emerson et al., 1987; Szulborski and Palmer, 1990; Livingstone and Conway, 2003), and is not predicted by the standard hierarchical model of cortical processing. How, then, can the rather different reports of complex cell behavior be reconciled? We have found in a detailed intracellular study of complex cells that both types of response patterns exist – in some complex cells, the interactions between stimuli in a pair clearly depended on the separation between stimuli and their polarity; in others, stimuli interacted in a MAX-like manner, independent of separation or polarity. The two types of cells lay at the ends of a continuum: Quantitative indices of MAX-like behavior showed a unimodal distribution, and were inversely correlated with the spatial frequency tuning bandwidth of the cells.

A hierarchical energy model (Adelson and Bergen, 1985) constructed using one or two pairs of simple cells qualitatively reproduced the complex-cell behavior observed here. Classical responses emerged when a single pair of matched spatial frequency simple cells was employed. More MAX-like behavior was observed when two pairs of cells with different spatial frequency

selectivities provided input to the model complex cell. Thus it may be the case that complex cells participate in a variety of image processing computations dependent in part on the spatial frequency preferences of the inputs they receive.

Experimental Procedures

Animal preparation: Adult female cats weighing between 2 and 3 kg were anesthetized with a ketamine/acepromazine mixture (30 mg/kg ketamine, 0.7 mg/kg acepromazine, i.m.). Cannuli were inserted into the femoral veins and anesthesia was subsequently maintained with intravenous infusion of sodium pentathol (1-2 mg/kg/hr). A trachea tube for artificial respiration and a vertebral clamp for suspension of the thorax were surgically inserted, and the animal was placed in a sterotaxic headholder. Small caliber holes were drilled in the cranium for placement of screws used to monitor the electroencephalogram (EEG), and a craniotomy measuring between 2-4 mm in width and 4-7 mm in length was made at Horsley-Clark coordinates centered 2 mm laterally and 6 mm posteriorly. After a suitable period during which the animal's vital signs remained stable, paralysis was induced with the perfusion of 3-4 ml of vecuronium bromide (1.5 mg/kg), following which the animal was artificially respirated at 30 breaths / minute. Heart rate and expired CO_2 were monitored, with the later adjusted periodically by the alteration of an administered room air / O_2 mixture to keep end tidal CO_2 in the range of 3.5-4.2%. Continuous perfusion of paralytic was administered at 0.2 mg/kg/hr for the duration of the experiment; anesthetic (pentothal) was set to perfuse automatically at a rate between 1-2 mg/kg/hr to maintain the animal in stage II sleep. Body temperature was monitored and maintained close to 38.3^{o} C with a feedback controlled heat lamp. Gas permeable hard contact lenses filled with a saline/atropine mixture, to effect pupil dilation, were inserted in both eyes after the nictating membranes were retracted by the application of 2-3 drops of 2% phenylephrine hydrochloride. Corrective external lenses were then placed in front of the eyes to focus the display screen onto the retinas. Focus was determined by imaging the retinas onto the

display screen with a fiber optic light source directed into the eyes. To minimize brain movement related to respiration, bilateral pneumothoracotomies were performed to limit respiration-induced changes in intrathoracic pressure. Prior to recording, a durotomy was made, typically with an area between 2-5mm^2, over which a layer of warm agar (3% in 0.9% saline) was applied to protect the cortex during and between electrode penetrations. All methods related to animal treatment during experiments have been approved by Northwestern University's Committee on Experimental Animal Research.

Stimulation and recording: Recordings were made with whole-cell patch microelectrodes pulled on a Flaming/Brown micropipette puller (Sutter Instruments model p87) from 1.2mm thin-wall borosilicate glass filaments. The electrodes, with resistances ranging between 7 and 11 MΩ, were filled with an internal solution, consisting of (in mM): 130 K^+-gluconate, 2 $MgCl_2$, 5 HEPES 1.1, EGTA, 0.1 $CaCl_2$, and 4 Mg^{2+}-ATP, which was subsequently buffered to pH 7.3 and adjusted (via dilution with ddH2O) to 285 mOsm. Membrane potentials were recorded with an Axoclamp-2A amplifier in current clamp mode and digitized at between 4 and 10 kHz. Spikes were identified by subtracting a low pass filtered version of the membrane potential from the original voltage trace. For voltage analyses spikes were removed by median filtering. Stimuli were generated on a Macintosh computer (Apple, Cupertino, CA) running Matlab (Mathworks, Natick, MA) with the Psychophysics toolbox libraries, which controls the output of a Viewsonic (Walnut, CA) video monitor (mean luminance 20cd/m^2) placed 50cm in front of the animal. All neurons were recorded from areas 17 or 18 and had receptive fields with eccentricities less than 10^o.

Receptive field characterization and cell classification: Receptive fields were initially identified by hand. Once localized, the cell's ocular dominance was determined and the non-dominant eye was blocked. Orientation preference was assessed with a protocol that pseudorandomly interleaves twelve four-second presentations, each preceded by 250ms of blank stimulation, of a drifting grating at 12 different orientations between 0 and 330^{o}. The preferred orientation was defined as that which produced the largest change in mean potential (F0) at the grating's temporal frequency. In order to position the visual stimulus more precisely over the receptive field center, a one dimensional map was made employing two sets of between eight and twelve bars of high, typically 90%, contrast: Bright bars ('ON') were presented at 38 cd/m^2 and dark bars ('OFF') were presented at 2 cd/m^2. The bars were between 0.2 and 0.6 degrees wide, and were presented individually and pseudorandomly for 60ms, with a subsequent 240ms of blank stimulation in order to allow the cell to return to near the resting potential between stimuli. Both ON and OFF bar maps were constructed, and those bars that produced a response different from the mean at 50-80ms latency were considered to lie within the receptive field. This protocol was repeated a number of times (between 5 and 25). Stimulation boundaries were set to encompass the extremes of each map. Overlap between maps made with ON and OFF stimuli, as well as the F1/F0 ratio derived from responses to preferred orientation drifting gratings, were used to classify cells as complex.

Spatial frequency measurements: Spatial frequency selectivity was measured by presenting eleven (including a blank trial) pseudorandomly interleaved drifting gratings of different spatial frequencies, all at the preferred orientation and temporal frequency and at 64% contrast. Tuning

curves were constructed from the average (DC) response to each spatial frequency. Preferred spatial frequency and bandwidth were obtained from fits to the equation:

$$V(sf) = V_{rest} + A * \exp(-(sf / lpF)^{lpE}) * (1/\sqrt{1 + (hpF / sf)^2})^{hpL}.$$

where V_{rest} is the resting membrane potential of the cell, A is amplitude, sf is spatial frequency, lpF and lpE are variables capturing the low pass behavior of the curve, and hpF and hpL are variables capturing the high pass behavior of the curve. The spatial frequency that elicited a maximal response was considered to be the preferred spatial frequency. Bandwidth was determined by extracting the high and low spatial frequency cutoffs (where the response dropped to one-half maximal) and taking the ratio of the two (Sceniak et al., 2002). The preferred spatial frequency determined in large part the number of stimuli used for paired-bar mapping – the bars were chosen such that 0.5/bar width, which is the highest frequency that can be resolved, was larger than the cell's peak spatial frequency.

Paired-bar measurements: All paired-bar experiments contained presentations of (1) single bars presented at each position with both ON and OFF polarity (±45% contrast), (2) every possible pair of ON and OFF polarity bars at ±45% contrast, (3) single bars of ±90% contrast at each position, representing the superposition of bars of ±45% contrast, and (4) blank stimuli.

Results

Complex cell responses to paired bar stimuli

Whole-cell intracellular recordings were obtained from complex cells in anesthetized cats. We first determined the cell's orientation and spatial frequency preferences using drifting gratings. The cell's receptive field center and extent were then mapped by measuring subthreshold responses to narrow, optimally oriented, single bright (ON) and dark (OFF) bar stimuli (±45% Weber contrast relative to background). For paired bar mapping, the receptive field, together with a small portion of the surrounding region, was divided into between 8 and 12 oriented bars, with the bar width chosen to be as narrow or narrower than ½ the period of the cell's preferred spatial frequency. Bars were presented for 20-60ms followed by blank periods of 160-240ms. We measured responses to a complete set of second order (paired) stimuli by flashing every combination of bar position and polarity. When the members of a pair with the same polarity were both located in the same position, the contrasts added to create a single bar with doubled (±90%) contrast.

The stimulation method and example responses recorded from a complex cell are shown in Figure 3.01. Here the stimuli were flashed for 60ms, followed by 240ms of a blank screen (Fig. 3.01A). All bar pairs were displayed once in random order during a stimulus trial. A subset of one trial is shown for an example complex cell in Figure 3.01B with the cell's responses above the stimuli that evoked them; note that this block of 10 stimuli contained a blank stimulus (at 2.4 seconds) in amongst the pairs and singleton bars. We generally presented the full stimulus set

between 10 and 30 times (each time with a different random order), and then averaged the responses to each stimulus after eliminating spikes with a median filter.

Figure 3.01C shows average responses for the cell in Figure 3.01B to high-contrast single bright bars (gray traces) and dark bars (black traces). At most spatial locations, the peak responses of the cell were very similar in amplitude and occurred at nearly the same latency, identifying the cell as being complex. To the right are space-time maps of the responses, again showing the receptive field similarity when probed with either ON or OFF polarity bars.

Figure 3.01. **The paired-bar stimulus protocol and example complex cell responses**.

A. Example stimulus frames for the paired-bar protocol. All possible pairs of ON and OFF bars were flashed at 45%-contrast along with individual bars of both polarities at 45% and 90% contrast.

B. Complex cell responses to part of a paired-bar protocol. Bars were flashed for 60ms with 240ms of succeeding blank time. This cell depolarized in response to almost all combinations of bars regardless of polarity.

C. Responses to ON and OFF bars at 90% contrast for the complex cell in B. The similarity between ON and OFF responses at each position clearly marked this cell as being complex.

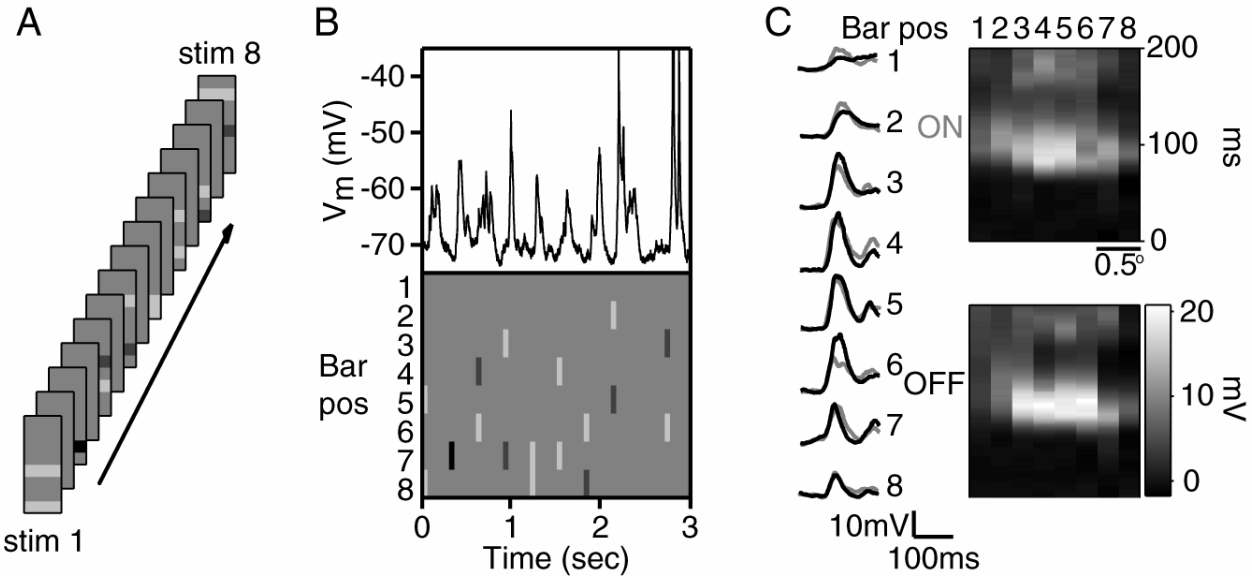
A
stim 8
stim 1
B
-40
-50
-60
-70
Vm (mV)
Bar pos
1
2
3
4
5
6
7
8
0
1
2
3
Time (sec)
C
Bar pos 1 2 3 4 5 6 7 8
ON
OFF
200
100
0
ms
0.5°
20
10
0
mV
10mV
100ms

A matrix representation of all the ON-ON bar pairs with which the cell was stimulated is shown in Figure 3.02A. Individual bar stimuli are arrayed along the top row and leftmost column. Across any given row *a,* one bar stimulus remains constant and is paired with its appropriate partner from column *b*. Thus, along the major diagonal (when $a = b$, marked "0" in the figure), the bars superimpose to produce a single bar stimulus that is doubled in contrast. In this representation, paired-bar stimuli with the same separation between bars fall along diagonals; looking along the diagonal marked by 2, for example, the center-to-center distance between paired bars is exactly two bar widths. Altogether, four different matrices could be constructed from our stimulus set, one for each combination of bar polarities – ON-ON, OFF-OFF, ON-OFF, and OFF-ON.

Average responses to each stimulus in Figure 3.02A are illustrated in Figure 3.02B for the complex cell from Figure 3.01. The responses are plotted relative to, and slightly elevated from, the resting membrane potential of the cell (horizontal lines). At all positions (a,b) where $a \neq b$, we placed the response to bar *a* alone (red), the response to bar *b* alone (red), the response to bars *a* and *b* together (blue), and the linear prediction generated by summing the responses to *a* and *b* alone (green). For major diagonal entries ($a = b$; bar separation = 0) in ON-ON and OFF-OFF matrices, we placed the response to *a* (red), the response to *a* at double (±90% contrast) contrast (blue; same as ON responses in Figure 3.01C), and twice the response to *a* (green). Above the example ON-ON matrix are four entries (gray box) that have been magnified so that the individual traces can more easily be identified. The coincidence of an ON bar and an OFF bar at the same spatial location is equivalent to showing no stimulus at all, and so the major diagonals in ON-OFF and OFF-ON matrices were not filled in.

Figure 3.02. **Stimulus organization**.

A. Responses to all stimuli were arranged in matrices as depicted. Stimuli falling along the arrows contained bars that were equidistant from each other. The major diagonal stimuli (bar separation=0) were at 90% contrast, mimicking the coincidence of two 45%-contrast stimuli at the same spatial location; all other stimuli were at 45%-contrast. Four variations of this stimulus arrangement were created for each cell: ON-ON (shown), ON-OFF, OFF-OFF, and OFF-ON. **B**. The full set of ON-ON responses for the complex cell from Figure 3.01. Responses to each bar presented individually at 45% contrast (R_a and R_b) are in red, the response to both bars together (R_{a+b}) is in blue, and the predicted response to both bars together ($R_a + R_b$) is in green. Along the main diagonal the blue responses are identical to the gray traces in Figure 3.01C. The gray box contains traces that have been magnified for easier viewing (arrow).

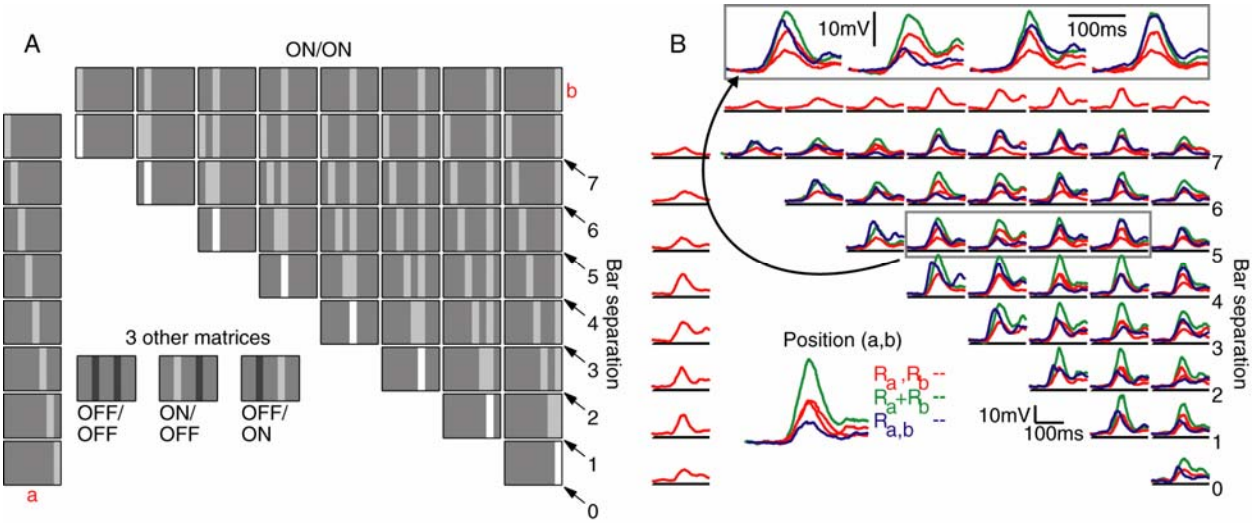
A
ON/ON
b
a
7
6
5
4
3
2
1
0
Bar separation
3 other matrices
OFF/
OFF
ON/
OFF
OFF/
ON
B
10mV
100ms
Position (a,b)
R_a, R_b --
R_a+R_b --
$R_{a,b}$ --
10mV
100ms
Bar separation

MAX-like and Classical complex cells

We recorded paired bar responses from 45 complex cells, two of which are shown in Figure 3.03. For cell M1 (Fig. 3.03A), we show the ON-ON matrix (left) and the ON-OFF matrix (right). For cell M2 (Fig. 3.03B), we show the OFF-OFF matrix (left) and OFF-ON matrix (right). The remaining two matrices for M1 and M2 can be found in Figures 3.05 and 3.06. The format of Figure 3.03 is identical to that of Figure 3.02B, except that the matrix entries with bar separation 0 are not shown. Here again the gray boxes depict entries that have been magnified for easier viewing. For M1, 8 bars of 0.31° width were flashed in pairs, with each stimulus on for 60ms and followed by 240ms of a blank screen. For M2, 8 bars of 0.39° width were flashed in pairs with each stimulus on for 40ms and followed by 180ms of a blank screen.

These two cells appeared to behave in a very MAX-like manner: Their responses to almost all bar pairs (blue traces) were consistently similar to the larger of the individual responses (red traces), and invariably smaller than the sum of the individual responses (green). The two cells in Figure 3.04, on the other hand, showed very different behavior, resembling more the classical complex cells described by Movshon et al. (1978b) in which the interactions between bars in a pair depended systematically on bar separation and polarity. In cell C1, for example, bright bars separated by a distance of 2 bar widths showed consistent suppression across the entire receptive field: For the violet-shaded traces in Figure 3.04A (left), the paired response (blue) was always smaller than the larger of the individual (red) responses. When the polarity of one of the bars was reversed, however, the suppression turned into facilitation: For the gray shaded traces in Figure 3.04A (right), the paired response to the bars (blue) was always greater than the larger of the individual responses (red). At further separations between bars (5 bar widths), the interaction

reversed, with summation seen between two bright bars (Fig. 3.04A, left, gray shading). Note that there is no obvious corresponding suppression for the larger bar separation visible in Figure 3.04A (right). This was a consistent finding in that we did not often observe suppression between bars of opposite polarity.

Similar separation- and polarity-dependent changes in bar interactions are shown for a second cell (C2) in Figure 3.04B; note that bar separations 0 and 1 are, however, not shown. The remaining two matrices for the cells in Figure 3.04 can be found in Supplementary Figures 7 and 8.

Figure 3.03. **MAX-like responses in complex cells**.

A. A complex cell ON-ON matrix and ON-OFF matrix filled as described in Figure 3.02. All of the paired responses are clearly sublinear, and the example magnified traces show that the paired responses are very similar to the maximum of the individual responses, indicating MAX-like behavior.

B. OFF-OFF and OFF-ON matrices for a second complex cell, both demonstrating mostly MAX-like responses independent of bar distance or polarity.

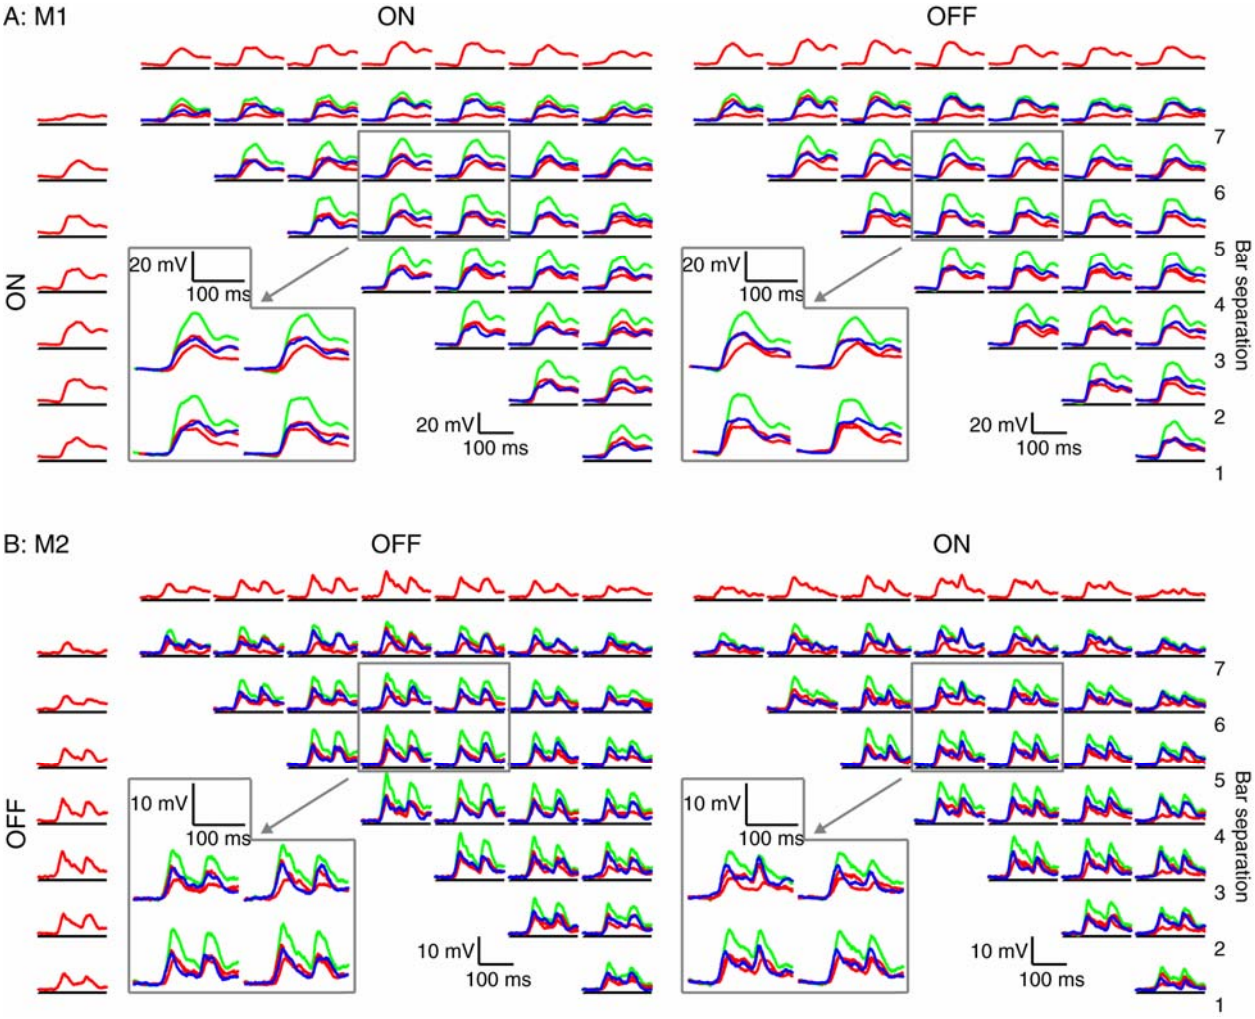
A: M1
ON
OFF
ON
20 mV
100 ms
20 mV
100 ms
20 mV
100 ms
20 mV
100 ms
Bar separation
1
2
3
4
5
6
7
B: M2
OFF
ON
OFF
10 mV
100 ms
10 mV
100 ms
10 mV
100 ms
10 mV
100 ms
Bar separation
1
2
3
4
5
6
7

Figure 3.04. **Classical responses in complex cells**.

A. ON-ON and ON-OFF matrices for a complex cell with varied responses to pairs of bars. At a bar distance of 2, the response to both bars together is consistently less than the maximum response to the bars alone for same-polarity stimuli (indicated by violet shading); the reverse is true for opposite-polarity stimuli (indicated by gray shading). At a bar distance of 5, same-polarity bars evoke a larger response together than individually.

B. A second complex cell's OFF-OFF and OFF-ON matrices showing the same behavior as in A.

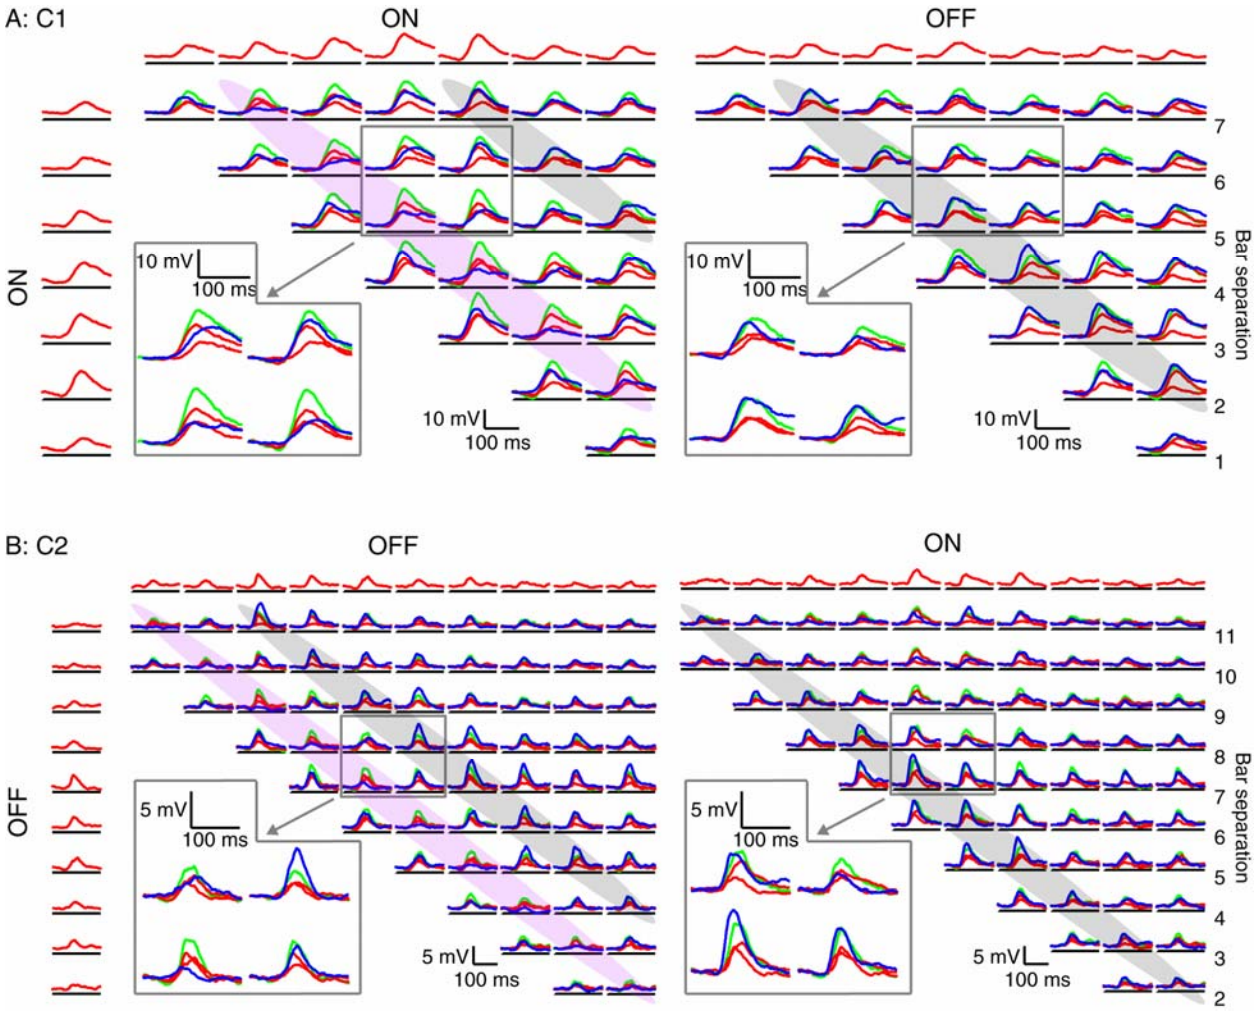
A: C1
ON
OFF
ON
10 mV
100 ms
10 mV
100 ms
10 mV
100 ms
10 mV
100 ms
Bar separation
1
2
3
4
5
6
7
B: C2
OFF
ON
OFF
5 mV
100 ms
5 mV
100 ms
5 mV
100 ms
5 mV
100 ms
Bar separation
2
3
4
5
6
7
8
9
10
11

Figure 3.05. **The remaining two matrices for cell M1**.

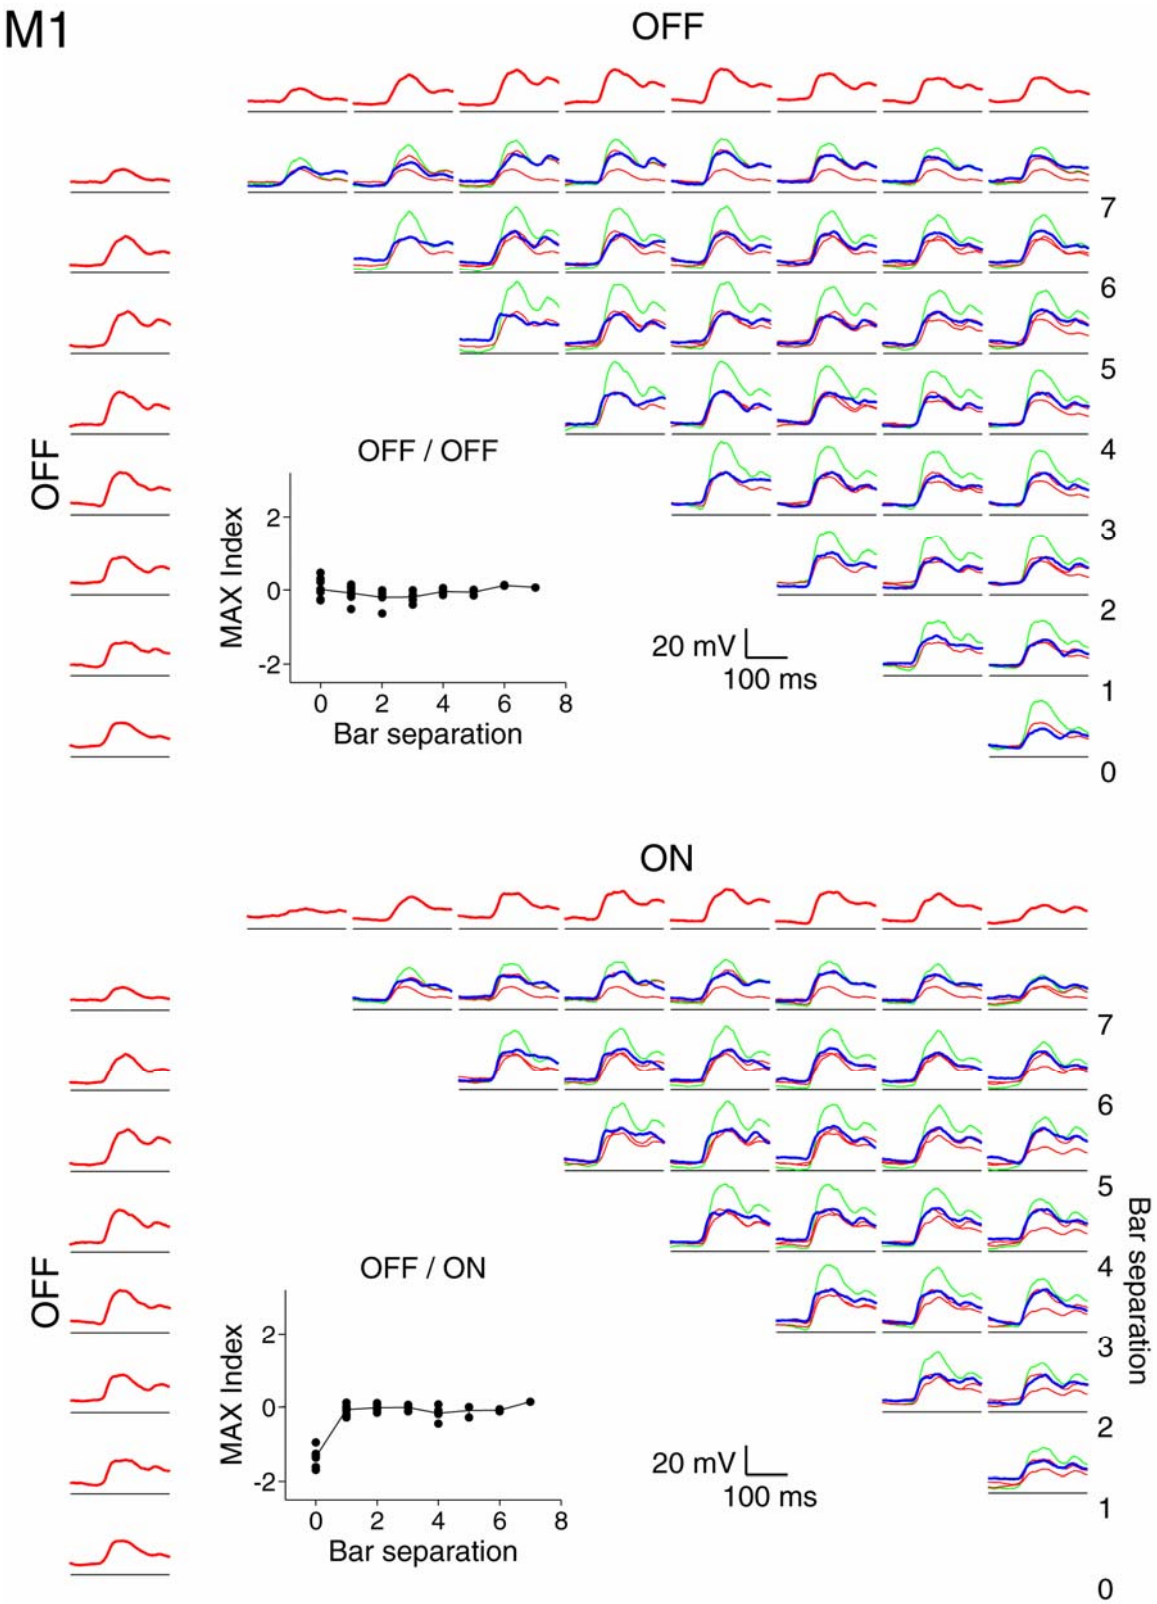
M1
OFF
OFF
OFF / OFF
MAX Index
2
0
-2
0 2 4 6 8
Bar separation
20 mV
100 ms
7
6
5
4
3
2
1
0
ON
OFF
OFF / ON
MAX Index
2
0
-2
0 2 4 6 8
Bar separation
20 mV
100 ms
Bar separation
7
6
5
4
3
2
1
0

Figure 3.06. **The remaining two matrices for cell M2**.

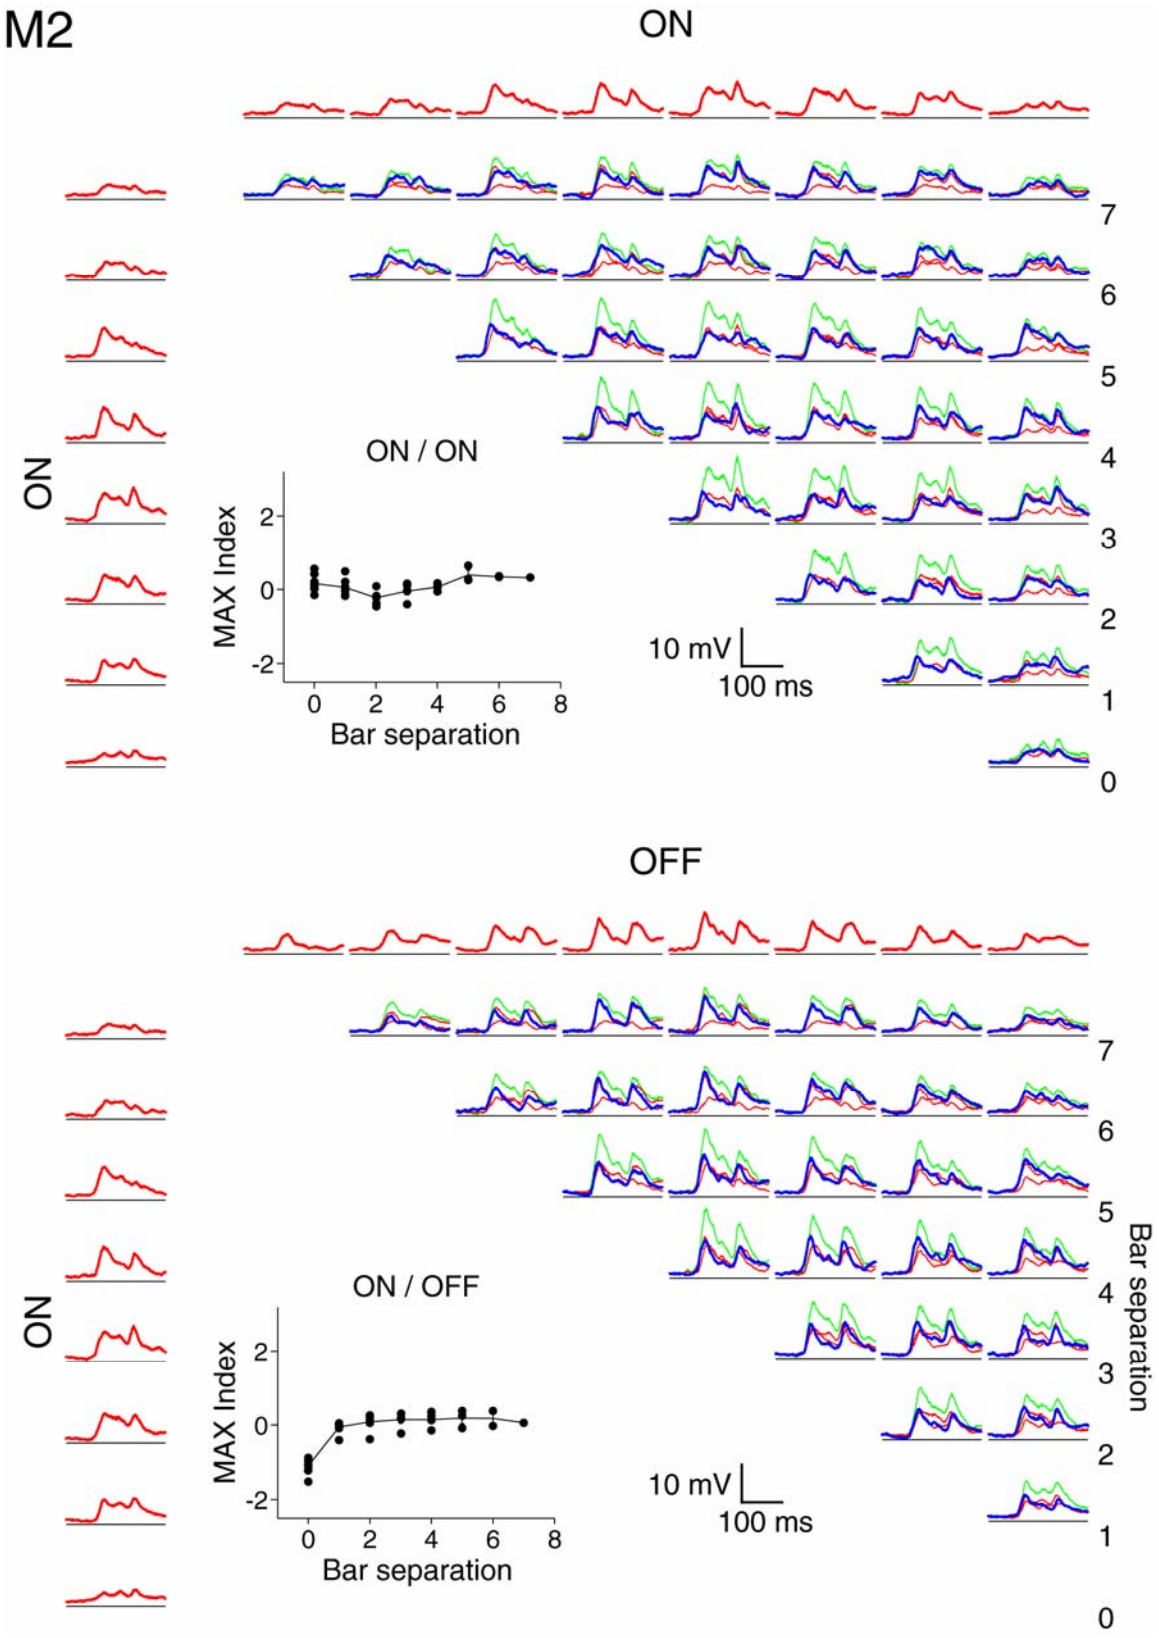
M2
ON
ON
ON / ON
MAX Index
Bar separation
10 mV
100 ms
OFF
ON
ON / OFF
MAX Index
Bar separation
10 mV
100 ms
Bar separation

Figure 3.07. **The remaining two matrices for cell C1**.

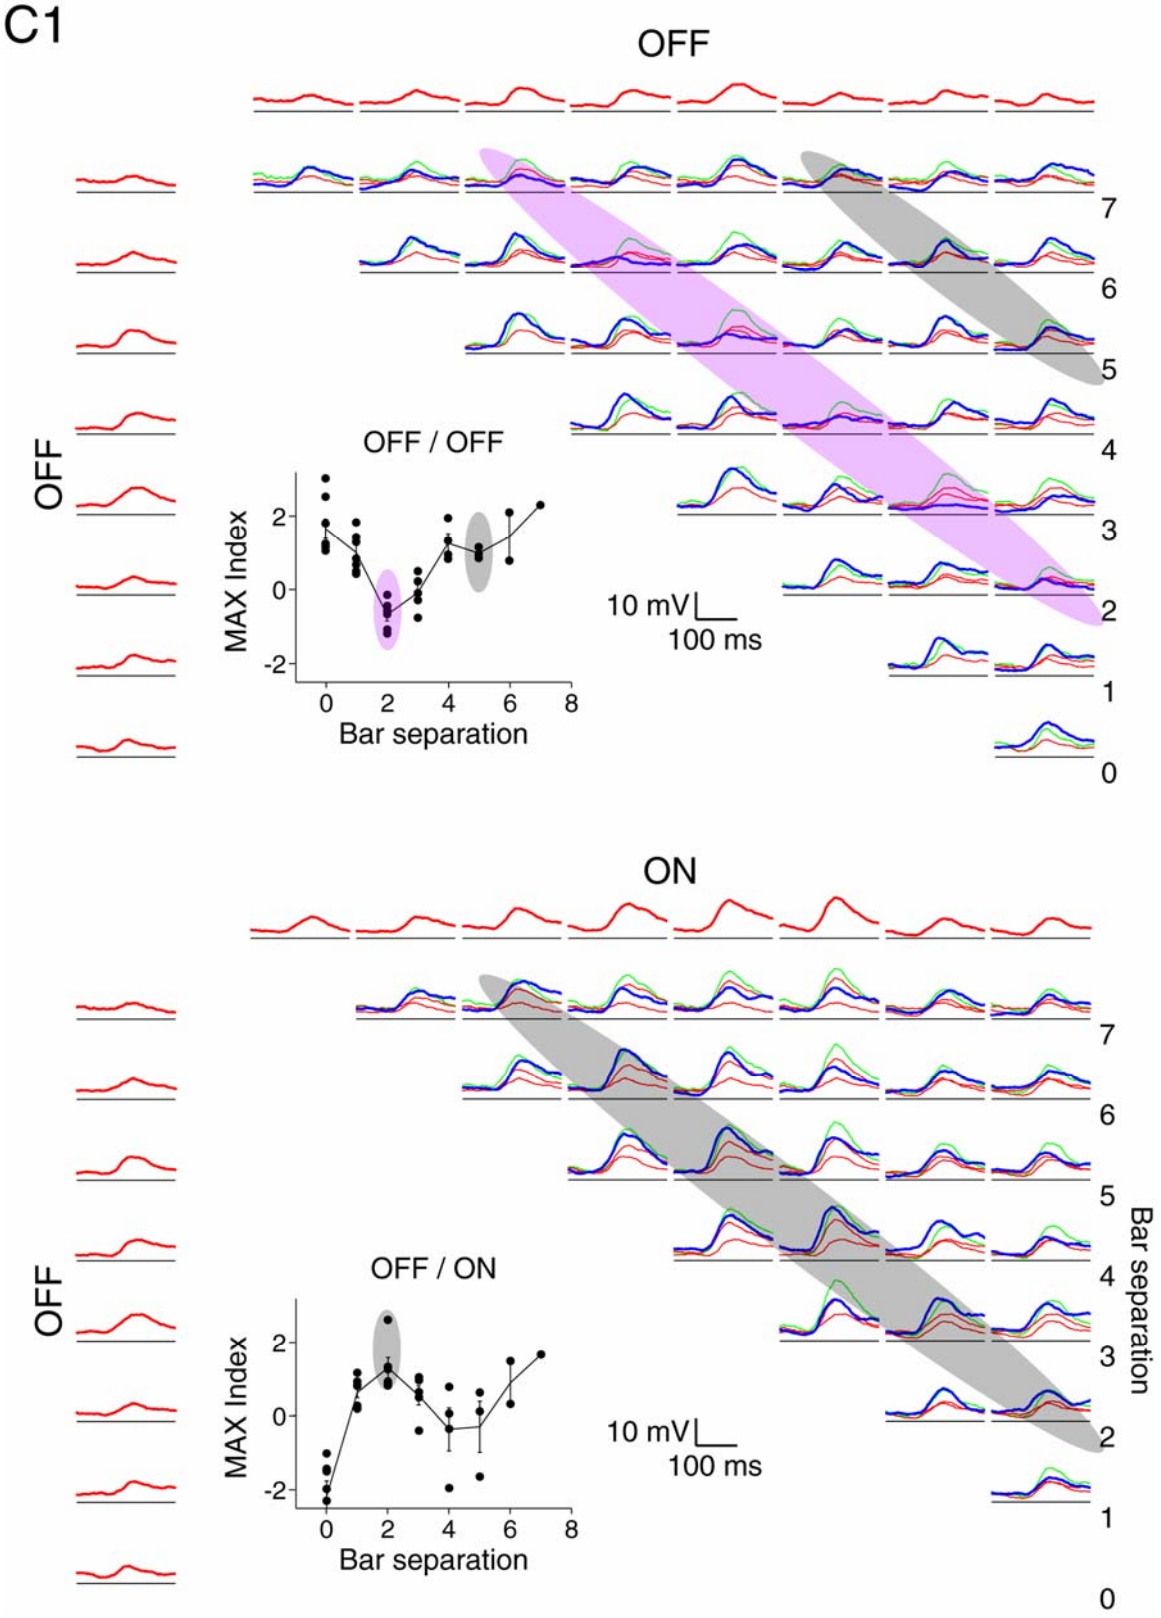
C1
OFF
OFF
OFF / OFF
MAX Index
Bar separation
10 mV
100 ms
ON
OFF
OFF / ON
MAX Index
Bar separation
10 mV
100 ms
Bar separation

Figure 3.08. **The remaining two matrices for cell C2**.

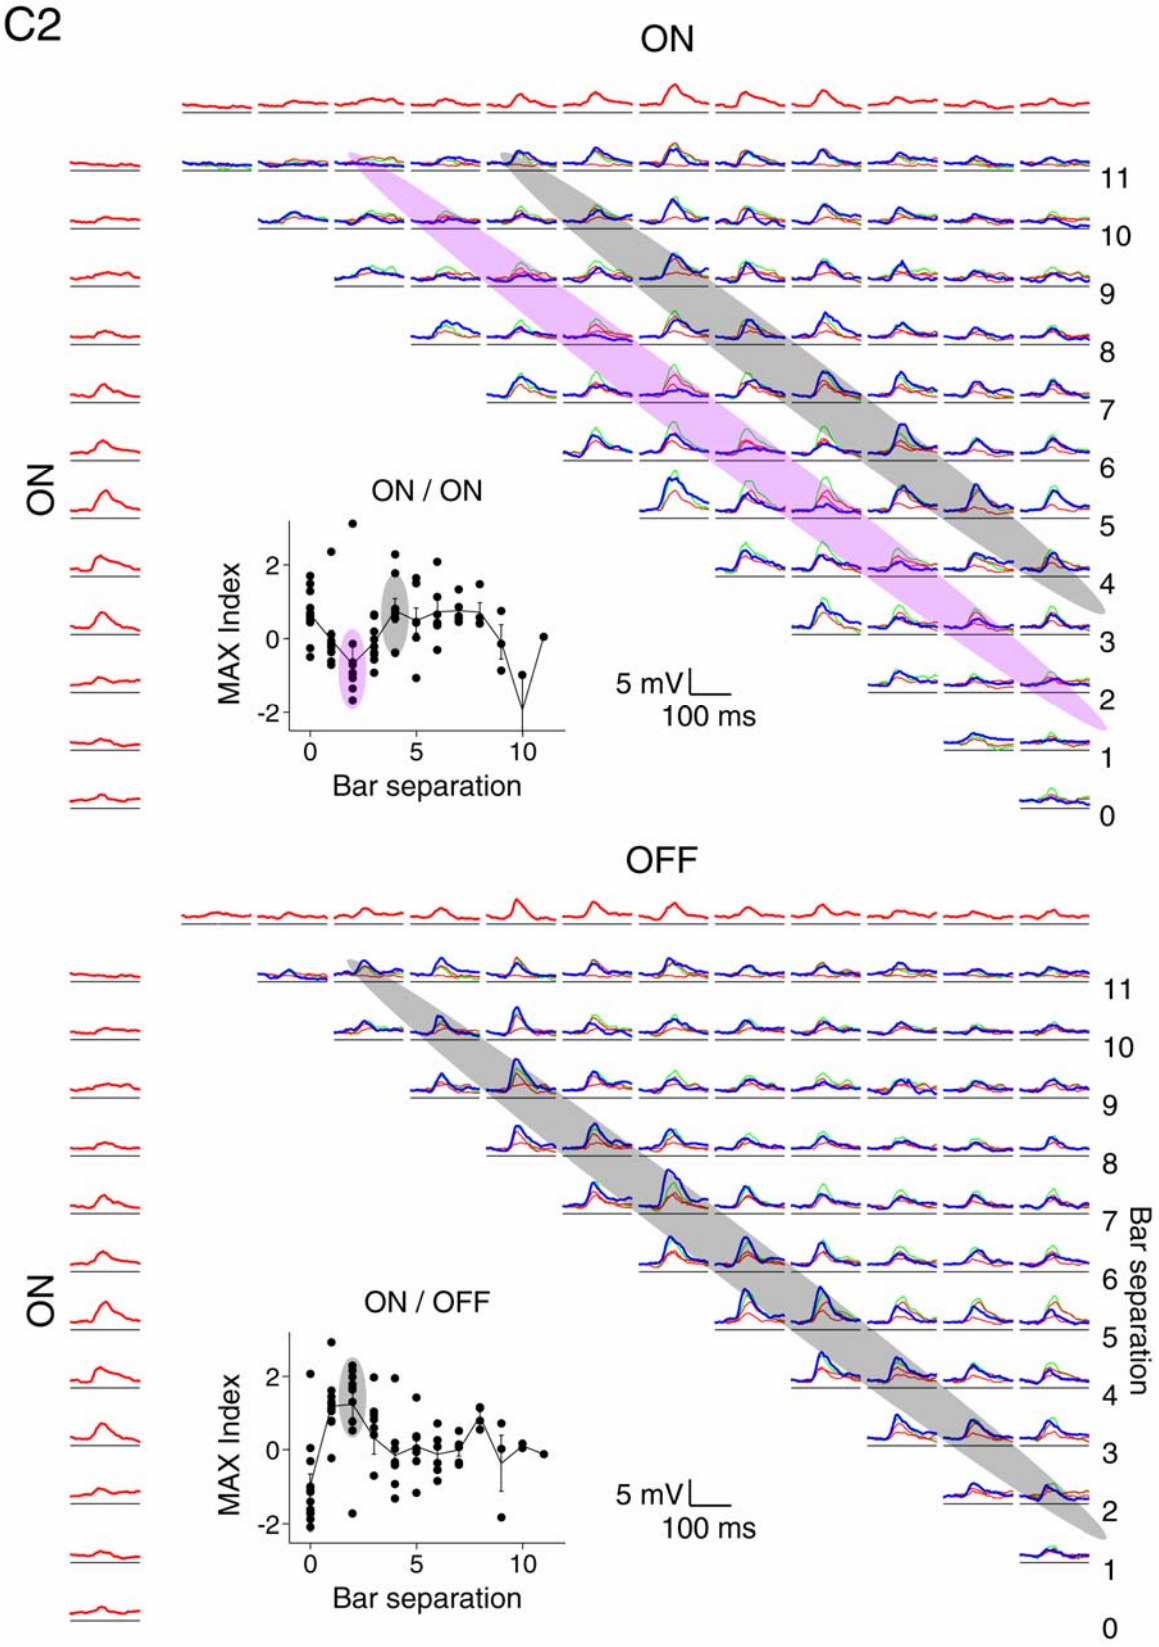
C2
ON
ON
ON / ON
MAX Index
Bar separation
5 mV
100 ms
OFF
ON
ON / OFF
MAX Index
Bar separation
5 mV
100 ms
Bar separation

Quantification of responses to flashed stimuli

For each cell we measured the amplitude of its responses to all flashed stimuli by calculating the average membrane potential within a 10ms window centered on the peaks of the responses. In separate graphs, we plotted the paired stimulus ($R_{a,b}$) amplitudes against the maximum, the minimum, and the linear sum of the component bar (R_a and R_b) amplitudes. Figure 3.09 shows the results of this analysis for three MAX-like cells (including M1 and M2 from Fig. 3.03) and three Classical cells (including C1 and C2 from Fig. 3.04, and C3 from Fig. 1 and 2).

For MAX-like cells, the maximum of the individual responses to the bars in a pair was a good predictor of the cell's response to the pair, while for Classical cells the prediction was poor. The MIN measure underpredicted the paired responses for both types of cells in almost all cases. As expected from the data in Figures 3.03 and 3.04, the LINEAR sum of the responses to individual stimuli generally overpredicted the paired response in the case of MAX-like cells, and did not accurately predict the paired response in the case of Classical cells.

To determine the MAX behavior of each stimulus pair, we computed a MAX index as:

$$MI(a,b) = \frac{R_{a,b} - MAX(R_a, R_b)}{MIN(R_a, R_b)}, \qquad (1)$$

This formulation was originally proposed by Sato (Sato, 1989), and has been used previously to study complex cells (Lampl et al., 2004). It has the benefit of being able to represent perfect MAX behavior (a max index of 0) as well as perfectly linear behavior (a max index of 1); in addition, facilitation (a paired response greater than the maximum of the individual responses) is

represented by values greater than 0, while suppression (a paired response less than the maximum of the individual responses) is represented by values less than 0.

The MAX index attempts to capture the relationships between three different quantities – R_a, R_b and $R_{a,b}$ – in a single number, and as such does not always represent all aspects of the relationships perfectly. For example, as a quotient, the MAX index can be overly sensitive to the size of the denominator (MIN response): If the MIN response is very small, the index will be artificially inflated, whereas if the MIN response is large, the index will be closer to perfect MAX (0) than would, intuitively, be expected. In our data the former problem arose infrequently, as 75% of the MIN responses we measured were larger than 25% of the corresponding MAX responses. In addition, Figure 3.09 illustrates that even for cells where the MAX and MIN responses were often comparable in amplitude (M1), the paired responses were, in general, better explained by the MAX responses than the MIN responses. These observations suggest that for our data the MAX index should serve well as a single, normalized index capable of differentiating between the spatial integration demonstrated by MAX-like cells and that demonstrated by Classical cells.

Figure 3.09. **Plots of complex cell response amplitudes**.

For MAX-like cells, the responses to bar pairs are most accurately predicted by the maximum of the responses to the individual bars (MAX column) across all spatial positions and bar polarities. The minimum and linear sum of the responses to individual bars do not predict well the paired responses (MIN and LINEAR columns). For Classical cells, the maximum, minimum, and linear sum of responses to individual bars all fail to predict the paired responses.

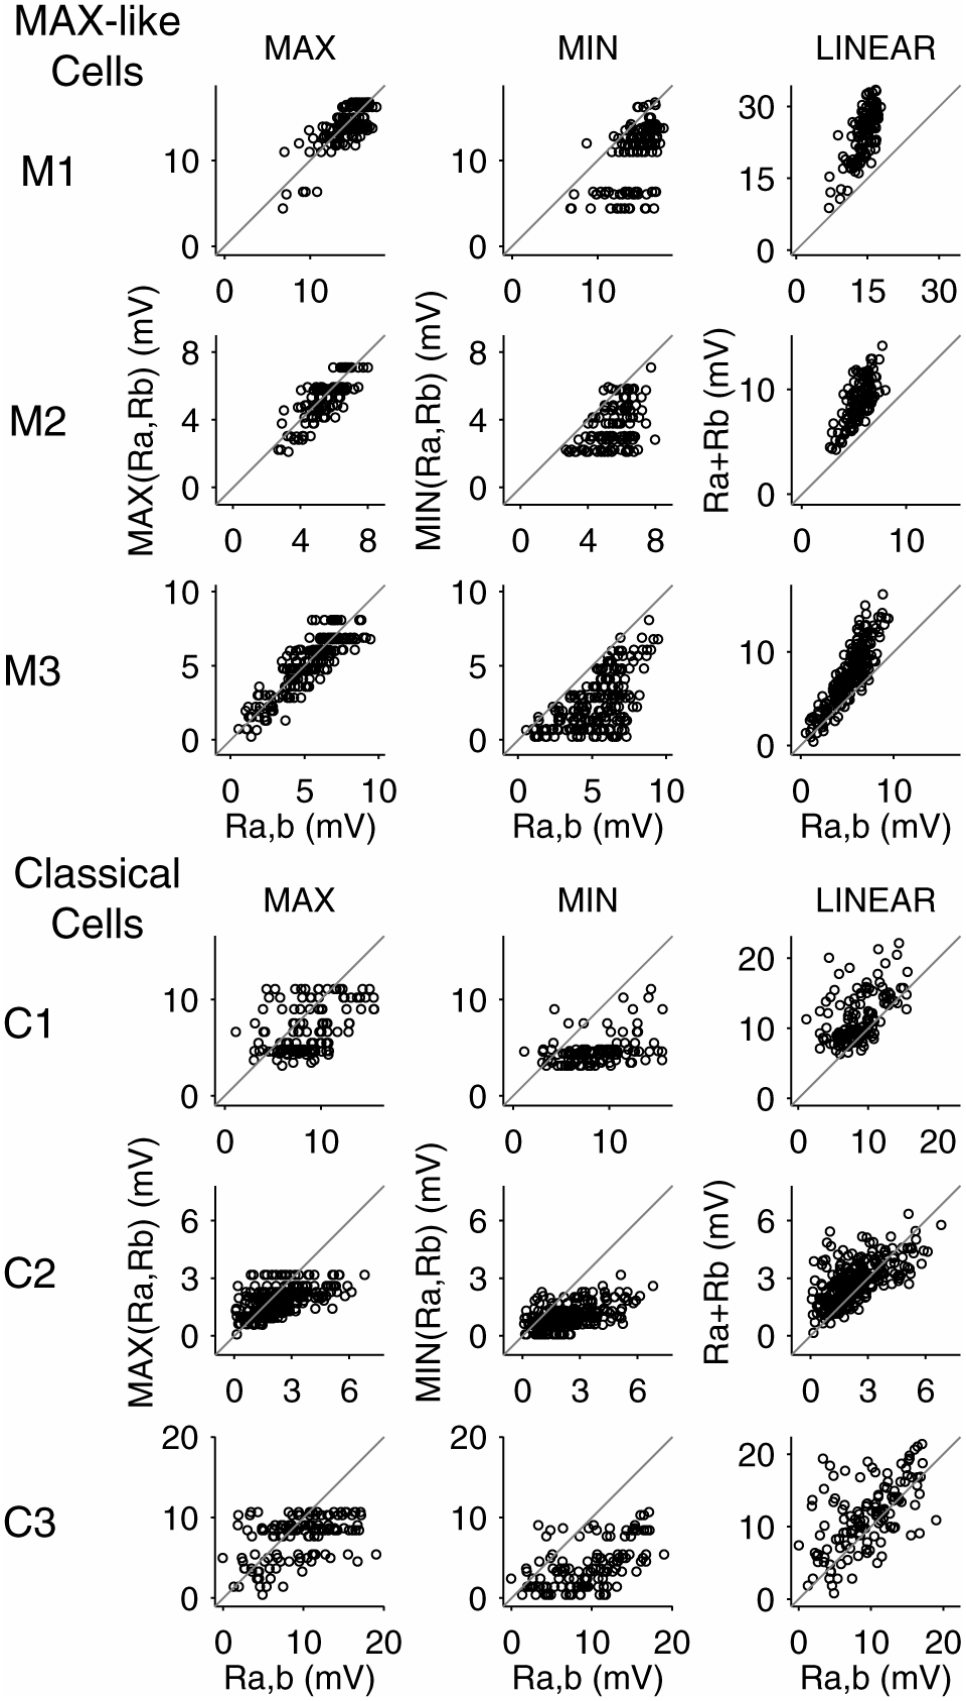
MAX-like Cells
MAX
MIN
LINEAR
M1
M2
M3
MAX(Ra,Rb) (mV)
MIN(Ra,Rb) (mV)
Ra+Rb (mV)
Ra,b (mV)
Classical Cells
C1
C2
C3

That the MAX index can faithfully represent the behavior seen in cells M1, M2, C1, and C2 is shown in Figure 3.10, where we plot MAX index against bar separation for all of the entries in the matrices from Figures 3.03 and 3.04. The MAX indices for M1 and M2 were all close to zero (individual points), as was the average MAX index at each bar separation (solid line). In these cells, the interaction between any two bars was almost completely independent of the separation and polarity of the bars. Cells C1 and C2, on the contrary, had strongly modulated MAX index profiles. The MAX indices clearly depended on the distance between bars, and the dependence flipped when one bar changed polarity (compare ON-ON with ON-OFF or OFF-OFF with OFF-ON); the shape of the dependence, however, was the same for matrices in which both bars had the same polarity (ON-ON compared with OFF-OFF) or in which the two bars had opposite polarity (ON-OFF compared with OFF-ON). The strong dependence of MAX index on bar separation thus explains why, for these cells, both the MAX and the LINEAR measures in Figure 3.09 failed to represent their behavior.

Figure 3.10. **MAX index plots for the matrices from Figures 3 and 4**.

The filled circles represent individual MAX indices, while the solid line tracks the mean of the indices at each bar separation. Cells M1 and M2 show relatively flat average MAX index profiles that do not vary greatly with bar separation or polarity. Cells C1 and C2 show modulated average MAX index profiles with suppression (MAX indices < 0) at small bar separations (purple ovals), and enhancement (MAX indices > 0) at larger bar separations (gray ovals) for same polarity bars. The pattern of suppression reverses when one of the bars switches polarity.

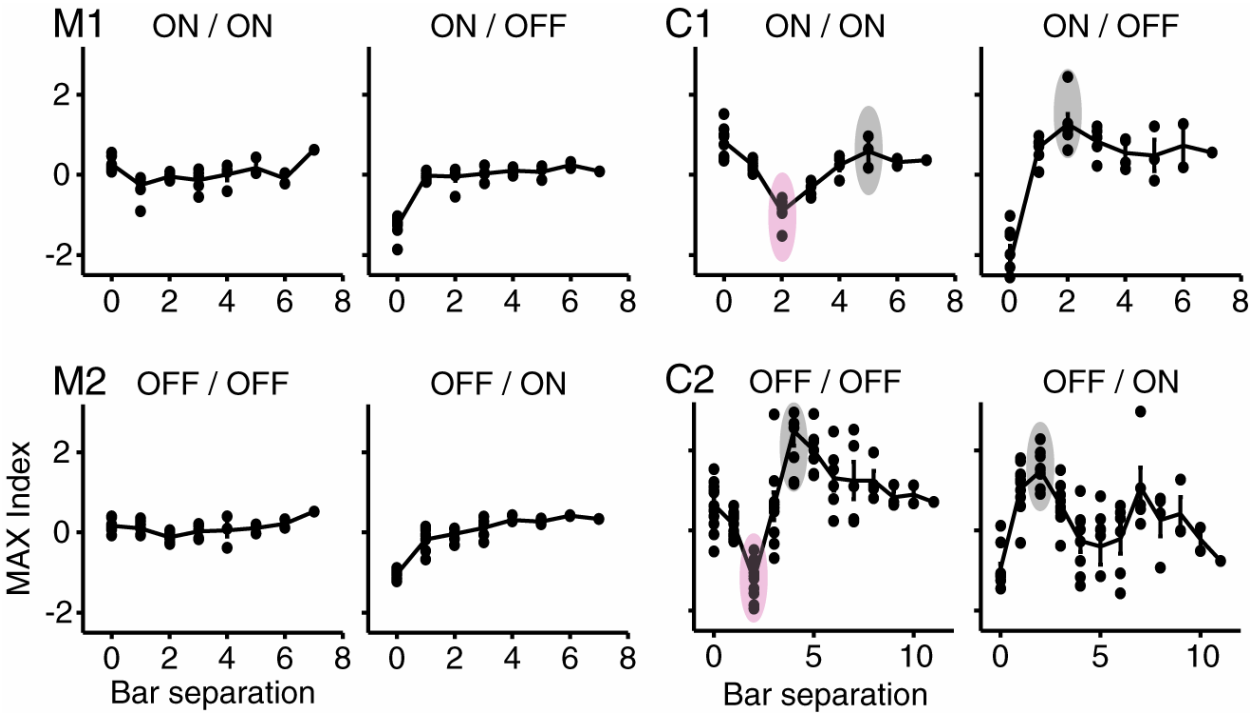
M1
ON / ON
ON / OFF
C1
ON / ON
ON / OFF
M2
OFF / OFF
OFF / ON
C2
OFF / OFF
OFF / ON
MAX Index
Bar separation
Bar separation

Quantification of overall MAX behavior

To determine whether the MAX-like behavior of cells such as M1 and M2 is distinct from the classical behavior of cells like C1 and C2, or whether there exists a continuum between MAX-like cells and Classical cells, we quantified the degree to which complex cells were sensitive to the separation between bar stimuli. For each cell, the indices from all four graphs of MAX index versus bar separation were combined (Fig. 3.11A and B, individual points). Indices for opposite-polarity pairs (ON-OFF and OFF-ON) were negated, and then at each bar separation, the points were averaged together to give a Composite MAX Index (Fig. 3.11A and B, solid lines).

Figures 3.11A and B show the composite MAX indices for a MAX-like cell (M1), and a more classical cell (C3). To quantify the sensitivity of the composite MAX index to bar separation, we took the difference between the maximum and minimum composite MAX indices, which we refer to as the spatial variation index or SVI. Cells like M1 and M2 with little or no modulation to their composite MAX index profiles had small SVI's (M1, 0.2; M2, 0.1). Cells with highly modulated composite MAX index profiles like C1 and C2 had larger SVI's (C1, 1.1; C2, 2.4). A histogram of the SVI for the 45 cells studied is plotted in Figure 3.11C.

We also applied a second method, one that does not rely on MAX indices, to evaluate the dependence of stimulus interactions on bar separation distance; this method, which is based on the difference between the average responses to same-polarity and opposite-polarity pairs at each bar separation, is detailed in Figure 3.12. The correlation between the two measurements of spatial variation was very strong ($R = 0.86$, $p < .001$). Neither metric gives any indication that

the population of complex cells is split into two distinct groups. Instead, both measures show a unimodal distribution (Fig. 3.12C).

Figure 3.11. **Method for measuring complex cell spatial variation**.

A. Average MAX indices (composite MAX index) for cell M1. The spatial variation index (SVI) is derived by taking the largest composite MAX index and subtracting the smallest composite MAX index.

B. Same as in A for the cell C3. **C**. Histogram of the SVI showing a continuum of MAX-like behavior over the population.

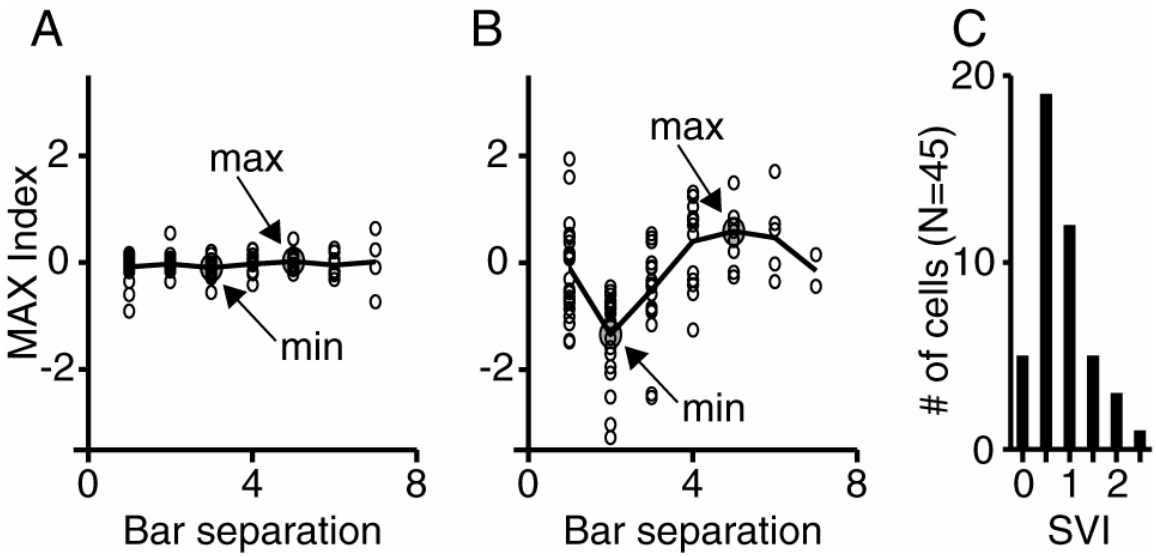
A
MAX Index
2
0
-2
max
min
0
4
8
Bar separation
B
2
0
-2
max
min
0
4
8
Bar separation
C
20
10
0
# of cells (N=45)
0 1 2
SVI

Figure 3.12. **A second method for measuring complex cell spatial variation**.

A. For each bar distance, we averaged together all of the responses to paired stimuli with the same polarities (ON-ON and OFF-OFF; gray traces), and separately averaged together all of the responses to paired stimuli with the opposite polarities (ON-OFF and OFF-ON; black traces). The differences between the peaks of the same-polarity and opposite-polarity average responses at each bar distance were measured and then normalized by the size of the largest average response (right). Max and Min indicate the bar distances that were subtracted to derive another spatial variation index (SVI_2). For this cell, the difference between same-polarity and opposite-polarity averages varied considerably with distance.

B. Same as in A for a cell with little difference between same-polarity and opposite polarity average responses at all bar separations.

C. Histograms of SVI_1 and SVI_2 showing that they had similar distributions.

D. Plot showing that the measures of spatial variation were highly correlated ($R = 0.86$, $p < .001$).

E. SVI_2 is also negatively correlated with spatial frequency bandwidth ($R = -.63$, $p < .001$).

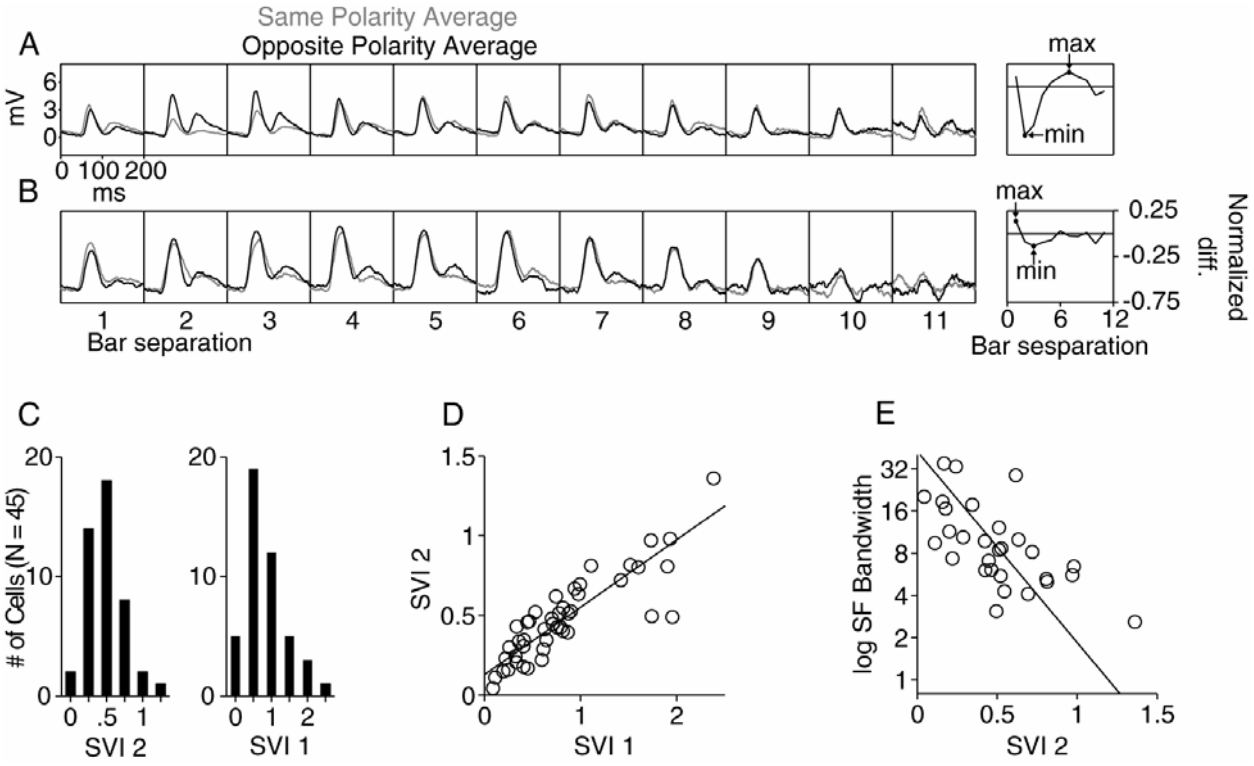
Same Polarity Average
Opposite Polarity Average
A
mV
6
3
0
0 100 200
ms
max
min
B
max
min
0.25
-0.25
-0.75
Normalized diff.
1 2 3 4 5 6 7 8 9 10 11
Bar separation
0 6 12
Bar sesparation
C
# of Cells (N = 45)
20
10
0
0 .5 1
SVI 2
0 1 2
SVI 1
D
SVI 2
1.5
1
0.5
0
0 1 2
SVI 1
E
log SF Bandwidth
32
16
8
4
2
1
0 0.5 1 1.5
SVI 2

Spatial frequency peak and bar-pair interactions

Movshon et al. (1978b) demonstrated that the spatial frequency tuning of some complex cells could be explained by the spatial interactions among paired bar stimuli. This relationship falls out naturally from the commonly used energy model of complex cell receptive fields (Adelson and Bergen, 1985). In the model, a single complex cell receives the squared output of two simple cells with overlapping receptive fields that are 90^{o} out of spatial phase from each other. The preferred spatial frequency of the complex cell derives from the preferred spatial frequencies of the presynaptic simple cells, which are assumed to be identical. Because simple cells are roughly linear in spatial summation, their preferred spatial frequencies are, in turn, determined largely by the distance between their subfields. At the same time, subfield separation determines the way in which the simple cells, and therefore the complex cells they project to, respond to paired bars. When bars of opposite polarity are separated by the same amount as the underlying subfields – that is, when a bright bar falls in a simple cell's ON region and a dark bar falls in the simple cell's OFF region – the simple cell (and consequently the complex cell) will respond strongly. Conversely, bars of the same polarity that fall one in an ON region and one in an OFF region will antagonize each other such that the simple cell (and the complex cell) will not respond strongly (Fig. 3.14).

According to the energy model, then, the separation between same-polarity bars that generates maximum suppression in a complex cell (or equivalently the separation between opposite-polarity bars that yields maximum facilitation), should be one-half the reciprocal of the cell's preferred spatial frequency. In terms of the MAX index, the reciprocal of preferred spatial frequency should equal the bar separation for which the composite MAX index is the smallest.

We tested this expectation by measuring the spatial frequency tuning of complex cells with drifting gratings. The cycle averaged responses of an example cell are shown in Figure 3.13A. Tuning curves were constructed from the mean (DC) depolarization at each spatial frequency and fit with a 6-parameter function (see methods). In Figure 3.13B tuning curves for 6 different cells are shown, 3 with very classical behavior (top) and 3 with MAX-like behavior (bottom).

The comparison between preferred spatial frequency and paired bar interactions is shown in Figure 3.13C. Here we plot the composite MAX index against bar separation for the same cells as in Figure 3.13B. The vertical arrow in each graph points to the bar separation at which maximum suppression should be observed considering the cell's preferred spatial frequency in response to drifting gratings. That is, the position of the arrow is at one-half of the reciprocal of the preferred spatial frequency. Note that the prediction is plotted in terms of the bar width used for stimulation (rather than degrees of visual angle).

For cells that were more classical in their behavior (strong dependence of the MAX index on bar separation; Fig. 3.13C, top), the predictions based on preferred spatial frequency clearly matched the bar separations at which the composite MAX index was minimal (filled circles). For MAX-like cells (cells with weak dependence of MAX index on bar separation; Fig. 3.13C, bottom), the match was more variable. To capture this observation for the population, we quantified the match by taking the difference between the location of the arrow in Figure 3.13C and the location of the minimum of the composite MAX index. We then plotted this difference against the spatial variation index for each cell (Fig. 3.13D). As can be seen from the graph, the more classical a

cell's behavior was (the larger the cell's SVI), the better the match between the predicted and measured bar separation yielding a minimal composite MAX index.

Figure 3.13. **Complex cell preferred spatial frequencies predict the distance of maximum paired-bar suppression**.

A. Cycle averaged responses of a complex cell to drifting gratings of varied spatial frequency.

B. Spatial frequency tuning curves for 6 cells with similar preferred spatial frequencies; their bandwidths, however, varied considerably. Smooth curve is a least-squares fit to the data (closed circles).

C. Plots of composite MAX indices for the cells in B and predictions of minimal composite MAX index (filled circles) from the preferred spatial frequency (vertical arrows).

D. Difference between the location of the smallest composite MAX indices and the predictions based on spatial frequency plotted against SVI. The predictions are worse for MAX-like cells.

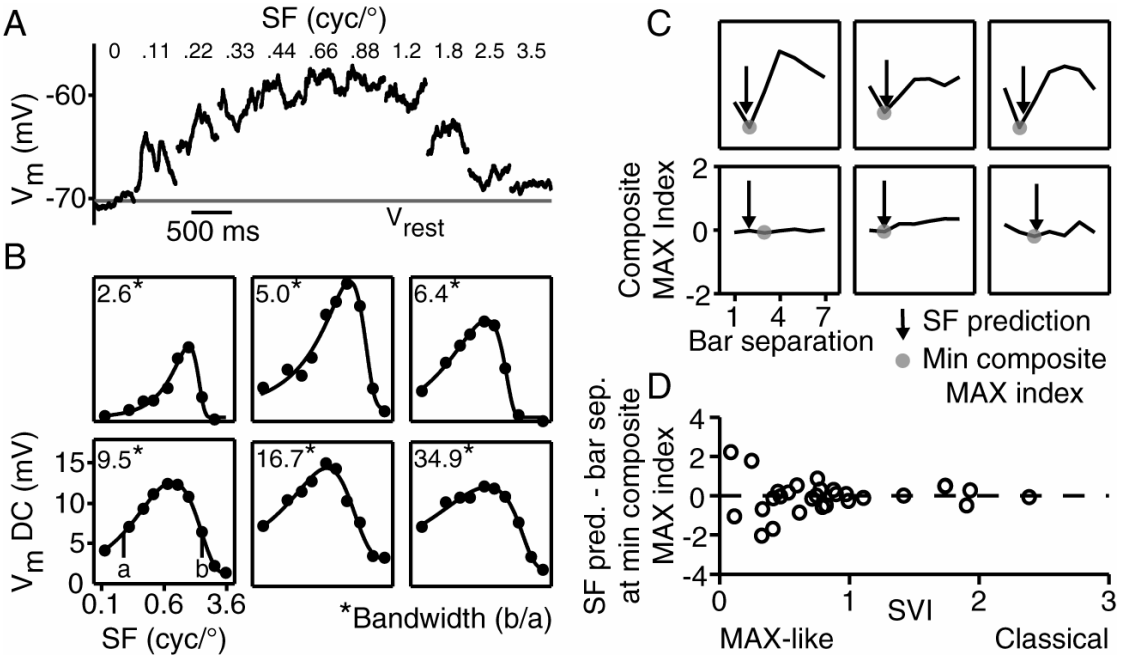
A
SF (cyc/°)
0 .11 .22 .33 .44 .66 .88 1.2 1.8 2.5 3.5
-60
-70
V_m (mV)
500 ms
V_{rest}
B
2.6*
5.0*
6.4*
9.5*
16.7*
34.9*
15
10
5
0
V_m DC (mV)
a
b
0.1 0.6 3.6
SF (cyc/°)
*Bandwidth (b/a)
C
Composite MAX Index
2
0
-2
1 4 7
Bar separation
SF prediction
Min composite MAX index
D
SF pred. - bar sep. at min composite MAX index
4
2
0
-2
-4
0 1 2 3
SVI
MAX-like
Classical

Figure 3.14. **Example of suppression in the energy model**.

In these cells bar A alone evokes no response (cell 1) and a hyperpolarization (cell 2); bar B alone evokes a depolarization in both of the cells. When their output is squared, both simple cells contribute excitation to the model complex cell. When bars A and B are presented together, they interfere in cell 2 so as to lead to no output, and as bar A evokes no response from cell 1, the only input to the complex cell is from bar B in cell 1. This total excitation is less than the total excitation provided by each bar individually.

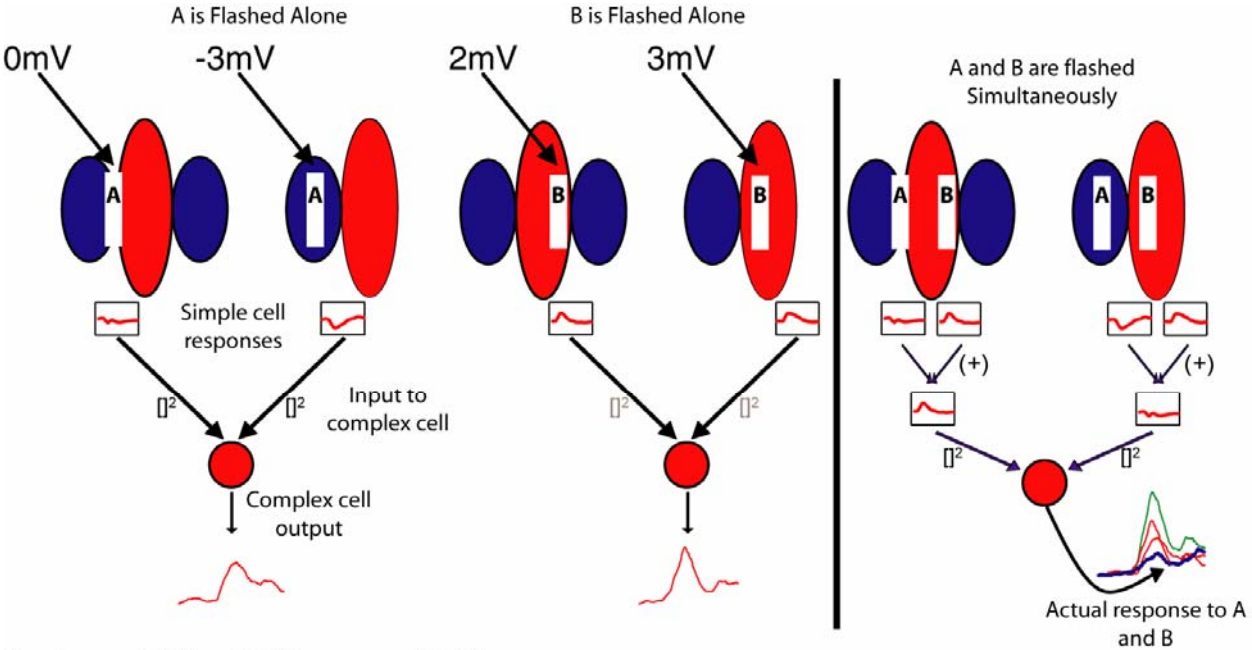

A alone: $(0)^2 + (-3)^2 \longrightarrow 9mV$

B alone: $(2)^2 + (3)^2 \longrightarrow 13mV$

A+B: $(0+2)^2 + (-3+3)^2 \longrightarrow 4mV$

Spatial frequency bandwidth and bar-pair interactions

An underlying assumption of the energy model is that the preferred spatial frequencies of the presynaptic simple cells are identical. In this way, bars of a given polarity and separation will interact regardless of their absolute position within the complex cell's receptive field. One possible difference between classical complex cells and more MAX-like complex cells, then, is that in the former the constituent simple cells match in spatial frequency, whereas in the latter they do not. The spatial frequency tuning curves of Figure 3.13B are suggestive in this regard, in that the classical complex cells (top row) are more narrowly tuned for spatial frequency than the MAX-like cells (bottom row). This relationship would be consistent with the simple-cell inputs to MAX-like cells having a range of preferred spatial frequencies, rather than matching preferred spatial frequencies.

To determine whether spatial frequency bandwidth is correlated with MAX-like behavior, we derived bandwidths for each cell from the fitted tuning curves by dividing the high and low spatial frequency cut-offs (at which half-maximal responses were observed). Bandwidths for curves in Figure 3.13B are indicated by the number in the upper left corner of each panel. In Figure 3.15A, bandwidth is plotted against the degree of MAX-like behavior for each cell (spatial variation index from Figure 3.11). The plot shows a significant negative correlation ($R = 0.65$; $p < .001$).

Figure 3.15C demonstrates explicitly the relationship between MAX behavior and spatial frequency tuning bandwidth. Here we chose two subsets of complex cells, the eight cells with the narrowest bandwidths (Fig. 3.15B, red), and the eight cells with the widest bandwidths (Fig.

3.15B, blue). For each subset, we averaged together the plots of MAX index versus bar separation distance (plots like those in Figure 3.11A and B). Before combining the plots, however, we rescaled the x-axis in units of optimal spatial period. That is, a value of 1 on the x-axis is equal to the reciprocal of the preferred spatial frequency of a given cell, as measured with drifting gratings. Using this normalization, the spacing of points on the x-axis became different for each cell, since the bar width in relation to the preferred spatial period was not the same over the population. To obtain the average curves, then, the entire set of points for each group of 8 cells was binned in intervals of ~¼ of the preferred period.

It is clear from the graphs in Figure 3.15C that cells with large spatial frequency bandwidths (blue line and S.E.M bars) show significantly less modulation in their MAX index profiles than cells with narrow bandwidths (red line and S.E.M bars). Both curves peak at one-half the preferred spatial period, as would be expected from the energy model. The curve for the large bandwidth (> 12) cells, however, has a peak amplitude less than half that of the small-bandwidth (< 6) cells, the peak is broader, and the curve falls nearly to 0 for larger bar separations. Overall, then the average MAX index depends only weakly on bar separation for cells with large spatial frequency bandwidths.

Figure 3.15. **The relationship between spatial frequency bandwidth and complex cell spatial variation**.

A. Plot showing negative correlation between the spatial variation index and spatial frequency bandwidth for 29 complex cells (R = 0.65).

B. Histogram of spatial frequency bandwidth over the population.

C. MAX indices for 8 cells with narrow spatial frequency bandwidths (< 6, red cells in B) were averaged after converting their raw bar distances to multiples of the preferred spatial period (red curve). Maximal suppression and facilitation for the average MAX indices occurred around 0.5, as predicted. MAX indices for 8 cells with broad spatial frequency bandwidths (> 12, blue cells in B) were also averaged (blue curve). The curve stayed much closer to the line y=0, or perfect MAX behavior, than the curve for cells with narrow spatial frequency bandwidths.

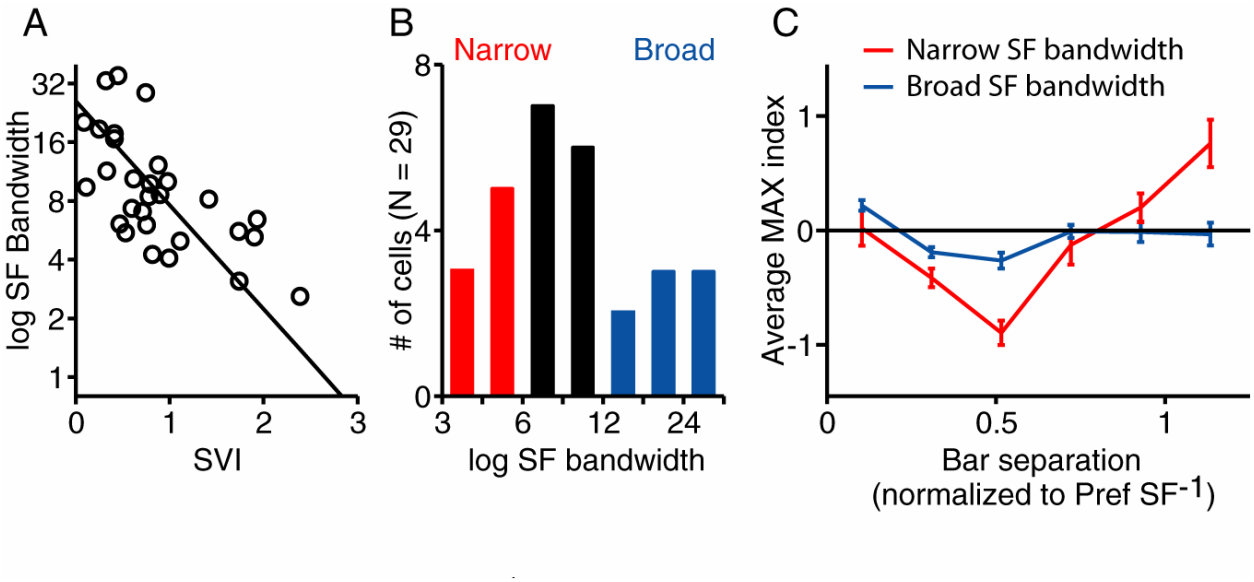
A
log SF Bandwidth
32
16
8
4
2
1
0
1
2
3
SVI
B
Narrow
Broad
# of cells (N = 29)
8
4
0
3
6
12
24
log SF bandwidth
C
Narrow SF bandwidth
Broad SF bandwidth
Average MAX index
1
0
-1
0
0.5
1
Bar separation
(normalized to Pref SF-1)

Spatial frequency bandwidth in the energy model

We proposed above that one way to generate MAX-like cells is to modify the energy model (Adelson and Bergen, 1985) by combining inputs from simple cells with a range of preferred spatial frequencies. To test the effect of this scenario on MAX-like behavior in complex cells, we created two versions of the energy model. In the classical version, the complex cells received input from a single pair of simple cells with identical preferred spatial frequencies and receptive fields 90° out of spatial phase with one another. The sensitivity of each simple cell varied (in the direction perpendicular to the preferred orientation) as a Gabor function of distance:

$$S_1 = \cos(1.5\pi x) \bullet e^{-x^2/0.23}$$

$$S_2 = \sin(1.5\pi x) \bullet e^{-x^2/0.23}$$

In the MAX-like version of the model, the complex cell received input from an additional pair of simple cells with a preferred spatial frequency of a little less than half that of those in Pair 1.

$$S_3 = \cos(0.6\pi x) \bullet e^{-x^2/0.23}$$

$$S_4 = \sin(0.6\pi x) \bullet e^{-x^2/0.23}$$

For each version of the model, we presented 6000 pairs of bars, either with the same or opposite polarity, that were one-eighth the size of the model cells' receptive fields, with the bar positions chosen randomly and constrained only to be non-overlapping. The response of each constituent simple cell to a bar was obtained by summing the extent of the simple cell's receptive field covered by the bar (taking the dot product of the stimulus and the receptive field sensitivity profile); the response to bar pairs was taken as the sum of the responses to each bar individually. All responses were measured relative to rest (0 mV). To obtain the response of the complex cells, the output of each simple cell was squared and then all of the outputs were summed, with the

result normalized by taking the square root (for an equivalent and more physiological formulation, see Discussion). MAX indices were calculated and then grouped according to the distance between bar centers in the same manner as the experimental data; note that here, as in Figures 3.11A and B and Figure 3.13C, the MAX indices for opposite polarity pairs were negated prior to averaging.

The average MAX indices for both versions of the model are plotted in Figure 3..16. The two simple models successfully captured many qualitative features of the curves in Figure 3.15C. With simple cell pair 1 alone providing input, the average MAX index for the model complex cell was strongly dependent on bar separation, with a pronounced trough at a separation of ½ the preferred spatial period, similar to the real data. At larger separations the average MAX index reached well above 0, again matching the behavior of real complex cells seen in Figure 3.15C. Most importantly, in the model complex cell with input from pair 1 and pair 2, both the trough in the average MAX index at small distances and the peak at larger distances were reduced, as was observed in the real complex cells with large spatial frequency bandwidths. Thus, even in this extremely simplified model, input from simple cells with different preferred spatial frequencies prevented the full expression of suppression and facilitation that is expected from an energy model with simple cells that have identical spatial frequency tuning.

Figure 3.16. **Multiple spatial frequency channels in an energy model of complex cells produce MAX-like responses**.

An energy model with two simple cells in quadrature phase (Pair 1), whose responses were measured relative to rest (0 mV) and squared (which is physiologically equivalent to using 4 simple cells with rectification at rest and half-squaring), generates MAX index plots similar to those for complex cells with narrow spatial frequency bandwidths (red curve). An energy model with two pairs (Pair 1 and Pair 2) of simple cells having different preferred spatial frequencies generates MAX index plots similar to those for complex cells with broad spatial frequency bandwidths (blue).

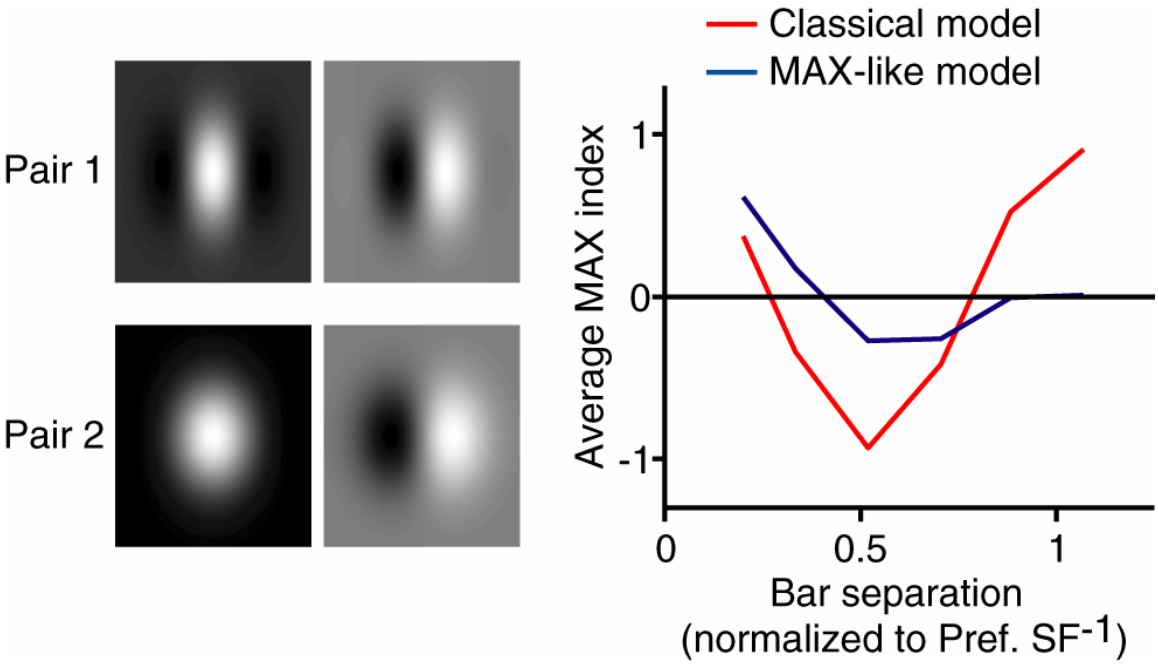
Pair 1
Pair 2
Classical model
MAX-like model
Average MAX index
1
0
-1
0
0.5
1
Bar separation
(normalized to Pref. SF-1)

Discussion

Lampl et al. (2004) observed MAX-like behavior in complex cells, but because single cells were not tested with a complete stimulus set and the indices measured from many cells were aggregated in their analysis, the authors could only draw conclusions about the average behavior of cortical complex cells. We have studied complex cells in more detail and observed that the manner in which individual complex cells integrate stimuli across their receptive fields is diverse. Many cells exhibit the classical patterns of interaction observed in previous extracellular experiments such as those of Movshon et al. Others perform an almost perfect MAX-like computation similar to that reported by Lampl et al. (2004)

A continuum of complex cell behavior

Although we make reference to these two patterns of response, it is clear that the distinction between MAX-like and classical complex cells is one of degree and not one of type. Quantitative measures of MAX-like behavior (the spatial variation index in Fig. 3.11 and a second index in Fig. 3.12) showed that complex cells fall along a continuum. This observation is in line with the characterization of complex cells as an extremely heterogeneous group which defies easy subcategorization. We have not attempted here to match various classification schemes for complex cells (for example, the A, B, and C subgroups of Henry et al, 1978, or the special complex cells many of which project to the superior colliculus (Palmer and Rosenquist, 1974; Gilbert, 1977)) with variation along the SVI axis. Further study will also be required to determine whether MAX-like behavior varies systematically among cortical layers. A preliminary analysis correlating electrode depth with SVI for our data would suggest, however,

that it does not. SVI also does not correlate with the F1/F0 metric for classifying simple and complex cells (Skottun et al., 1991). For most of the complex cells in this study, F1/F0 lay between 0.2 and 0.3.

Evidence for models of complex cell formation

Hubel and Wiesel (1962) proposed a hierarchical model in which simple cells converge onto complex cells, preserving orientation selectivity but generating phase insensitivity. Some complex cells, however, particularly those in lower layer 3 and below, receive direct input from relay cells of the lateral geniculate nucleus. Alternatives to the hierarchical model have made use of this connection to construct complex cell receptive fields. Mel. et al. (1998), for example, proposed that individual complex cell dendrites could act similarly to simple cell subfields by integrating input from LGN relay cells aligned in space; complete simple cell-like input would then be mimicked by the integration of multiple dendrites in the complex cell soma. There is as yet, however, no clear evidence that this scheme is employed in cortex.

Evidence for the hierarchical model, both direct and indirect, has come in many forms. Based on the matches obtained between two-bar profiles and predictions of those profiles derived from spatial frequency tuning, Movshon et al. (1978b) concluded that spatial frequency tuning for some complex cells could be explained by the convergence of input from relatively linear subunits similar to simple cells. Anatomical experiments (Gilbert and Wiesel, 1979) show a strong projection from layer 4, which contains a large number of simple cells, to layers 2 and 3, where complex cells predominate. Alonso and Martinez (1998) showed that complex cells in the upper layers receive monosynaptic excitatory input from simple cells with overlapping and

similarly oriented receptive fields.

Complex cell energy models

The energy model of Adelson and Bergen (1985) demonstrated theoretically that a complex cell could be built from the squared output of two linear filters (approximations to simple cells) with receptive fields offset by 90° in spatial phase, and indirect evidence for such a mechanism has been obtained from analyses of 1-D (bars) and 2-D (spots) white noise experiments on complex cell receptive fields (Touryan et al., 2002; Rust et al., 2005; Touryan et al., 2005). We note that the energy model as traditionally formulated has some physiological correlates, but is not altogether realistic. For example, the squaring of simple cell input implies that simple cells are only sensitive to the magnitude of a contrast change but not its sign. A more realistic model would substitute two simple cells for each one simple cell in the model, with the new simple cells having identical receptive field positions but opposite subunit signs. Each simple cell would then give half-squared output with rectification at rest, the half-squaring arising from the power-law relationship between membrane potential and spike rate as observed previously (Hansel and van Vreeswijk, 2002; Miller and Troyer, 2002; Priebe and Ferster, 2005).

In the original energy model, the output of the complex cells is un-normalized, being simply the squared sum of the inputs. Because squaring is an expansive nonlinearity, the range of MAX indices produced by such a model is far larger than those observed in real complex cells. The square-root normalization of the complex cells' output we have employed in our model is designed to make the relationship between stimulus contrast and response amplitude more physiological and to bring the MAX indices into a more realistic range, between approximately -

1 and 1. While the square root is a computationally convenient way of performing this normalization, it carries no theoretical significance. Other compressive or saturating nonlinear functions, such as the soft-MAX function or divisive normalization would serve equally well (Riesenhuber and Poggio, 1999, 2002). Physiological mechanisms that could contribute to the normalization step in the model include synaptic depression, depolarization-induced changes in driving force on synaptic currents, and synaptic inhibition (Heeger, 1992).

Conclusions

In this work we have characterized complex cells on the basis of their membrane potential responses. Though we have not made a systematic study of the MAX-like behavior of complex cells spike responses, it is likely that complex cells would show a range of behaviors in spiking similar to what we find in the membrane potential responses. Lampl et al. (2004) reported that complex cell spike responses were MAX-like on average. And two example cells, one MAX-like cell and one classical, in which the membrane potential and spike responses were similar in character, are shown in Figure 3.17. A quantitative characterization of the range of complex cell MAX behavior in spiking would require a comprehensive extracellular study.

On the face of it, MAX-like behavior might seem contradictory to that expected from a hierarchical or energy model, for in such models the interactions between stimuli at each stage propagate in a manner that should reflect the combination of stimuli as opposed to favoring one stimulus over the other. As we have shown, however, incorporating simple cells with different preferred spatial frequencies into the energy model can account for important aspects of MAX-like complex cells, including broader spatial frequency tuning and a minimal dependence of

MAX indices on bar separation distance or polarity. The model as presented was designed to replicate the average behavior of MAX-like complex cells (Figure 8), which shows some variation in the MAX index with bar separation. Changing the saturation function for complex cell output and adding more simple cells with different preferred spatial frequencies can produce responses that are closer to the almost perfectly invariant behavior seen in some individual cells, such as in Figure 3.11A.

The energy model is parsimonious in that it can account for both MAX-like and classical complex cells, the only difference being the preferred spatial frequencies of its component simple cells. Other models are possible, however. As suggested in Lampl et al. (2004), a MAX response to a bar pair would be expected if the bar that evoked the larger response also evoked strong shunting inhibition. The authors further suggest that measuring the conductance changes evoked by paired bar stimuli could potentially determine whether inhibitory mechanisms contribute to MAX-like behavior. A second possibility is that direct relay-cell inputs could be combined onto the complex cell dendrites to mimic the energy model (Mel et al., 1998), but with the dendritic subunits each having different preferred spatial frequencies. Experiments to determine whether MAX-like and classical complex cells receive different amounts of monosynaptic input from geniculate relay cells could help to address this question.

Whether MAX-like behavior arises from the convergence of multiple spatial frequency channels or through a different network mechanism (Ohzawa et al., 1990; Mel et al., 1998; Chance et al., 1999; Tao et al., 2004; Serre et al., 2007), our observation that a subset of complex cells in primary visual cortex compute a MAX operation over a wide variety of stimuli affirms the

conclusion of Lampl et al. (2004). The authors wrote that the visual system may employ MAX computing complex cells to achieve robust object recognition as in the model proposed by Riesenhuber and Poggio (1999). Given their model's recent successes in high-level object recognition (Serre et al., 2007), the potential is great for a computational approach to continue to inform biological experimentation, and vice-versa, in helping to understand the functioning of neo-cortex.

Figure 3.17. **Intracellular MAX-like and classical responses are reflected in spike rate**.

A and **B**. Example voltage (A) and spike rate (B) plots for a classical complex cell, showing that the shape of the MAX index dependence is similar.

C and **D**. Example voltage (C) and spike-rate (D) plots for a MAX-like complex cell. Note that for both of the spike-rate plots, MAX indices were excluded from the average if the individual bar responses were very small (< 10 % of the maximum response).

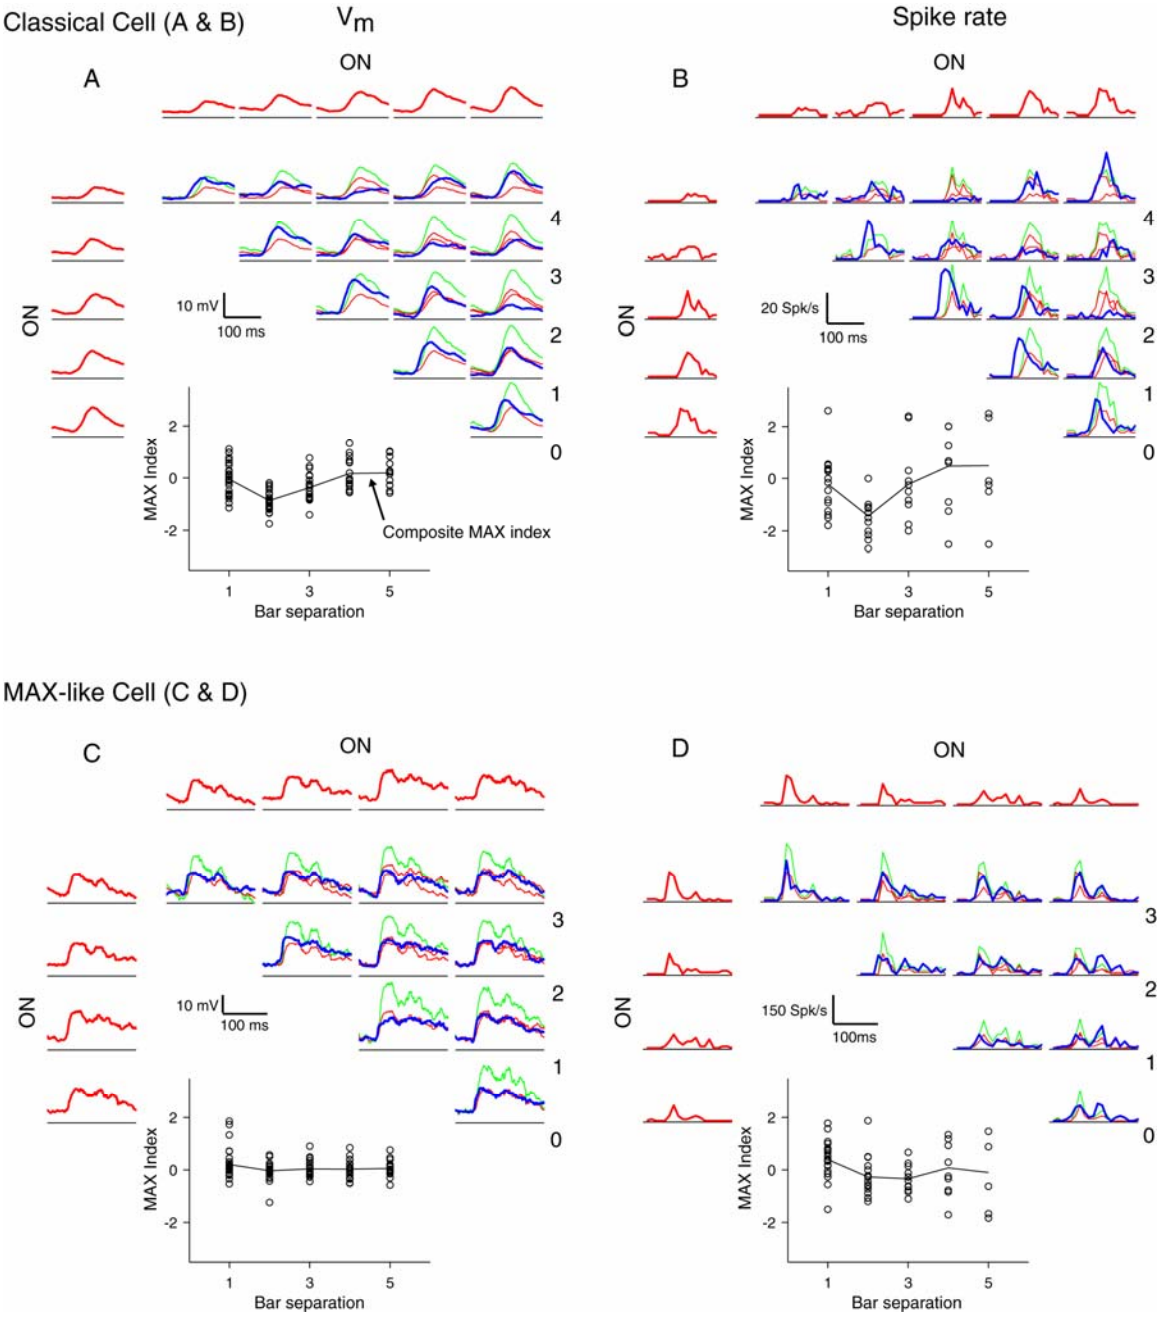
Classical Cell (A & B)
Vm
Spike rate
A
B
ON
ON
ON
10 mV
100 ms
20 Spk/s
100 ms
4
3
2
1
0
MAX Index
Composite MAX index
Bar separation
MAX-like Cell (C & D)
C
D
ON
ON
ON
10 mV
100 ms
150 Spk/s
100ms
MAX Index
Bar separation

Chapter IV

Inhibitory stabilization of the cortical network underlies visual surround suppression

Hirofumi Ozeki,[1*] Ian M. Finn,[1*] Evan S. Schaffer,[2]
Kenneth D. Miller,[2] and David Ferster[1†]

[1]Department of Neurobiology and Physiology, Northwestern University, Evanston, Illinois 60208, USA.

[2] Center for Theoretical Neuroscience and Center for Neurobiology and Behavior, Columbia University, College of Physicians and Surgeons,
New York, New York 10032, USA.

[*] These authors contributed equally to this work.

Acknowledgements: This work was supported by grants from the National Institute of Health to D.F. (R01 EY04726) and to K.M. (R01 EY11001). H.O. was partly supported by a JSPS Postdoctoral Fellowship. We thank Nicholas Priebe, Elizabeth Boudreau, Hiromichi Sato, Larry Abbott, Michael Eisele, and Taro Toyoizumi for useful comments.

Abstract

Lateral inhibition is thought to be a universal process that sharpens the tuning of sensory neurons and enhances apparent contrast (Hartline, 1949). According to this principle, a stimulus lateral to a neuron's preferred stimulus, either in physical space or in feature space, is expected to increase inhibition onto that neuron. In visual cortex, one form of lateral inhibition is surround suppression: stimuli in the receptive field surround suppress the response to stimuli in the receptive field center (Blakemore and Tobin, 1972; DeAngelis et al., 1994; Li and Li, 1994; Levitt and Lund, 1997; Anderson et al., 2001; Cavanaugh et al., 2002a, b; Ozeki et al., 2004). Contrary to expectation, we find that suppressive stimuli evoke only a transient increase in synaptic inhibition, after which both inhibition and excitation decrease to below their initial levels. These observations suggest that cortex forms an inhibition-stabilized network (Tsodyks et al., 1997), in which strong, recurrent excitatory connections are, on their own, unstable, and therefore require inhibition for stabilization. Lateral inhibition may thus be implemented through a reduction in local recurrent excitation rather than an increase in synaptic inhibition.

Experimental Procedures

Animal preparation. Acute experiments were performed on adult female cats (2-2.5 kg) (Boudreau and Ferster, 2005). Anesthesia was induced with ketamine (30 mg/kg i.m.) and acepromazine (0.3 mg/kg i.m.), and maintained with sodium thiopental (20 mg/kg i.v. initial; 1-2 mg/kg/hr i.v. maintenance). The head was fixed in a stereotaxic head holder. To minimize eye movements, paralysis was maintained with vecuronium bromide (0.2 mg/kg/hr i.v.), and the animal was artificially respirated at a rate to maintain end-tidal CO_2 at 3.7%. Body temperature, EEG, and ECG were continuously monitored. The nictitating membranes were retracted with phenylephrine hydrochloride, and the pupils were dilated with atropine sulfate. The corneas were protected by contact lenses with 4-mm artificial pupils. All procedures were approved by the Northwestern University Animal Care and Use Committee.

Visual stimulation. Monocular, circular sinusoidal drifting gratings were generated by a Macintosh computer using the Psychophysics Toolbox (Brainard, 1997) and Matlab (MathWorks), and presented on a ViewSonic CRT monitor (mean luminance, 20 cd/m^2; refresh rate, 100 Hz; resolution, 1024 × 768 pixels) placed 50 cm in front of the cat's eyes. Grating orientation, direction and spatial frequency were adjusted to be optimal (2 Hz temporal frequency; 64% contrast). Gratings were centered on a neuron's receptive field, as determined by presenting small gratings or flashing bars at different locations in the visual field. To measure size tuning of surround suppression in the steady state, we presented circular gratings of different diameters and annular gratings of 20-degree outer diameter with varying inner diameter. Stimuli, including a 0%-contrast blank, were presented in pseudorandom order for 4 sec. The classical

receptive field size was taken to be the largest stimulus size that evoked the strong firing response, but for which the surround alone (an annulus of the same inner diameter) did not evoke firing. To measure the contrast response function, the contrast of the center grating was varied from trial to trial, with and without surround stimulus present.

To measure orientation tuning of surround suppression in the steady state, a center grating of optimal size and preferred orientation was combined with 20-degree surround annuli of different orientations extending from the edge of the center grating (Fig. 4.01). To measure the transient response to the onset of surround stimuli (Fig. 4.05), the center and surround portions of the grating came on asynchronously (4 Hz temporal frequency): The center came on; 500 ms later, the surround came on; 500 ms later, both gratings were turned off. For simple cells, the starting spatial phase of the grating was varied from trial to trial, so that the surround grating came on during different phases of the cell's response. Stimuli were again interleaved in pseudorandom order, including center-alone stimulation and a blank.

Intracellular recordings. Intracellular whole-cell patch recordings were made from cells in area V1 of the cat visual cortex (400-850μm depth), at <5 degrees eccentricity. Borosilicate glass electrodes were filled with a K^+-gluconate solution including Ca^{2+} buffers, pH buffers, and cyclic nucleotides. In some experiments for conductance measurements (see below), K^+-gluconate was replaced by either Cs^+-gluconate or Cs^+-methanesulfonate to block K^+-channels, and QX-314 (7-8 mM, Sigma-Aldrich) was added to block Na^+-channels. Cells were classified as simple or complex on the basis of the ratio of modulation (F1) and mean (DC) components of firing rate response to an optimal grating (Skottun et al., 1991).

The amplitude of membrane potential and conductance responses (e.g., Fig. 4.01F-I) were derived from their peaks, which were measured as the F1 + DC components of the cycle-averaged responses. Spikes were removed from membrane potential traces before averaging by median filtering.

Conductance measurements. Stimulus-evoked changes in conductance were measured by injecting steady currents of 3 different amplitudes (Anderson et al., 2000b). Electrode resistance, measured by injection of brief current pulses (-0.1 nA; 250 ms), was compensated for off-line. At each point during the visual responses, membrane conductance was derived from the slope of the I-V curve. We derived the excitatory and inhibitory components of the visually evoked conductance from the membrane equation:

$$V_{visual}(t) = \left[g_e(t) \cdot V_e + g_i(t) \cdot V_i + g_{rest} \cdot V_{rest}\right] / g(t) \ ,$$

where $V_{visual}(t)$ is the response without injected current, V_{rest} is resting potential, $g(t)$ is the total conductance, and g_{rest} is the resting conductance. $g_e(t)$ and $g_i(t)$ are the visually evoked changes in excitatory and inhibitory conductances relative to the resting, unstimulated level, and can be either positive or negative. V_e and V_i are reversal potentials for excitatory and inhibitory conductances. V_i is assumed to arise from $GABA_A$- and $GABA_B$-mediated inhibition. For K^+-gluconate solution, V_e and V_i were 0 mV and -80 mV; for Cs^+-based solution (which blocks $GABA_B$ receptors), V_e and V_i were 0 mV and -70 mV.

Extracellular recordings from LGN. Single-unit recordings were made from the A layers of the LGN using lacquer-coated or glass-coated tungsten electrodes. Receptive fields of the recorded cells were located at a retinal eccentricity comparable to that of recorded V1 cells. Like cortical

intracellular recordings, we first used small drifting gratings at different positions to determine the precise location of the receptive field, and then tested the size tuning by varying stimulus inner and outer diameters of drifting gratings with optimal spatial frequency. Stimulus parameters were varied on a finer scale than those for cortical recordings because of the small diameter of geniculate receptive fields. Then, we tested the orientation tuning for the surround stimulus.

Electrical stimulation and latency measurements. Thalamocortical connectivity was determined from the latency of the response to stimulation of the LGN through a lacquer-coated tungsten electrode placed in layer A of the LGN in the retinotopic location matching that of the cortical recording electrode (500 μA electrode negative, 200 ms duration). Latencies <2.3 ms indicated monosynaptic excitation from the LGN; latencies >2.8 ms indicated exclusive polysynaptic input (Chung and Ferster, 1998).

Introduction / Results

The effects of surround stimulation on spike rate, membrane potential, and the underlying synaptic conductances are illustrated for a simple cell in the cat primary visual cortex in Figure 4.01. A 10-fold increase in stimulus size beyond the classical receptive field center reduced the spike response by 50% (Fig. 4.01A), but only when the surround stimulus was at the optimal orientation (DeAngelis et al., 1994; Li and Li, 1994; Levitt and Lund, 1997; Angelucci et al., 2002; Cavanaugh et al., 2002b; Bair et al., 2003; Ozeki et al., 2004; Webb et al., 2005). The membrane potential changes underlying the firing-rate responses were similarly selective for surround orientation, though smaller in amplitude (Fig. 4.01B, top traces). Thus, spike threshold significantly amplifies surround suppression, as it does for other types of visual selectivity such as end stopping (Anderson et al., 2001), direction selectivity (Jagadeesh et al., 1997), orientation selectivity (Carandini and Ferster, 2000; Volgushev et al., 2000) and cross-orientation suppression (Priebe and Ferster, 2006).

To measure the changes in synaptic input that underlie the membrane potential responses, we recorded the responses to visual stimulation while injecting different levels of steady current into the cell (Fig. 4.01B). Membrane conductance (Fig. 4.01C) was derived from the I-V relationship at each time point, and changes in excitatory and inhibitory conductance (Fig. 4.01, D and E) were derived from estimates of synaptic reversal potentials (Anderson et al., 2000b). The surround stimulus reduced both excitation and inhibition in an orientation selective manner, both for this cell (Fig. 4.01, D to F) and for a significant fraction of the recorded population (Fig. 4.01G and H). The magnitude of surround suppression in membrane potential is closely

correlated with the reduction in excitation, and not with an increase in inhibition (Fig. 4.01I).

Figure 4.01. **Steady-state measurements of surround suppression**.

A. Cycle-averaged firing-rate responses of a cortical simple cell to a blank stimulus, a stimulus covering the classical receptive field (2-degree diam), and the center stimulus plus a surround stimulus (20-degree diam.) at 3 different orientations (K^+-gluconate solution in the recording pipette). Maximal suppression occurred when the center and surround orientations matched.

B. Membrane potential responses (cycle-averaged after removing spikes) with 3 different levels of current injected into the cell. Gray traces are reconstructed from conductance measurements.

C-E. Stimulus-evoked changes in total membrane conductance, and excitatory and inhibitory conductance derived from the responses in (B). Horizontal lines represent the mean of the blank responses.

F. Peak firing rate, membrane potential, and changes in excitatory (red) and inhibitory (blue) conductances as a function of surround orientation relative to center orientation. Blank and center-only responses are shown as horizontal lines (shading = s.e.m.).

G. Changes in peak excitatory and inhibitory conductance evoked by the center-plus-surround stimulus plotted against changes evoked by the center stimulus alone.

H. Same as (G) with the surround stimulus at 90° to the preferred orientation.

I. Suppression ratio (response to center-plus-surround divided by the response to center alone) of membrane potential against that in excitation and inhibition.

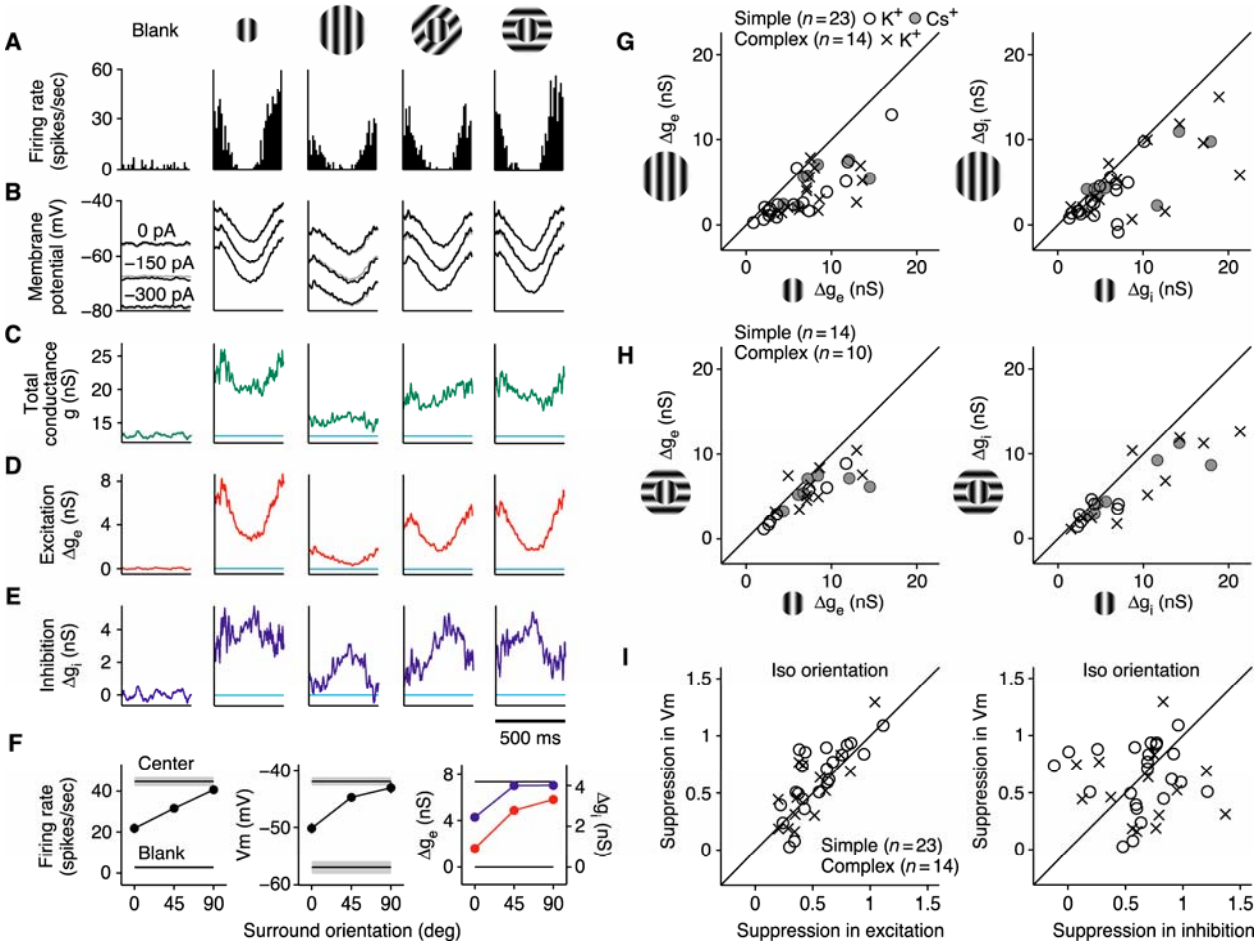
A
Blank
Firing rate (spikes/sec)
B
Membrane potential (mV)
0 pA
−150 pA
−300 pA
C
Total conductance g (nS)
D
Excitation Δg_e (nS)
E
Inhibition Δg_i (nS)
500 ms
F
Center
Blank
Firing rate (spikes/sec)
Vm (mV)
Δg_e (nS)
Surround orientation (deg)
G
Simple (n = 23) ○ K+ ● Cs+
Complex (n = 14) × K+
Δg_e (nS)
Δg_i (nS)
H
Simple (n = 14)
Complex (n = 10)
I
Iso orientation
Suppression in Vm
Simple (n = 23)
Complex (n = 14)
Suppression in excitation
Suppression in inhibition

One potential source for the surround-evoked reduction of excitation and inhibition is a withdrawal of excitatory input from the lateral geniculate nucleus (LGN). Three features of the cortical responses make this possibility seem unlikely. First, when tested with the same stimuli as those used in the cortex (Fig. 4.02A, B), average surround suppression in the spike responses of geniculate neurons (20%) was smaller than the suppression in membrane potential responses of cortical neurons (44%) and the suppression in excitatory and inhibitory conductances (54% and 48%). Second, in cortical neurons with strong surround suppression, the suppression was much more orientation selective than that observed in geniculate neurons (Fig. 4.02C, D). Third, in many simple cells, the withdrawal of excitation and inhibition (e.g., Fig. 4.01D, E) was accompanied by a downward shift in the trough of the sinusoidally modulated membrane potential (Fig. 4.03). Geniculate neurons fire little during this trough, so the underlying withdrawal of excitation is likely of cortical origin.

Figure 4.02. **Orientation tuning of surround suppression in geniculate relay cells and cfortical cells**.

A. Normalized and averaged tuning curves for firing rate in 18 LGN cells in response to center plus surround grating, plotted as a function of surround orientation (relative to center orientation). Horizontal line (at a value of 1) indicates the normalized response to center stimulus alone. The size of the center stimulus used was identical to those used for cortical cells: 2 degrees for LGN cells with receptive fields smaller than 1 degree, or 4 degrees for LGN cells with receptive fields larger than 1 degree. The suppression ratios for the iso-oriented grating and for the cross-oriented grating were (mean ± s.e.m.): 0.80 ± 0.04 and 0.89 ± 0.03.

B. Normalized and average tuning curves for simple cells: left, firing rate; center, membrane potential; right, changes in excitatory (red) and inhibitory (blue) conductances. Center stimulus was 2 degrees for 18 cells, or 4 degrees for 8 cells. Suppression ratio for iso-oriented surround: 0.29 ± 0.04 (firing rate), 0.56 ± 0.05 (membrane potential), 0.46 ± 0.05 (excitation), and 0.52 ± 0.13 (inhibition). Suppression ratio for the cross-oriented surround: 0.73 ± 0.06, 0.91 ± 0.03, 0.71 ± 0.04, and 0.75 ± 0.07. As reported previously in primates (Solomon et al., 2002), compared to cortical cells, LGN showed far weaker orientation tuning for surround orientation.

C. Orientation selectivity index (OSI) of surround suppression plotted against the strength of suppression for firing rate responses (F1 component) of geniculate neurons. The OSI was calculated as (% suppression by iso-oriented grating - % suppression by cross-oriented grating) / (% suppression by iso-oriented grating + % suppression by cross-oriented grating).

D. Orientation selectivity index of suppression plotted against the strength of suppression (peak response) for membrane potential (left) and excitatory and inhibitory conductances (right) of cortical cells.

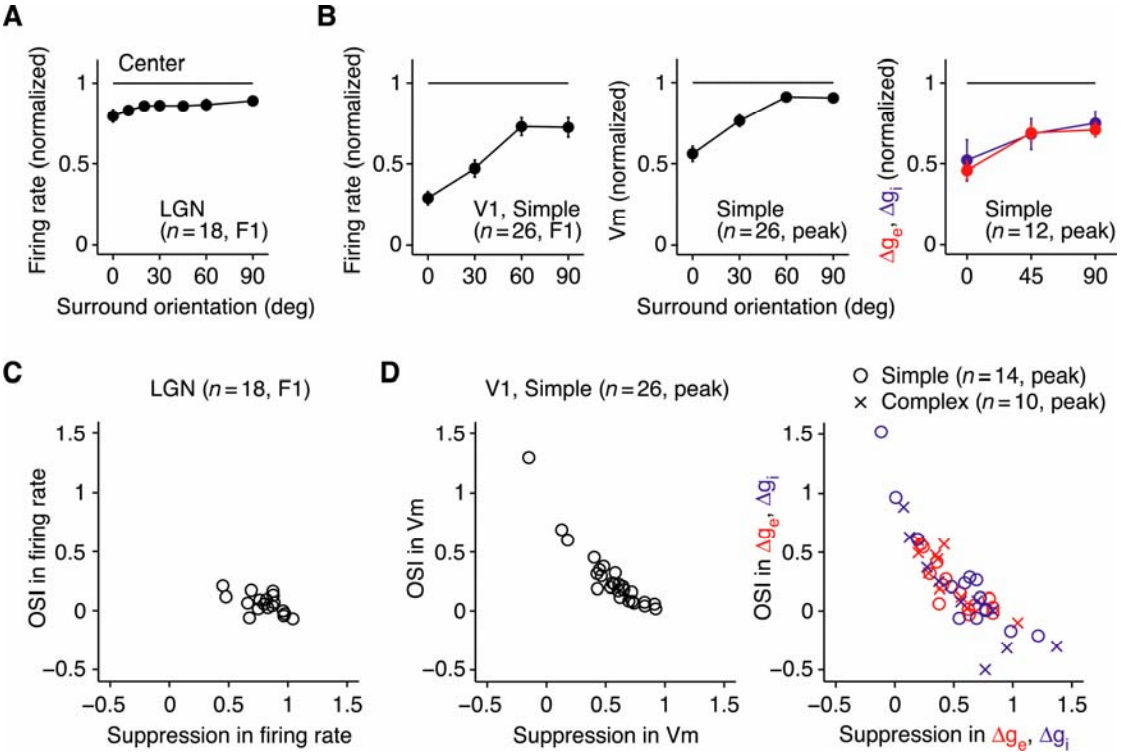
A
Center
Firing rate (normalized)
LGN
(n = 18, F1)
Surround orientation (deg)
B
Firing rate (normalized)
V1, Simple
(n = 26, F1)
Vm (normalized)
Simple
(n = 26, peak)
Surround orientation (deg)
Δge, Δgi (normalized)
Simple
(n = 12, peak)
C
LGN (n = 18, F1)
OSI in firing rate
Suppression in firing rate
D
V1, Simple (n = 26, peak)
OSI in Vm
Suppression in Vm
Simple (n = 14, peak)
Complex (n = 10, peak)
OSI in Δge, Δgi
Suppression in Δge, Δgi

Figure 4.03. **Evidence that reduction of cortical input underlies surround suppression**.

In many cells, the surround stimulus caused a downward shift of the entire membrane potential response (see *inset* at right).

A. The minimum of the membrane potential response to the iso-oriented grating plotted against that to the center grating (relative to resting potential) for 19 simple cells.

B. In the presence of the surround stimulus, the trough of the sinusoidally shaped response to the center grating was hyperpolarized by up to 5 mV (mean = 2.1 mV). This shift was observed when the trough of the center response was above the resting potential (A), suggesting that the downward shift presents evidence for the withdrawal of cortical excitation as a mechanism of surround suppression because geniculate relay cells should be contributing little excitation at the trough of the response of simple cells.

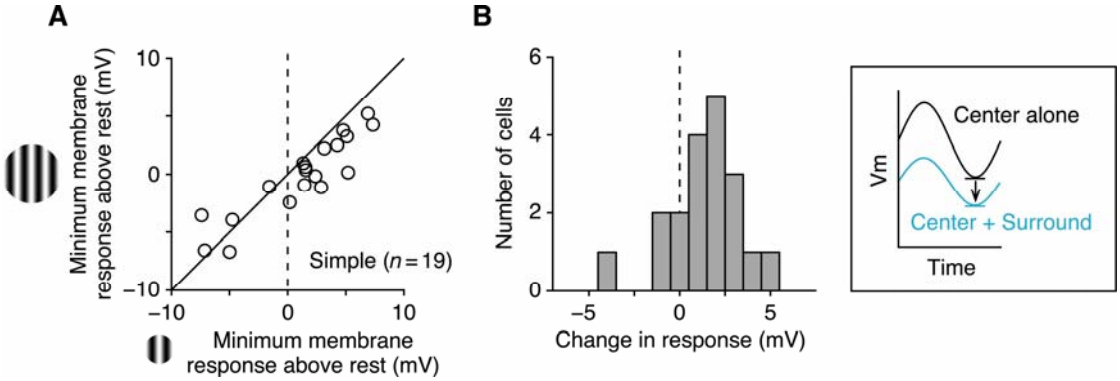
A
Minimum membrane response above rest (mV)
10
0
−10
−10
0
10
Simple (n = 19)
Minimum membrane response above rest (mV)
B
Number of cells
6
4
2
0
−5
0
5
Change in response (mV)
Center alone
Vm
Center + Surround
Time

A cortical mechanism that is consistent with these observations is an inhibition-stabilized network, or ISN (Tsodyks et al., 1997). In this model (Fig. 4.04A), two populations of cells in a cortical column, one excitatory (E) and one inhibitory (I), make recurrent and reciprocal connections. Each population also receives excitatory input from two external sources: 1) a feed-forward pathway that defines each cell's receptive field center, originating from geniculate relay neurons or from other neurons within the column; 2) a laterally-projecting surround pathway that synapses most strongly on the inhibitory population, and that originates from cortical neurons outside the column – either in V1 (Buzas et al., 2001; Angelucci et al., 2002) or in extrastriate areas (Angelucci et al., 2002; Bair et al., 2003).

To be an ISN, a network must satisfy two properties. First, recurrent connections among the excitatory neurons must be strong enough to be unstable. That is, without inhibition, the feedback excitation will drive the excitatory neurons either to saturation or to very low firing rates, making it impossible for the network to respond to stimuli with sustained, moderate firing rates. Second, feedback inhibition must be strong enough to stabilize the network, so that in the presence of inhibition, moderate rates become possible.

When these two constraints are met, the network responds as shown in Figure 4.04B. At t = -50 ms, the network is in the steady state evoked by the center stimulus alone, with both the E and I populations active. At t = 0 ms, the surround stimulus is added at the preferred orientation, activating the surround pathway, which evokes an increase in the external excitation to inhibitory cells (Fig. 4.04Ba). The firing rates of inhibitory cells increase, increasing inhibition onto the excitatory population (Fig. 4.04Bb). This inhibition decreases the firing rates of the excitatory

population, thereby reducing recurrent excitation onto excitatory cells, and further decreasing their activity (Fig. 4.04Bc). Decreased firing in the excitatory cells decreases excitation to the inhibitory cells (Fig. 4.04Bc). By virtue of the intrinsic properties of the ISN (unstable excitatory recurrence stabilized by inhibition), this reduction in excitation to the inhibitory cells becomes larger than the initial increase in excitation from the surround pathway. In the steady state, then, activity of the both E and I populations are reduced (Fig. 4.04Bd).

The network behavior can be understood from a phase plane analysis, in which the state of the network is represented by plotting the average firing rate of the I and E populations, r_I vs. r_E (Fig. 4.04C, top). The inhibitory nullcline (blue) shows the activity in the I population, r_I, that results when r_E is clamped at different values. The inhibitory subnetwork is stable in that after a perturbation away from the nullcline, r_I moves vertically back towards the nullcline along the vertical arrows. The excitatory nullcline (red) shows the activity in the E population, r_E, that results when r_I is clamped at different values. The excitatory subnetwork alone is stable around the portions of the nullcline with negative slope (horizontal black arrows). In the region of positive slope, however, when the network is perturbed even slightly away from the nullcline, excitatory feedback drives the network activity even further away, toward very high or low levels (gray arrows). Thus, in the region of positive slope, the first requirement of an ISN is met, i.e., that the excitatory subnetwork, on its own, is unstable.

When inhibition and excitation are both free to vary, the network's steady state, or fixed point, lies at the intersection of the two nullclines (black point), the point at which, in the absence of perturbations, network activity will not change over time. As dictated by the second requirement

for an ISN (network stability), when the network is transiently perturbed away from the fixed point, it settles back to the fixed point along trajectories such as those shown in green. These trajectories are determined by the network trends in the vicinity of the nullclines (black and gray arrows). Note that a stable fixed point can only emerge when the inhibitory nullcline has a larger slope than the excitatory nullcline at the point of intersection.

When the surround pathway is activated by increasing the stimulus size, the inhibitory nullcline moves upwards (Fig. 4.04C, bottom, dashed line): for any fixed value of r_E, there is now additional external excitation onto the I population, making the value of r_I that results larger than before. Because the inhibitory nullcline has a greater positive slope than the excitatory nullcline, raising the inhibitory nullcline shifts the fixed point downward and to the left. Thus, at this new fixed point, activity in both the E and I populations decreases, as we have observed in cortical cells.

Figure 4.04. **An inhibition-stabilized network (ISN) as applied to surround suppression in the visual cortex**.

A. Two populations of cells, one excitatory (E) and one inhibitory (I), synapse on themselves and on each other. Each population receives excitatory, feed-forward inputs from a center pathway that defines the receptive field centers of the neurons. Lateral excitatory input from a surround pathway projects more strongly to inhibitory than to excitatory cells.

B. The sequence of events that follow when a surround stimulus is added to a pre-existing center stimulus. For simplicity, the surround pathway is assumed to stimulate only the inhibitory population, and the center pathway, though active, is not shown. After a transient increase in the activity of the inhibitory cells (b, light blue), activity in both excitatory and inhibitory cells decreases relative to the level evoked by center stimulus alone (d, green). The temporal sequence of changes in the activities of E and I populations are shown in the bottom panel.

C. Phase-plane diagrams of the network activity (see text).

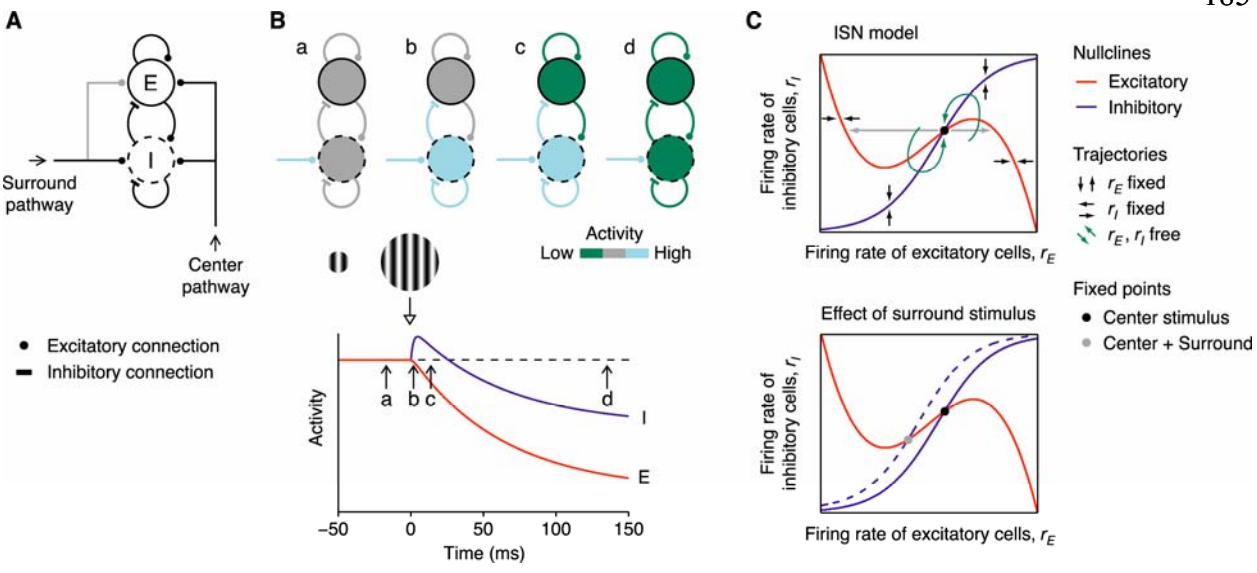
A
E
I
Surround pathway
Center pathway
Excitatory connection
Inhibitory connection
B
a
b
c
d
Activity
Low
High
Activity
a b c
d
I
E
−50
0
50
100
150
Time (ms)
C
ISN model
Firing rate of inhibitory cells, r_I
Firing rate of excitatory cells, r_E
Nullclines
Excitatory
Inhibitory
Trajectories
r_E fixed
r_I fixed
r_E, r_I free
Fixed points
Center stimulus
Center + Surround
Effect of surround stimulus
Firing rate of inhibitory cells, r_I
Firing rate of excitatory cells, r_E

Can the surround-evoked decrease in excitation and inhibition arise from an excitatory subnetwork that is stable on its own? In other words, is an ISN required to explain the experimental data? In one non-ISN scenario, surround stimulation could reduce external excitation coming to the local network either from the center or surround pathways. As noted above, however, surround suppression in geniculate neurons is too small and too weakly orientation selective to account for the cortical data. That the surround stimulus reduces activity in the surround pathway seems unlikely because it functions as a center stimulus for the cells in the pathway, and should therefore increase their activity. The remaining possibilities – including the classical model of lateral inhibition, where the surround stimulus evokes external inhibition directly onto the excitatory cells – all require the network to operate in the inhibition-stabilized regime (see supporting online text).

In addition to a steady-state reduction in excitation and inhibition, the ISN model makes three predictions. First, prior to reaching its steady state, the network should show a transient increase in inhibition (Fig. 4.04Bb). To test this prediction, we presented a stimulus in the receptive field center and then abruptly increased the stimulus size to cover the surround. A transient increase in inhibitory conductance was observed in 6 cells. In the complex cells of Figure 4.05A and B, the transient was only evoked by a surround stimulus of the preferred orientation (* blue traces, 2nd and 3rd columns). Here, the surround stimulus transiently increased excitation as well as inhibition. For the simple cell in Figure 4.05C, we varied the starting spatial phase of the drifting grating. No matter whether the increase occurred during the depolarizing phase of the response (2nd and 3rd columns) or during the hyperpolarizing phase (4th and 5th columns), a brief (50-100

ms) increase in inhibitory conductance was observed before both conductances decreased to steady state levels.

Figure 4.05. **Transient increase in inhibitory conductance evoked by addition of a surround stimulus**.

A and **B**. The responses of 2 complex cells to a blank stimulus (1st column) and to a sudden onset of the surround stimulus (arrows) with either the iso-orientation (2nd column) or cross-orientation (3rd column). The initial stimulus covered the receptive field center and began 250 ms prior to the start of the traces. Black, membrane potential recorded with different currents injected; red and blue, derived changes in excitatory and inhibitory conductance. A transient increase in conductance is indicated by asterisks.

C. A simple cell tested with 4 iso-oriented surround stimuli. For each of the 4 stimuli, the starting phase of the grating is shifted by 90° so that the increase in stimulus size occurs during a different phase of the response. These cells were recorded with K^+-gluconate solution in the recording pipette.

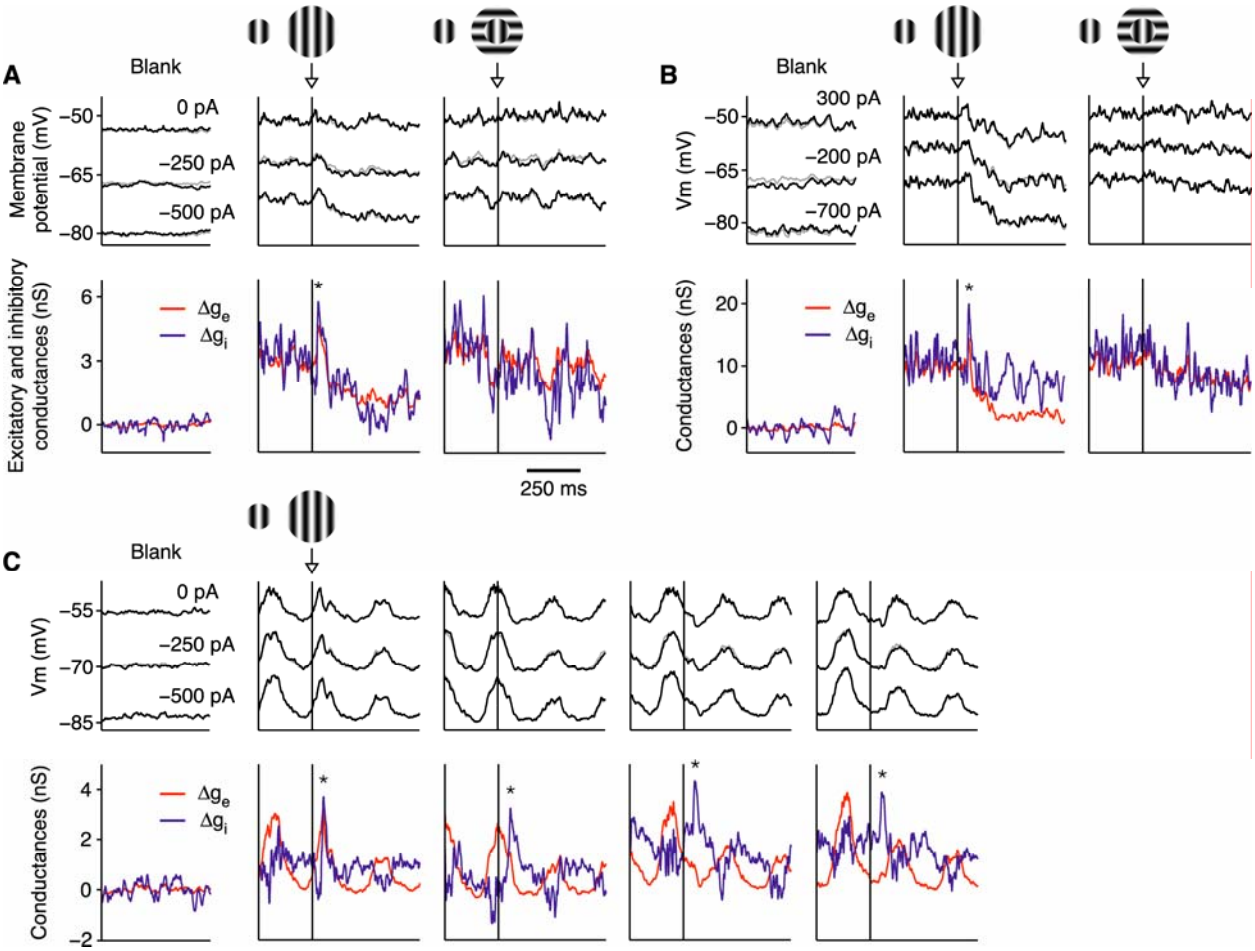
A
Blank
0 pA
−250 pA
−500 pA
Membrane potential (mV)
−50
−65
−80
Excitatory and inhibitory conductances (nS)
6
3
0
Δg_e
Δg_i
*
250 ms
B
Blank
300 pA
−200 pA
−700 pA
Vm (mV)
−50
−65
−80
Conductances (nS)
20
10
0
Δg_e
Δg_i
*
C
Blank
0 pA
−250 pA
−500 pA
Vm (mV)
−55
−70
−85
Conductances (nS)
4
2
0
−2
Δg_e
Δg_i
*

The second prediction of the ISN model is that neurons that receive all of their input from geniculate relay cells should show little surround suppression. The orientation-selective component of surround suppression arises from a withdrawal of intracortical excitation, and so in cells without significant cortical input, what suppression there is should be similar to the small, weakly orientation-selective suppression observed in geniculate cells. By comparison, in a classical model of lateral inhibition (Hartline, 1949; Hubel and Wiesel, 1965), suppression arises from an increase in synaptic inhibition; any neuron that receives this inhibition is therefore subject to surround suppression, regardless of the source of its excitation.

In our experiments, we have not made a direct measurement of the proportion of input each cell receives from the LGN. We have shown previously, however, that this proportion correlates closely with the ratio of the mean depolarization evoked by center gratings of the null and preferred orientations, $DC_{null}/DC_{preferred}$ (Finn et al., 2007). Cells with high ratios receive most of their excitatory input from the LGN; cells with low ratios receive most of their excitatory input from other cortical cells. To gauge the orientation selectivity of surround suppression, we calculated an orientation selectivity index (OSI) as % suppression of the peak response evoked by an iso-oriented grating minus the % suppression evoked by a cross-oriented grating, divided by sum of the two. Indices of 0 and 1 correspond to orientation non-selective suppression and completely selective suppression. As predicted by the ISN model, cells with higher $DC_{null}/DC_{preferred}$ ratios (> 0.7), and therefore a higher proportion of LGN input, rarely showed strongly orientation-selective suppression (Fig. 4.06A).

The strength of suppression of cortical cells was also correlated with the latency of response to electrical stimulation of the LGN (Fig. 4.06B). Cells with latency below 2.3 ms, which necessarily receive some proportion of their input directly from the LGN (Chung and Ferster, 1998), exhibited only small suppression (Walker et al., 2000; Akasaki et al., 2002). The suppression was stronger for those cells with latencies longer than 2.8 ms, which receive all of their input from other cortical cells. The amount of orientation selective surround suppression was also correlated with latency (Fig. 4.06C): Cells with short latencies had less orientation-selective surround suppression than those with longer latencies. These correlations are not expected to be as strong as that in Figure 4.06A, since latency only indicates whether or not a cell receives input from the LGN, and not the proportion of geniculate input.

The third prediction of the ISN model is that surround suppression should be little affected by blocking synaptic inhibition in a small number of cells. Doing so will have a minimal effect on the overall network behavior, and therefore the surround stimulus will still reduce the net excitation received by neurons in the region of the blockade. In agreement with this prediction, applying bicuculline to a small area surrounding a recorded cell has little effect on surround suppression (Ozeki et al., 2004). In contrast, when applied to the whole network relatively small attenuation of inhibition can yield sudden instability (Chagnac-Amitai and Connors, 1989).

Because it represents only the average behavior of the E and I populations, the model in Figure 4.04 cannot account for one important aspect of the data: the large cell-to-cell variability in surround suppression of excitation and inhibition (Fig. 4.01G-H). This diversity can be captured, however in a multi-neuron ISN model. In this model, variability in surround-induced suppression

of synaptic conductances (Fig. 4.06D) is comparable to that seen in the data (Fig. 4.06E), except for the relative lack of cells below the diagonal in Figure 4.06E. Such cells, with larger suppression in inhibition than in excitation, would likely show surround facilitation in membrane potential. Although we found surround facilitation in some cortical cells (a monotonic increase in response with an increase in stimulus size) (Li and Li, 1994; Walker et al., 2000), we did not record their synaptic currents, so such cells are missing from our data set. Figure 4.06D, however, suggests that cells demonstrating surround suppression or facilitation represent different segments of the continuum of surround effects produced by the ISN circuit, rather than fundamentally different populations.

Figure 4.06. **Comparison of population data with predictions of the ISN model**.

A. Orientation selectivity index (OSI) of surround suppression for simple cells plotted against the ratio of the mean potential evoked by center gratings of the null and preferred orientations. This ratio is highly correlated with the proportion of input each cell receives from the LGN, with higher ratios indicating a larger proportion of LGN input (Finn et al., 2007). OSI = (% suppression by iso-oriented surround - % suppression by cross oriented surround) / (% suppression by iso-oriented surround + % suppression by cross-oriented surround).

B and **C**. Suppression ratio (B) and orientation selectivity index (C) of surround suppression plotted against the latency of response to electrical stimulation of the LGN (B, $r = -0.52$, $P < 0.01$; C, $r = 0.55$, $P < 0.02$). The arrows show the mean OSI (A, C) and mean suppression ratio (B) for 18 LGN cells.

D. The relationship between the surround-evoked reductions of excitation and inhibition in a linear, multi-neuron model (supporting online text).

E. Corresponding data from intracellular records (from Fig. 4.01G). Colors indicate the strength of suppression in peak membrane potential.

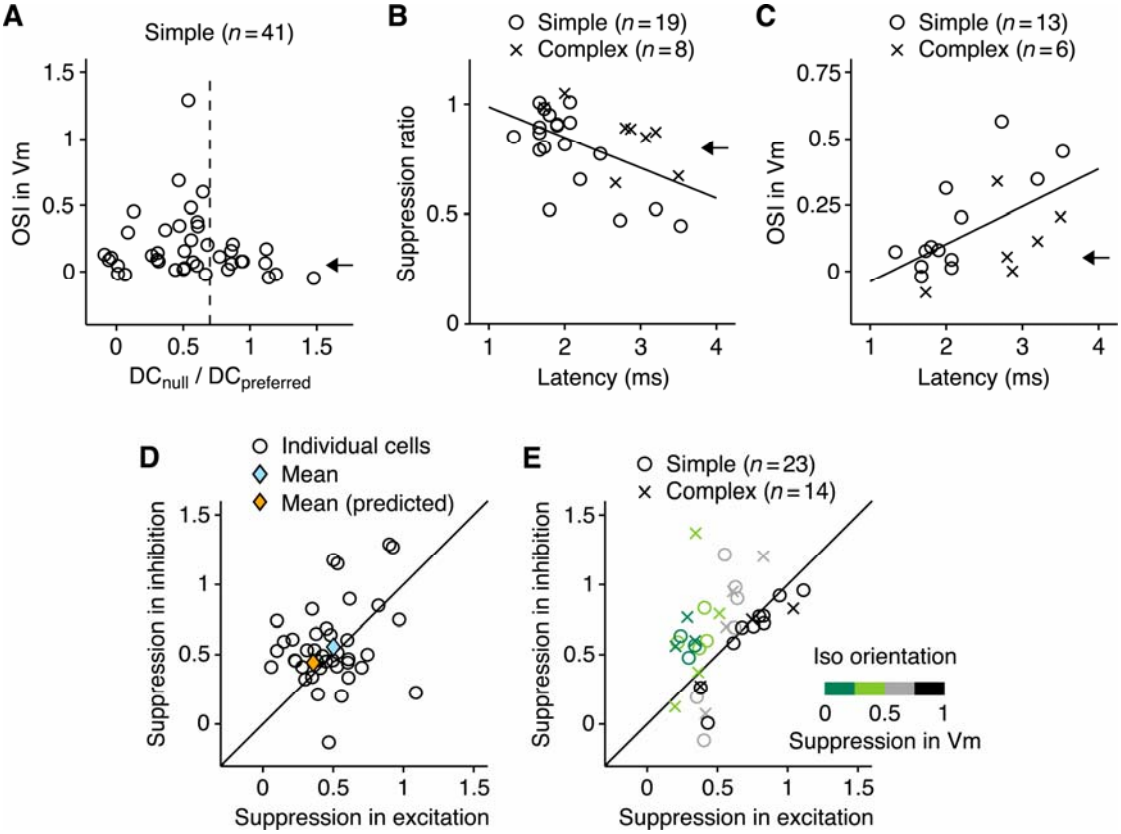
A
Simple (n = 41)
OSI in Vm
DC_null / DC_preferred
B
Simple (n = 19)
Complex (n = 8)
Suppression ratio
Latency (ms)
C
Simple (n = 13)
Complex (n = 6)
OSI in Vm
Latency (ms)
D
Individual cells
Mean
Mean (predicted)
Suppression in inhibition
Suppression in excitation
E
Simple (n = 23)
Complex (n = 14)
Iso orientation
Suppression in Vm
Suppression in inhibition
Suppression in excitation

Discussion

We note that our results are similar in some respects to those of Anderson et al. (2001) who showed a reduction in excitatory and inhibitory conductances for length tuning. The source of this reduction is controversial, however, because geniculate cells also exhibit strong length tuning for the stimuli used in their experiments (Cleland et al., 1983a; Nolt et al., 2004). In contrast, only a small proportion of orientation-dependent surround suppression can be explained by the properties of the LGN.

Cerebral cortical circuits are characterized by massive recurrent excitation, which presumably is critical to the computations they perform. To reach threshold, a neuron requires synchronous excitatory activity in fewer than 100 of its synaptic inputs (Bruno and Sakmann, 2006; Waters and Helmchen, 2006), and yet each cell receives synapses from thousands of other excitatory cells (Elston, 2003). When excitatory cells are active, this strong recurrent excitation would seem to require strong inhibition to balance it, and such a balance can also account for the high levels of variability observed in cortical activity (van Vreeswijk and Sompolinsky, 1998; Haider et al., 2006; Haider et al., 2007). Computational models have recognized the difficulty of constructing networks with strong recurrence that will operate at moderate firing rates, and have suggested that stabilization of an unstable excitatory subnetwork by feedback inhibition may provide a robust solution (Latham and Nirenberg, 2004). There are few areas of the brain, however, where the function of synaptic inhibition in local circuits has been fully characterized. In a study of the hippocampus, Tsodyks et al. (Tsodyks et al., 1997) first defined the ISN regime, explored its behavior theoretically, and provided experimental evidence that the local circuit might operate in

this regime. We now provide evidence that at least one area of the neocortex, V1, operates in this same regime.

The properties of surround suppression that we have observed fit well with predictions of the ISN model, and are difficult to explain with any other plausible circuit architecture. We also suggest that inhibitory stabilization gives rise to the classical property of lateral inhibition, but does so in an unexpected and paradoxical way. Whether this mechanism lies at the root of other receptive field properties in the visual cortex and elsewhere remains to be determined.

Chapter V

Conclusions

The work presented in this thesis was aimed at discerning the contributions made by feed-forward and recurrent processing to suppressive response properties in cat visual cortex. An examination of the detailed behavior underlying contrast invariant orientation tuning in simple cells confirmed key aspects of the Hubel-Wiesel feed-forward model for the genesis of orientation selectivity, and argued against a role for lateral inhibition. Studying the spatial integration of complex cells revealed a range of behavior consistent with hierarchical energy models. Finally, a careful investigation of surround suppression in both simple and complex cells suggested that, while inhibition may not underlie response specificity in V1, it is critical for maintaining stability in the V1 network.

Contrast-invariant orientation tuning

The achievement of contrast-invariant orientation tuning (Sclar and Freeman, 1982; Skottun et al., 1987; Alitto and Usrey, 2004) is considered an important stage in visual information processing, and a computation that requires substantial contributions from the intracortical network. Lateral inhibition is one type of intracortical contribution that is commonly suggested, as it would seem ideally suited to counteract the purely excitatory, contrast-dependent feed-forward input that leads to a breakdown of contrast-invariance in Hubel-Wiesel type models (Sompolinsky et al., 1990; Heeger, 1992; Somers et al., 1995; Troyer et al., 1998; Hirsch et al.,

2003; Lauritzen and Miller, 2003; McLaughlin et al., 2003). Although we did not directly measure inhibitory and excitatory influences on simple cells in Chapter II, our results strongly suggest that lateral inhibition does not contribute, as has traditionally been thought, to the computation of contrast-invariant orientation tuning.

One of the enduring challenges to Hubel and Wiesel's feed-forward model, and a strong argument for the role of lateral inhibition, has been the lack of evidence, either direct or indirect, for an orientation-independent rise in the mean membrane potential of simple cells when responding to drifting grating stimuli. Measured extracellularly, the bandwidth of orientation tuning does not change as a function of contrast (Sclar and Freeman, 1982; Skottun et al., 1987; Anderson et al., 2000c; Alitto and Usrey, 2004). On the face of it, this observation would appear to signal a substantial failure of the feed-forward model, which predicts contrast-dependent changes in orientation tuning bandwidth due to contrast-dependent changes in mean membrane potential at all orientations. Intracellular measurements have also failed to reflect orientation-independent changes in mean membrane potential, as the majority of published simple cell data shows contrast dependent depolarizations in response to preferred stimuli, but not to orthogonal stimuli (Anderson et al., 2000c; Lampl et al., 2001; Monier et al., 2003; Marino et al., 2005), .

In Chapter II we reported observing a subset of simple cells in V1 that demonstrate precisely the behavior expected from a feed-forward model. These cells respond to orthogonally oriented stimuli in a constant, unmodulated fashion, and their average depolarizations from rest are nearly identical across orientation. The existence of such simple cells argues very strongly that orientation selectivity, at least for some V1 neurons, arises from the convergence of pre-synaptic

LGN cells with spatially aligned receptive fields, as predicted by Hubel and Wiesel (Hubel and Wiesel, 1962).

The data in Chapter II also indicate that not all simple cells derive their orientation selectivity from the thalamus. Instead, some simple cells rely on intracortical input to set up their receptive field properties. By virtue of their origin, these disynaptic or second order simple cells demonstrate smaller membrane potential responses to orthogonally oriented stimuli than their network predecessors, which do depolarize in response to orthogonal stimuli but generally do not spike.

Extracellular contrast invariant orientation tuning in disynaptic simple cells arises in a manner that has previously been well described (Anderson et al., 2000c). The explanation hinges on a crucial observation that the membrane potential to spike-rate transformation more closely resembles a smooth power-law function than a threshold-linear function (Hansel and van Vreeswijk, 2002; Miller and Troyer, 2002). When operating on membrane potential tuning curves that are themselves nearly invariant and, importantly, Gaussian shaped, the power-law power law transformation significantly enhances (and/or preserves) invariance, and generally narrows tuning width, for the resulting spike-rate tuning curves.

Given that substantial violations of extracellular contrast invariant orientation tuning have not yet been observed, our data in Chapter II raises the question of how monosynaptic simple cells, which intracellularly are not invariant, could become invariant extracellularly. The explanation lies in part with an observation we made that response variability, on a trial-to-trial basis, is

indirectly correlated with contrast. Multiple cycles of a drifting grating presented at high contrast evoke stereotyped membrane potential changes with similar peak amplitudes. Multiple cycles of the same grating at lower contrast evoke membrane potential responses with peak amplitudes that vary considerably. As a consequence, orientations far from the preferred at high contrast can produce fewer spikes than orientations closer to the preferred at low contrast, even though the membrane potential responses evoked in the former case might, on average, be larger than in the latter. Decreases in membrane potential variability may thus underlie classic observations that at high contrast, non-preferred stimuli appear to suppress firing activity relative to normal or elevated background rates, and do so more than low-contrast stimuli of the same orientation (Morrone et al., 1982; Sclar and Freeman, 1982; Ramoa et al., 1986). Together, these observations imply that intracellular tuning curves, constructed as they are from average peak responses, do not contain all the necessary information to predict the properties of extracellular tuning curves. Instead, the variability of peak membrane potential responses must be taken into account as well.

The data in Chapter II present many avenues for future studies. In particular, figure 2.09 raises questions about the origin of large contrast-dependent differences in membrane potential variability, indicating that they may arise from feed-forward LGN input. Similarly, figure 2.04 is suggestive, though far from conclusive, regarding a laminar segregation for simple cells with large DC_N/DC_P ratios and simple cells with small DC_N/DC_P ratios. Determining conclusively whether simple cells that show large depolarizations to orthogonal stimuli are clustered in layer IV would further extend the story for how orientation tuning arises in primary visual cortex.

MAX behavior in complex cells

Intracellular measurements made by Lampl et al. (2004) offered a view of spatial integration in complex cells that seemed to be at odds both with previous findings (Movshon et al., 1978b) and with Hubel and Wiesel's hierarchical model for information processing in V1. In their study, Lampl et al. (2004) observed that some complex cells respond to pairs of stimuli within their receptive fields as if only one of the stimuli were present, in particular the stimulus that alone evokes a larger response (MAX behavior). The authors proposed that MAX behavior could be explained either by intracortical mechanisms similar to those for gain control (Carandini and Heeger, 1994), or by a shunting mechanism whereby the total conductance evoked by the stronger stimulus significantly suppresses the impact of the weaker stimulus.

These explanations depend heavily on contributions from inhibition, in the first case to scale back the neuron's response in proportion to the total excitatory drive, and in the second to increase the neuron's conductance enough to shunt excitatory input provided by the weaker stimulus. In addition, both of these possibilities implicitly assume that the two stimuli always activate independent populations of neurons pre-synaptic to the complex cell. Thus MAX behavior in these scenarios reflects the operation that the complex cell uses to pool its spatially segregated inputs.

While it is likely that some form of response normalization occurs on input to complex cells, mediated by the intracortical network or by more immediate means, such as by synaptic depression or driving force changes, our data support a more traditional feed-forward origin for the MAX behavior reported by Lampl et al. (2004). Similar to Movshon et al. (1978b), many of

the neurons we recorded from demonstrated distance and polarity dependent interactions in response to pairs of bar stimuli, which could be explained if the stimuli interacted with each other in the receptive fields of pre-synaptic simple cells. Further, evidence suggests that complex cell receptive fields are at most only slightly larger than simple cell receptive fields (Gilbert, 1977; Livingstone and Conway, 2003), which argues against the idea that pairs of bars could consistently fall in the receptive fields of spatially disjoint simple cells that project to a single complex cell. Stimulating different pre-synaptic subunits with paired stimuli is potentially plausible if complex cell receptive fields derive from convergent LGN input, but there is no strong evidence that this is the case.

If the majority of paired bar stimuli do interact at the level of pre-synaptic simple cells, how can we then explain the emergence of MAX behavior, which in this case reflects the actual membrane potential response of a complex cell to pairs of bars, and not the pooling operation the complex cell performs on its inputs? We suggested and tested a model that involves the spatial frequency tuning of simple cell inputs to complex cells. The model grew out of an observation we made that the spatial frequency tuning bandwidth of highly MAX-like complex cells tended to be larger than the bandwidth for non-MAX-like cells (Classical cells).

We incorporated this observation into a standard energy model (Adelson and Bergen, 1985) for complex cell formation, which assumes that 2 (or 4, depending upon implementation) simple cells, identical in all respects except for their relative spatial phases, provide input to a single complex cell. Complex cells with (relatively) narrow spatial frequency tuning curves were created using the standard prescription, while complex cells with large spatial frequency tuning

bandwidths were mimicked by combining input from multiple pairs of simple cells with different preferred spatial frequencies. When tested with many thousands of pairs of bar stimuli, the two types of model complex cells qualitatively reproduced the average behavior of MAX-like and Classical cells observed in our experimental population.

Our results suggest that MAX-like behavior in response to pairs of bar stimuli might emerge from the particular combination of pre-synaptic subunits that converge on a complex cell. Physiologically it seems reasonable to assume that there is some amount of variability in the properties of simple cells that provide input to a complex cell, a hypothesis that to date has not been explored in detail. Investigating behaviors that might arise in models of complex cells using a greater range of variability in simple cell input is one possible avenue for future work.

A clear question raised by this work is whether the differences in spatial frequency bandwidth between MAX-like and Classical complex cells do in fact arise from differences in the spatial frequency tuning of the inputs to these cells. This is a difficult question to address experimentally, but the answer would be highly illuminating. Questions more amenable to experiment include determining whether MAX-like and Classical cells segregate according to cortical layer, or according to other, previously defined classes of complex cells (Palmer and Rosenquist, 1974; Gilbert, 1977; Henry et al., 1978). Another important question not addressed by our study concerns the behavior of simple cells in response to paired bar stimuli. A thorough investigation of intracellular, and extracellular, spatial integration in simple cell receptive fields could inform future models of complex cell formation, and perhaps shed additional light on the origins of MAX-like and Classical response patterns in complex cells.

Lateral inhibition in primary visual cortex

Chapter II of this thesis presents strong evidence that intracortical lateral inhibition is not necessary for the genesis of contrast invariant orientation tuning. Chapter III makes the case for a purely excitatory, feed-forward origin of MAX-like behavior in complex cells. These results, when considered in light of previous work suggesting that non-linearities of the feed-forward, excitatory pathway give rise to cross-orientation inhibition (Priebe and Ferster, 2006) and sharp orientation tuning (Lampl et al., 2001), would seem to leave little room for inhibition to shape response properties in V1. It is clear from many intracellular studies (Ferster, 1986; Anderson et al., 2000b; Anderson et al., 2001; Hirsch et al., 2003; Monier et al., 2003; Marino et al., 2005; Priebe and Ferster, 2005, 2006; Tucker and Fitzpatrick, 2006), however, that strong inhibitory conductances are evoked by visual stimulation. The question thus becomes: What function does inhibition subserve in V1?

Chapter IV of this thesis proposes that a major role for inhibition is to stabilize the excitatory subnetwork in V1. This idea was suggested by intracellular studies of visual surround suppression, which is functionally a form of lateral inhibition. As with retinal lateral inhibition, stimuli in one location suppress the response to stimuli in adjacent locations, presumably as a mechanism for enhancing the discriminability of adjacent visual stimuli. Our data, however, are not consistent with a simple increase in inhibition between neurons with adjacent receptive fields, which is the mechanism for lateral inhibition observed in the retina and often proposed for the cortex.

Our ISN model of visual cortex is highly similar to one proposed by Tsodyks et al. (1997) for hippocampal circuits. Both models give rise to paradoxical behavior whereby increasing excitatory drive to inhibitory neurons through an external pathway can actually lead to a reduction in the inhibitory neurons' activity, mediated by a decrease in excitatory drive through an internal pathway. That both the hippocampus and an area of neo-cortex implement inhibitory stabilization argues strongly for the critical importance of this mechanism.

Why might V1 or other areas of neo-cortex employ inhibitory stabilization? In general, cortical computations require neurons to communicate with each other (recurrently) by changing each other's activity levels. If recurrent excitation is too strong, however, the network will suffer an excess of positive feedback and become unstable. One mechanism to prevent excessive positive feedback is to keep recurrent excitation weak. Given the massive number of excitatory inputs received by each cortical cell, however, this mechanism implies that each connection would be so weak as to severely restrict the computations accessible to cortex. An alternative is to allow powerful recurrent excitation in the network, but to modulate excitatory activity with dynamic feedback inhibition. This solution allows computations to be performed that require a non-trivial amount of excitation, while keeping the network activity at low or moderate levels relative to biophysically maximal firing rates (Latham et al., 2000; Latham and Nirenberg, 2004).

We observed in Figure 4.06 that putative layer IV cells that receive large amounts of feed-forward, geniculate input do not generally demonstrate surround suppression. These cells could either be considered a subpopulation of the ISN in which excitatory and inhibitory conductances do not change in response to surround stimulation, or they may be thought of as outside the ISN

altogether. This latter perspective would suggest that non-surround suppressed cells might provide "external" input to other cortical cells, whose behavior would then be subject to the ISN. In visual cortex then, complex image processing computations might take place after the emergence of orientation tuning. Additionally, according to Hubel and Wiesel's hierarchical model some surround suppression in V1 should be due to network inheritance. Complex cells receiving input from surround suppressed simple cells would, for example, appear surround suppressed themselves, with both excitatory and inhibitory conductances decreasing in the steady-state.

Although we did not identify the cells in our study as being excitatory or inhibitory, it is likely we recorded from predominantly, if not exclusively, excitatory cells. One possibility for extending the work performed here is to target inhibitory interneurons (Tsodyks et al. 1997) specifically and record their activity, either intracellularly or extracellularly, in response to surround stimulation. Other avenues for future investigation include exploring in more detail the nature of the transient conductance evoked by surround stimulation. Determining the precise surround strength necessary to evoke a transient conductance increase could provide further insight into the role of inhibitory stabilization. Finally, extending the ISN model to account for known features of the V1 network, including highly thalamorecipient layer IV simple cells and complex cells that surround suppress due to inheritance, might generate new predictions that could be tested experimentally.

In summary, we have elucidated a number of mechanisms that underlie the suppression of responses in V1 simple and complex cells. Three of the mechanisms identified – contrast-

dependent trial-to-trial variability (Chapter II), orientation-dependent intracortical excitation (Chapter II), and the convergence of input from cells with variable spatial frequency tuning (Chapter III) – do not require intracortical inhibition. Surround suppression does involve intracortical inhibition, but the suppression ultimately derives from decreases in excitation (Chapter IV). Our findings raise new questions about how cortical response properties arise, suggesting that both the anatomical organization of visual cortex and the interplay of excitation and inhibition may make important contributions to determining cortical computations.

References

Abbott LF, Varela JA, Sen K, Nelson SB (1997) Synaptic depression and cortical gain control. Science 275:220-224.

Adelson EH, Bergen JR (1985) Spatiotemporal energy models for the perception of motion. J Opt Soc Am 2:284-299.

Ahmed B, Anderson JC, Douglas RJ, Martin KA, Nelson JC (1994) Polyneuronal innervation of spiny stellate neurons in cat visual cortex. J Comp Neurol 341:39-49.

Akasaki T, Sato H, Yoshimura Y, Ozeki H, Shimegi S (2002) Suppressive effects of receptive field surround on neuronal activity in the cat primary visual cortex. Neurosci Res 43:207-220.

Alitto HJ, Usrey WM (2004) Influence of contrast on orientation and temporal frequency tuning in ferret primary visual cortex. J Neurophysiol 91:2797-2808.

Alonso JM, Martinez LM (1998) Functional connectivity between simple cells and complex cells in cat striate cortex. Nat Neurosci 1:395-403.

Anderson J, Lampl I, Reichova I, Carandini M, Ferster D (2000a) Stimulus dependence of two-state fluctuations of membrane potential in cat visual cortex. Nat Neurosci 3:617-621.

Anderson JS, Carandini M, Ferster D (2000b) Orientation tuning of input conductance, excitation, and inhibition in cat primary visual cortex. J Neurophysiol 84:909-926.

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

Anderson JS, Lampl I, Gillespie DC, Ferster D (2001) Membrane potential and conductance changes underlying length tuning of cells in cat primary visual cortex. J Neurosci 21:2104-2112.

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.

Azouz R, Gray CM (2000) Dynamic spike threshold reveals a mechanism for synaptic coincidence detection in cortical neurons in vivo. Proc Natl Acad Sci U S A 97:8110-8115.

Bair W, Cavanaugh JR, Movshon JA (2003) Time course and time-distance relationships for surround suppression in macaque V1 neurons. J Neurosci 23:7690-7701.

Bannister NJ, Nelson JC, Jack JJ (2002) Excitatory inputs to spiny cells in layers 4 and 6 of cat striate cortex. Philos Trans R Soc Lond B Biol Sci 357:1793-1808.

Bishop PO, Coombs JS, Henry GH (1971) Interaction effects of visual contours on the discharge frequency of simple striate neurones. J Physiol (Lond) 219:659-687.

Bishop PO, Coombs JS, Henry GH (1973) Receptive fields of simple cells in the cat striate cortex. J Physiol (Lond) 231:31-60.

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

Blakemore C, Carpenter RH, Georgeson MA (1970) Lateral inhibition between orientation detectors in the human visual system. Nature 228:37-39.

Borg-Graham LJ, Monier C, Fregnac Y (1998) Visual input evokes transient and strong shunting inhibition in visual cortical neurons. Nature 393:369-373.

Boudreau CE, Ferster D (2005) Short-term depression in thalamocortical synapses of cat primary visual cortex. J Neurosci 25:7179-7190.

Brainard DH (1997) The Psychophysics Toolbox. Spat Vis 10:433-436.

Bruno RM, Sakmann B (2006) Cortex is driven by weak but synchronously active thalamocortical synapses. Science 312:1622-1627.

Buzas P, Eysel UT, Adorjan P, Kisvarday ZF (2001) Axonal topography of cortical basket cells in relation to orientation, direction, and ocular dominance maps. J Comp Neurol 437:259-285.

Callaway EM (1998) Local circuits in primary visual cortex of the macaque monkey. Annu Rev Neurosci 21:47-74.

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 primate visual cortex. Science 264:1333-1336.

Carandini M, Ferster D (1997) A tonic hyperpolarization underlying contrast adaptation in cat visual cortex [see comments]. 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, 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.

Chagnac-Amitai Y, Connors BW (1989) Horizontal spread of synchronized activity in neocortex and its control by GABA-mediated inhibition. J Neurophysiol 61:747-758.

Chance FS, Nelson SB, Abbott LF (1999) Complex cells as cortically amplified simple cells. Nat Neurosci 2:277-282.

Chance FS, Abbott LF, Reyes AD (2002) Gain modulation from background synaptic input. Neuron 35:773-782.

Chung S, Ferster D (1998) Strength and orientation tuning of the thalamic input to simple cells revealed by electrically evoked cortical suppression. Neuron 20:1177-1189.

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

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) Organization of supression in receptive fields of neurons in cat visual cortex. J Neurophysiol 68:144-163.

Elston GN (2003) Cortex, cognition and the cell: new insights into the pyramidal neuron and prefrontal function. Cereb Cortex 13:1124-1138.

Emerson RC, Citron MC, Vaughn WJ, Klein SA (1987) Nonlinear directionally selective subunits in complex cells of cat striate cortex. J Neurophysiol 58:33-53.

Ferster D (1986) Orientation selectivity of synaptic potentials in neurons of cat primary visual cortex. J Neurosci 6:1284-1301.

Ferster D (1988) Spatially opponent excitation and inhibition in simple cells of the cat visual cortex. J Neurosci 8:1172-1180.

Ferster D, Lindström S (1983) An intracellular analysis of geniculocortical connectivity in area 17 of the cat. J Physiol (Lond) 342:181-215.

Ferster D, Miller KD (2000) Neural mechanisms of orientation selectivity in the visual cortex. Annu Rev Neurosci 23:441-471.

Ferster D, Chung S, Wheat H (1996) Orientation selectivity of thalamic input to simple cells of cat visual cortex. Nature 380:249-252.

Finn IM, Priebe NJ, Ferster D (2007) The emergence of contrast-invariant orientation tuning in simple cells of the cat visual cortex. Neuron 54:137-152.

Fitzpatrick D (1996) The functional organization of local circuits in visual cortex: insights from the study of tree shrew striate cortex. Cereb Cortex 6:329-341.

Freeman TC, Durand S, Kiper DC, Carandini M (2002) Suppression without inhibition in visual cortex. Neuron 35:759-771.

Gardner JL, Anzai A, Ohzawa I, Freeman RD (1999) Linear and nonlinear contributions to orientation tuning of simple cells in the cat's striate cortex. Vis Neurosci 16:1115-1121.

Gawne TJ, Martin JM (2002) Responses of primate visual cortical V4 neurons to simultaneously presented stimuli. J Neurophysiol 88:1128-1135.

Gilbert CD (1977) Laminar differences in receptive field properties of cells in cat primary visual cortex. J Physiol 268:391-421.

Gilbert CD, Wiesel TN (1979) Morphology and intracortical projections of functionally characterized neurons in the cat visual cortex. Nature 280:120-125.

Haider B, Duque A, Hasenstaub AR, McCormick DA (2006) Neocortical network activity in vivo is generated through a dynamic balance of excitation and inhibition. J Neurosci 26:4535-4545.

Haider B, Duque A, Hasenstaub AR, Yu Y, McCormick DA (2007) Enhancement of visual responsiveness by spontaneous local network activity in vivo. J Neurophysiol 97:4186-4202.

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

Hartline HK (1949) Inhibition of activity of visual receptors by illuminating nearby retinal areas in the Limulus eye. Federal Proceedings 8:69.

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

Henry GH, Goodwin AW, Bishop PO (1978) Spatial summation of responses in receptive fields of single cells in cat striate cortex. Exp Brain Res 32:245-266.

Hirsch JA, Alonso JM, Reid RC, Martinez LM (1998) Synaptic integration in striate cortical simple cells. J Neurosci 18:9517-9528.

Hirsch JA, Martinez LM, Pillai C, Alonso JM, Wang Q, Sommer FT (2003) Functionally distinct inhibitory neurons at the first stage of visual cortical processing. Nat Neurosci 6:1300-1308.

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

Hubel DH, Wiesel TN (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. J Physiol (Lond) 160:106-154.

Hubel DH, Wiesel TN (1965) Receptive Fields and Functional Architecture in Two Nonstriate Visual Areas (18 and 19) of the Cat. J Neurophysiol 28:229-289.

Jagadeesh B, Wheat HS, Kontsevich L, Tyler CW, Ferster D (1997) Direction selectivity of synaptic potentials in simple cells of the cat visual cortex. J Neurophysiol 78:2772-2789.

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

Kayser A, Priebe NJ, Miller KD (2001) Contrast-dependent nonlinearities arise locally in a model of contrast-invariant orientation tuning. J Neurophysiol 85:2130-2149.

Lampl I, Anderson JS, Gillespie DC, Ferster D (2001) Prediction of orientation selectivity from receptive field architecture in simple cells of cat visual cortex. Neuron 30:263-274.

Lampl I, Ferster D, Poggio T, Riesenhuber M (2004) Intracellular measurements of spatial integration and the MAX operation in complex cells of the cat primary visual cortex. J Neurophysiol 92:2704-2713.

Latham PE, Nirenberg S (2004) Computing and stability in cortical networks. Neural Comput 16:1385-1412.

Latham PE, Richmond BJ, Nelson PG, Nirenberg S (2000) Intrinsic dynamics in neuronal networks. I. Theory. J Neurophysiol 83:808-827.

Lauritzen TZ, Miller KD (2003) Different roles for simple-cell and complex-cell inhibition in V1. J Neurosci 23:10201-10213.

Levitt JB, Lund JS (1997) Contrast dependence of contextual effects in primate visual cortex. Nature 387:73-76.

Li B, Thompson JK, Duong T, Peterson MR, Freeman RD (2006) Origins of cross-orientation suppression in the visual cortex. J Neurophysiol 96:1755-1764.

Li CY, Li W (1994) Extensive integration field beyond the classical receptive field of cat's striate cortical neurons--classification and tuning properties. Vision Res 34:2337-2355.

Livingstone MS, Conway BR (2003) Substructure of direction-selective receptive fields in macaque V1. J Neurophysiol 89:2743-2759.

MacLean JN, Watson BO, Aaron GB, Yuste R (2005) Internal dynamics determine the cortical response to thalamic stimulation. Neuron 48:811-823.

Marino J, Schummers J, Lyon DC, Schwabe L, Beck O, Wiesing P, Obermayer K, Sur M (2005) Invariant computations in local cortical networks with balanced excitation and inhibition. Nat Neurosci 8:194-201.

Martinez LM, Alonso JM (2001) Construction of complex receptive fields in cat primary visual cortex. Neuron 32:515-525.

Martinez LM, Alonso JM (2003) Complex receptive fields in primary visual cortex. Neuroscientist 9:317-331.

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

Martinez LM, Wang Q, Reid RC, Pillai C, Alonso JM, Sommer FT, Hirsch JA (2005) Receptive field structure varies with layer in the primary visual cortex. Nat Neurosci 8:372-379.

McGuire BA, Gilbert CD, Rivlin PK, Wiesel TN (1991) Targets of horizontal connections in macaque primary visual cortex. J Comp Neurol 305:370-392.

McLaughlin D, Shapley R, Shelley M (2003) Large-scale modeling of the primary visual cortex: influence of cortical architecture upon neuronal response. J Physiol Paris 97:237-252.

Mel BW, Ruderman DL, Archie KA (1998) Translation-invariant orientation tuning in visual "complex" cells could derive from intradendritic computations. J Neurosci 18:4325-4334.

Miller KD (1994) A model for the development of simple cell receptive fields and the ordered arrangement of orientation columns through activity-dependent competition between ON- and OFF-center inputs. J Neurosci 14:409-441.

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

Monier C, Chavane F, Baudot P, Graham LJ, Fregnac Y (2003) Orientation and direction selectivity of synaptic inputs in visual cortical neurons: a diversity of combinations produces spike tuning. Neuron 37:663-680.

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

Movshon JA, Thompson ID, Tolhurst DJ (1978a) Spatial summation in the receptive fields of simple cells in the cat's striate cortex. Journal of Physiology (London) 283:53-77.

Movshon JA, Thompson ID, Tolhurst DJ (1978b) Receptive field organization of complex cells in the cat's striate cortex. J Physiol 283:79-99.

Nolt MJ, Kumbhani RD, Palmer LA (2004) Contrast-dependent spatial summation in the lateral geniculate nucleus and retina of the cat. J Neurophysiol 92:1708-1717.

Ohzawa I, DeAngelis GC, Freeman RD (1990) Stereoscopic depth discrimination in the visual cortex: neurons ideally
suited as disparity detectors. Science 249:1037-1141.

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.

Palmer LA, Rosenquist AC (1974) Visual receptive fields of single striate cortical units projecting to the superior colliculus in the cat. Brain Res 67:27-42.

Palmer LA, Jones JP (1984) Quantitative analysis of cat simple receptive fields in two spatial and two temporal frequency dimensions. Soc Neurosci Abstr 10:800.

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

Priebe NJ, Ferster D (2005) Direction selectivity of excitation and inhibition in simple cells of the cat primary visual cortex. Neuron 45:133-145.

Priebe NJ, Ferster D (2006) Mechanisms underlying cross-orientation suppression in cat visual cortex. Nat Neurosci 9:552-561.

Priebe NJ, Mechler F, Carandini M, Ferster D (2004) The contribution of spike threshold to the dichotomy of cortical simple and complex cells. Nat Neurosci.

Ramoa AS, Shadlen M, Skottun BC, Freeman RD (1986) A comparison of inhibition in orientation and spatial frequency selectivity of cat visual cortex. Nature 321:237-239.

Reid RC, Alonso JM (1995) Specificity of monosynaptic connections from thalamus to visual cortex. Nature 378:281-284.

Riesenhuber M, Poggio T (1999) Hierarchical models of object recognition in cortex. Nat Neurosci 2:1019-1025.

Riesenhuber M, Poggio T (2002) Neural mechanisms of object recognition. Curr Opin Neurobiol 12:162-168.

Rust NC, Schwartz O, Movshon JA, Simoncelli EP (2005) Spatiotemporal elements of macaque v1 receptive fields. Neuron 46:945-956.

Sanchez-Vives MV, Nowak LG, McCormick DA (1997) Cellular and network mechanisms generating adaptation to contrast in the visual cortex: An in vivo and in vitro study. Soc Neurosci Abs 23:1944.

Sato T (1989) Interactions of visual stimuli in the receptive fields of inferior temporal neurons in awake macaques. Exp Brain Res 77:23-30.

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.

Sclar G, Freeman RD (1982) Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. Exp Brain Res 46:457-461.

Serre T, Wolf L, Bileschi S, Riesenhuber M, Poggio T (2007) Robust object recognition with cortex-like mechanisms. IEEE Trans Pattern Anal Mach Intell 29:411-426.

Shu Y, Hasenstaub A, McCormick DA (2003) Turning on and off recurrent balanced cortical activity. Nature 423:288-293.

Skottun BC, Bradley A, Sclar G, Ohzawa I, Freeman R (1987) The effects of contrast on visual orientation and spatial frequency discrimination: a comparison of single cells and behavior. J Neurophysiol 57:773-786.

Skottun BC, De Valois RL, Grosof DH, Movshon JA, Albrecht DG, Bonds AB (1991) Classifying simple and complex cells on the basis of response modulation. Vision Res 31:1079-1086.

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.

Somers DC, Nelson SB, Sur M (1995) An emergent model of orientation selectivity in cat visual cortical simple cells. J Neurosci 15:5448-5465.

Sompolinsky H, Shapley R (1997) New perspectives on the mechanisms for orientation selectivity. Current Opinion in Neurobiology 7:514-522.

Sompolinsky H, Golomb D, Kleinfeld D (1990) Global processing of visual stimuli in a neural network of coupled oscillators. PNAS 87:7200-7204.

Stettler DD, Das A, Bennett J, Gilbert CD (2002) Lateral connectivity and contextual interactions in macaque primary visual cortex. Neuron 36:739-750.

Stratford KJ, Tarczy HK, Martin KA, Bannister NJ, Jack JJ (1996) Excitatory synaptic inputs to spiny stellate cells in cat visual cortex. Nature 382:258-261.

Szulborski RG, Palmer LA (1990) The two-dimensional spatial structure of nonlinear subunits in the receptive fields of complex cells. Vision Res 30:249-254.

Tao L, Shelley M, McLaughlin D, Shapley R (2004) An egalitarian network model for the emergence of simple and complex cells in visual cortex. Proc Natl Acad Sci U S A 101:366-371.

Tolhurst DJ, Heeger DJ (1997) Comparison of contrast-normalization and threshold models of the responses of simple cells in cat striate cortex. Vis Neurosci 14:293-309.

Touryan J, Lau B, Dan Y (2002) Isolation of relevant visual features from random stimuli for cortical complex cells. J Neurosci 22:10811-10818.

Touryan J, Felsen G, Dan Y (2005) Spatial structure of complex cell receptive fields measured with natural images. Neuron 45:781-791.

Toyama K, Kimura M, Tanaka K (1981) Organization of cat visual cortex as investigated by cross-correlation technique. J Neurophysiol 46:191-201.

Troyer TW, Krukowski AE, Priebe NJ, Miller KD (1998) Contrast-invariant orientation tuning in cat visual cortex: thalamocortical input tuning and correlation-based intracortical connectivity. J Neurosci 18:5908-5927.

Tsodyks MV, Markram H (1997) The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability [published erratum appears in Proc Natl Acad Sci U S A 1997 May 13;94(10):5495]. Proc Natl Acad Sci U S A 94:719-723.

Tsodyks MV, Skaggs WE, Sejnowski TJ, McNaughton BL (1997) Paradoxical effects of external modulation of inhibitory interneurons. J Neurosci 17:4382-4388.

Tucker TR, Fitzpatrick D (2006) Luminance-evoked inhibition in primary visual cortex: a transient veto of simultaneous and ongoing response. J Neurosci 26:13537-13547.

van Vreeswijk C, Sompolinsky H (1998) Chaotic balanced state in a model of cortical circuits. Neural Comput 10:1321-1371.

Volgushev M, Pernberg J, Eysel UT (2000) Comparison of the selectivity of postsynaptic potentials and spike responses in cat visual cortex. Eur J Neurosci 12:257-263.

Walker GA, Ohzawa I, Freeman RD (2000) Suppression outside the classical cortical receptive field. Vis Neurosci 17:369-379.

Waters J, Helmchen F (2006) Background synaptic activity is sparse in neocortex. J Neurosci 26:8267-8277.

Webb BS, Dhruv NT, Solomon SG, Tailby C, Lennie P (2005) Early and late mechanisms of surround suppression in striate cortex of macaque. J Neurosci 25:11666-11675.

Wiesel TN (1959) Recording inhibition and excitation in the cat's retinal ganglion cells with intracellular electrodes. Nature 183:264-265.