Signal transmission, feature representation and computation in areas V1 and MT of the macaque monkey

by

Nicole C. Rust

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Center for Neural Science

New York University

September, 2004

_____

J Anthony Movshon

Acknowledgements

I have had the great fortune of working with two mentors during my graduate studies, Tony Movshon and Eero Simoncelli. I hold profound admiration for the excitement, insight, integrity, creativity, and rigor that these two men bring to science and I consider the lessons that I have learned from them invaluable. If I had to do it over, I would chose to work with both of them again; I cannot imagine two pairs of footsteps I would rather attempt to follow.

I am thankful to Mal Semple, the chair of my committee, for many years of wisdom, advice and encouragement. I also thank my examiners, Ken Miller and Jonathan Victor, for helpful comments and discussion of the manuscript. Simon Schultz, Najib Majaj and Adam Kohn helped me shape my vague notions into concrete ideas and Jonathan Pillow was a valuable resource for much of this work. I am indebted to Leanne Chukoskie, Hannah Bayer, Anita Disney, Simon Schultz, Najib Majaj, Stu Greenstein and Lynne Kiorpes for support, advice, and the patience to listen as I worked things through.

The excellent education I received at the Center for Neural Science is the product of many efforts, particularly those of Sam Feldman. I would also like to thank the teachers and scientists who have nurtured my interests in science through their enthusiasm and example: Dean Lindstrom, Richard Schreiber, Doug England, Nick Hoffman, Thomas Bitterwolf, Michael Laskowski, Ann

Norton, Allan Caplan, Douglas Baxter, Steve Munger, David Weinshenker, Richard Palmiter, James Austin, Simon Schultz, Bob Shapley, Mike Hawken, Dan Sanes, Bill Bialek, Paul Glimcher, and E. J. Chichilnisky.

I greatly appreciate the support and encouragement of my mother, Christine, my grandmother, Blanche Permoda, and my brother, Ferris. I would like pay a special thanks to my father, David Rust. Throughout my life, he has shown me the beauty of curiosity, the power of explanation, and the work ethic needed to make it all happen.

Preface

The second chapter of this thesis was a collaborative effort between myself, Simon Schultz (currently at Imperial College, London), and my advisor J. Anthony Movshon (New York University). This work has been published (Rust et al, 2002). Chapter three arose from a collaboration between myself, Odelia Schwartz (currently at the Salk Institute), my second advisor Eero Simoncelli (New York Unviersity), and J. Anthony Movshon; portions of this work have been published as well (Rust et al, 2004). Eero Simoncelli and J. Anthony Movshon were involved in the work presented in chapter four.

Table of contents

List of figures

1 Introduction

Within each sensory modality, processing begins by decomposing the physical world into the most fundamental components (e.g. light intensity or sound frequency) by sensory receptors. All the information available to an organism is determined at this early stage; the data processing inequality maintains that from this point on information can not increase. However, the form in which the information is stored and organized can change. In all sensory systems, the representation of sensory information becomes increasingly complex as one ascends the processing hierarchy. Consequently, early sensory processing can be examined from two perspectives. First, to what degree is information preserved as it is propagated through the brain? The problem of signal transmission can be viewed as a problem of "how much" information is present at each stage of processing independent of "what" is being encoded. Alternatively, one can ask: what features of the world are represented in the firing patterns of neurons in a particular subcortical structure or cortical area? Questions related to the representation of information are closely related to questions regarding the computations neurons perform to achieve those representations. This chapter focuses on a review of signal transmission, feature representation and computation at the early stages of visual processing.

1.1 Signal transmission

The nervous system is noisy. When presented with the same stimulus on repeated trials, neurons respond with a variable number of spikes. In the cortex, response variance increases linearly with mean firing rate; the ratio between the variance and mean rate, referred to as the Fano factor, is often used to quantify the variability of these neurons. On average, cortical neurons respond to stimuli with a response variance 1-2 times the mean response rate (Bradley et al 1987, Britten et al 1993, Scobey & Gabor 1989, Snowden et al 1992, Softky & Koch 1993, 1983, Vogels et al 1989).

The source of noise in these neurons is unclear. Intrinsic noise (e.g. spike generation) was once thought to be a significant source of unreliability. Intracellular current injections have since determined that the transformation of the intracellular potential into a spike train occurs with a higher fidelity than previously appreciated, suggesting that noise also arises from alternate sources such as synaptic transmission or alternate intracellular processes (Mainen & Sejnowski 1995). Noise appears to increase as signals propagate through subcortical structures (Kara et al 2000) but remains approximately constant across cortical areas (Softky & Koch 1993). Within the highly interconnected networks found in the cortex, balanced excitation and inhibition may play a key role in maintaining constant ratios of signal and noise (Shadlen & Newsome 1998).

Noise limits how well neurons can report information about a stimulus and consequently the amount of information available to an organism for perception and action. Signal detection theory has been used as a method to compare the ability of an observer to discriminate different stimuli with the ability of single neurons to perform the same task (Britten et al 1996, Britten et al 1992, Prince et al 2000, Prince et al 2002, Shadlen et al 1996). A similar measure of discriminability is taken from Shannon's information theory (Shannon 1948). Information theory has advantages over signal detection theory when working with large stimulus sets or stimuli that are difficult to parameterize (e.g. stochastic or naturalistic stimuli) and when one suspects that the response probability distributions are non-Gaussian (Buracas & Albright 1999).

1.1.1 Shannon's information theory

Shannon's information theory was introduced as a general technique for evaluating the transmission of a signal in the presence of noise, thus making it applicable to neural systems. *Mutual information* quantifies how much one can discern about a message after it has been passed down a noisy channel and received at the other end; this metric has been applied to various sensory systems to quantify how much can be determined from the responses of a particular neuron or neuronal population (Bialek et al 1991, Buracas et al 1998, Rieke et al 1995, Rolls et al 1997, Theunissen et al 1996, Warland et al 1997).

Consider the responses of a neuron to repeated presentations of a stimulus at seven contrasts (figure 1-1). A typical visual neuron's firing rate will increase with increasing contrast up to a point at which the response saturates. Likewise, the variability will increase in proportion to the mean. To quantify the information this neuron reports about contrast, one begins by constructing probability distribution histograms for the number of spikes elicited on each trial at each contrast level, the probability of a response (spike count) given a stimulus, $P(r|s)$ (fig 1-1 right, black and gray). Calculation of the probability distribution across all stimuli, $P(r)$ is also required (figure 1-1, red).

The task of the neuron is to report the contrast of a stimulus with the magnitude of its firing rate. If the seven contrasts are presented with equal probability, the probability of making a correct guess in the absence of firing rate information is 1/7. Mutual information about the stimuli given the responses, $I(S;R)$, quantifies the reduction in uncertainty when the firing rate of the neuron is taken into account. Information is typically measured in quantities of bits where a bit is the amount of information required to distinguish two binary alternatives (0 or 1). If the response distributions corresponding to different contrast levels were completely nonoverlapping, the neuron would be a perfect discriminator and the information reported by the neuron would be $\log_2(7)= 2.8$ bits. If the response distributions were completely overlapping, the contrast corresponding to a given response would be ambiguous and the

information would be zero bits. Partially overlapping distributions result in intermediate values.

The entropy of a distribution *P(x)* measures the amount of information required to specify the variability in that distribution and is computed by:

$$H = \sum_{x} P(x) \log_2 P(x)$$

Computation of the information the example neuron reports about contrast requires computation of two entropies. First, one needs to compute the amount of noise present in the responses at each contrast, the noise entropy *(H_{noise})* of *P(r|s)* (figure 1-1, black and gray distributions). Similarly, one needs to compute the entropy across all responses *(H_{total})* from *P(r)* (figure 1-1, red curve). The mutual information is the difference between the total entropy and the mean noise entropy (Cover & Thomas 1991):

$$I(R;S) = \sum_{r} P(r) \log_2 P(r) - \sum_{s} P(s) \sum_{r} P(r \mid s) \log_2 P(r \mid s)$$

Mutual information computed in this way quantifies the average discriminability between contrasts.

Although information has an advantage over signal detection theory in that the shape of the probability distributions are not assumed but rather calculated directly, the process of estimating probability distributions is known to lead to systematic biases in information measurements (Carlton 1969). If the number of samples is small and the number of bins used for the histogram is large, the bins will be sparsely sampled and systematic overestimations in

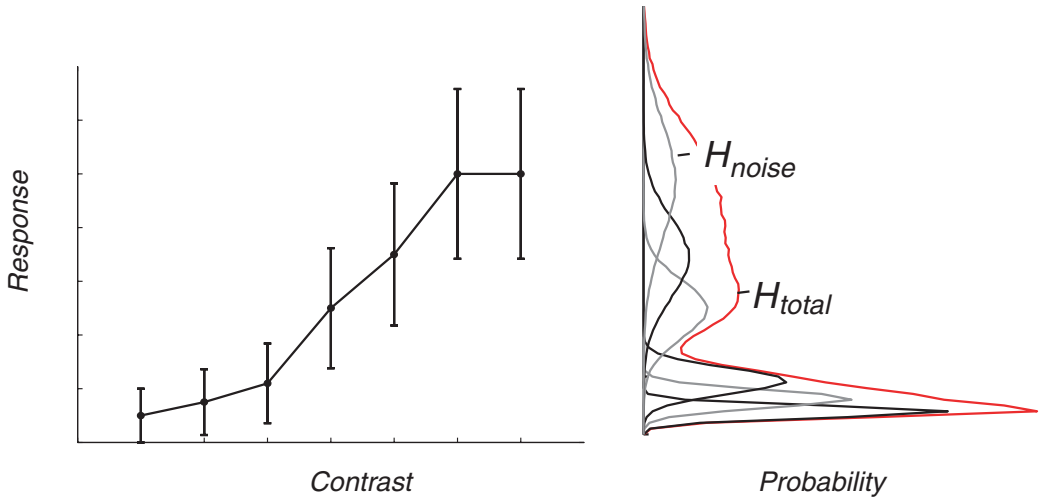

Figure 1-1. Computation of the mutual information about the contrast of a stimulus from the distribution of responses. Left: the contrast versus response function for a toy neuron. Error bars indicate the mean and standard deviation of firing rate across multiple presentations of stimuli at seven different contrasts. Note that the variance of the response increases in proportion to the mean rate. Right: the number of spikes elicited across trials for each contrast plotted as a normalized histogram (a probability distribution) for each of the seven contrasts (gray, black), referred to in the text as the probability of a response given a stimulus, *P(r|s)*. The information the neuron reports about contrast decreases with the amount of overlap of these distributions. Also shown is a histogram of the responses across all contrasts, *P(r)* (red). Mutual information is calculated as the difference in entropy between *P(r)*, labeled H_{total}, and the mean entropy of *P(r|s)*, labeled H_{noise}.

information will result. A number of methods have been developed to correct for this bias, including analytical estimation and correction of the bias (Panzeri & Treves 1996), neural network techniques (Hertz et al 1995) and Monte Carlo methods (Abbott et al 1996).

Chapter 2 is devoted to a characterization of response variability in visual area V1 during different stages of development. Both classical and information theoretic techniques are presented to characterize variability and discriminability in infant and adult neurons.

1.2 Representation and computation in early visual processing

How is the external world represented in the firing patterns of neurons? Visual processing begins by decomposing the world into a set of spatially localized light intensities. However, we perceive the visual world not in terms of light intensity but rather in terms of objects and their relative positions in space. One approach toward understanding the implementation of this sophisticated internal representation is to trace visual signals as they are transformed from the rudimentary description found in the retina through each stage of the processing hierarchy.

When pursuing a system from the "bottom-up", the question of representation in the brain is inextricably linked to the computations performed by neurons to achieve that representation. To determine the computations that occur at each stage of sensory processing, it is useful to build functional models

of neurons that describe the transformation of a stimulus into a neuron's firing pattern in terms of mathematical operations performed upon the stimulus. Such models describe not only processing performed by the neuron in question, but also include operations executed by neurons preceding it. As such, these models are not strictly biophysical, but rather serve to provide a compact description of sensory computation up to a particular stage of the processing hierarchy. In this section, I begin by reviewing the rich history of functional models in early vision and the linear and nonlinear systems analysis techniques used in their characterization. I then focus on one computation performed in early vision: the computation of motion direction within visual areas V1 and MT.

1.2.1 Functional models in early vision

Linear characterization of retinal ganglion cells:

The first efforts to describe visual receptive fields quantitatively through linear systems analysis were made by Rodieck and Stone (1965a, 1965b). They introduced the concept of considering the responses of visual neurons in terms of a linear sum of the light intensities falling on their receptive fields. Rodieck and Stone demonstrated that the linear weighting function of a retinal ganglion cell could be constructed by presenting small flashes of light to different locations of the receptive field and calculating a histogram of the responses at each position. The responses of the neuron to different stimuli (e.g. a moving bar) could then be predicted from this linear receptive field map. Description of

a neuron's response to a stimulus in terms of spatiotemporal linear filtering has proven to be a powerful predictor of a neuron's response at many stages of visual processing.

Rodieck (1965) also introduced the ideal of describing these receptive field maps with parametric mathematical models whose parameters could be adjusted to fit different neurons. These models, he proposed, should include the simplest possible mathematical description of the neurons in question. In the retina, Rodieck demonstrated that the center-surround organization of the retinal ganglion cell was efficiently and accurately described by a model containing a difference of two Gaussians (figure 1-2a).

Following the theme of analyzing retinal ganglion cell receptive fields as linear filters, Enroth-Cugell and Robson (1966) introduced the powerful technique of Fourier analysis to visual physiology. The power of Fourier theory rests on the fact that any visual stimulus can be decomposed into a linear combination of sinusoidal gratings. Thus the response of a linear system to a stimulus can be predicted by its responses to the individual grating components of that stimulus. Conversely, deviations from linearity can be identified through discrepancies between the linear prediction and the actual response. In their report, Enroth-Cugell and Robson identified two classes of retinal ganglion cells in the cat retina based upon the responses of these cells to stationary sinusoidal gratings presented at different phases. The responses of X-cells were in accord with the linear mapping of their line weighting functions; their responses

depended predicatively on the alignment of the phase of the stimulus and the center-surround organization of the receptive field. In contrast, the responses of Y-cells could not be predicted in the same manner; Y-cells responded to a stationary grating regardless of phase.

In the same paper, Enroth-Cugell and Robson illustrated the relationship between the linear weighting function of a cell and its Fourier spectra. The minimal contrast required to evoke a criterion response from a retinal ganglion cell (its contrast sensitivity) depends on the spatial frequency of a sinusoidal grating. These authors demonstrated that this relationship was in fact predicted by a Fourier transform of the difference-of-Gaussians model. The direct relationship between the linear weighting function of a neuron and its response to drifting sinusoidal gratings has been used to characterize visual neurons at many different levels of processing.

Characterization of retinal nonlinearities:

The linear description of a cell is only effective to the degree that the neuron behaves linearly. Characterization of the linear properties of a neuron's response can be constructed with small impulses and sinusoidal gratings, as described above. Characterization of the nonlinear properties of a neuron's response can be much more difficult. Wiener kernel analysis is a generalized technique for characterizing a system, regardless of the nature of the nonlinearities contained therein (Wiener 1958). This analysis effectively

provides a series expansion of a neuron's response properties by describing the dependence of a neuron's response on increasingly higher order correlations between stimulus dimensions with each successive term. Wiener kernels are characterized by presenting Gaussian white noise to a neuron. Gaussian white noise has the property of producing every possible stimulus combination (to a given resolution) given an infinite duration, hence infinite order correlations between stimulus dimensions can theoretically be computed. To compute a Wiener model for a neuron, the responses to a Gaussian white noise stimulus are recorded and the spike-triggered stimulus distribution, the stimulus history before each spike, is collected. Increasingly higher order statistical descriptions of this distribution are then calculated for inclusion in the model. A Wiener model of a cell would include terms that can approximately be described as: a baseline response (zeroth order term), the mean stimulus before a spike (first order term), the dependency of spiking on the covariance between pairs of dimensions (the second order term) and so on. In practice, each successive term requires an exponential amount of data to compute and terms beyond the second or third require more data than are accessible by current experimental techniques.

Marmarelis and Naka (1972) were the first to apply Wiener kernel analysis to neurons in the visual system. They recovered Wiener models of the temporal tuning properties of catfish retinal ganglion cells by injecting a Gaussian white noise current into their horizontal cell inputs. In addition to

recovering the first order temporal responses for these neurons, they identified significant second-order dependencies, indicative of nonlinearities in these neurons.

The temporal nonlinearity in Y-cells first described by Enroth-Cugell and Robson (1966) was further characterized by Hochestein and Shapley (1976) using stationary, sine-reversing (counterphase) gratings. In response to a counterphase grating stimulus, Y-cells produce a large nonlinear response at double the temporal frequency of the grating (the 2^{nd} harmonic). Hochstein and Shapley demonstrated that the ratio of the 2^{nd} and 1^{st} harmonic could reliably characterize neurons as X or Y-cells (cells with a $2^{nd}/1^{st}$ harmonic > 1 were classified as Y; cells with a $2^{nd}/1^{st}$ harmonic<1 were classified as X). Victor and Shapley (1979) further characterized this nonlinearity using a technique analogous to Wiener kernel analysis but applied in the frequency domain. Stimuli used in the characterization were comprised of sums of 6-8 temporally modulated (spatially stationary) sinusoids and the actual responses of retinal ganglion cells were compared with the linear prediction to identify nonlinearities in these neurons. While X-cells response's were primarily linear, in Y-cells they found large second-order dependencies between gratings at different temporal frequencies, indicative of nonlinear processing in these cells. Because these nonlinearities were present when the gratings were presented at spatial frequencies too high for the Y-cell center to resolve, they concluded that the nonlinearity must act before spatial pooling in the Y-cell. A second nonlinearity

was identified by Shapley and Victor (1978) in retinal ganglion cells. Using the same sum-of-sinusoids technique, they demonstrated a contrast gain control mechanism in these neurons: with increasing contrast the first-order behavior shifted toward higher temporal frequencies and was more sharply tuned.

Primary visual cortex (V1):

Retinal ganglion cells project to the lateral geniculate nucleus of the thalamus (LGN); receptive field properties in the LGN are nearly indistinguishable from those of retinal ganglion cells. Cells in the LGN provide the inputs to neurons in primary visual cortex (V1). Hubel and Wiesel demonstrated in the cat (1962) and later in the monkey (1968) that most V1 neurons are tuned for the orientation of bars passed across their receptive field and a subset are tuned for the direction of bar motion along this axis. In addition, they identified two classes of cells: simple and complex. Simple cells respond in a sign-dependant fashion to the polarity of a bar and its position on the receptive field whereas complex cells respond in a polarity or position insensitive manner. The functional models describing each of these computations in V1 are explained below.

Linear characterization of simple cells:

Hubel and Wiesel (1962) suggested that orientation tuning in V1 cells could be conferred by appropriately arranging LGN inputs (with different spatial

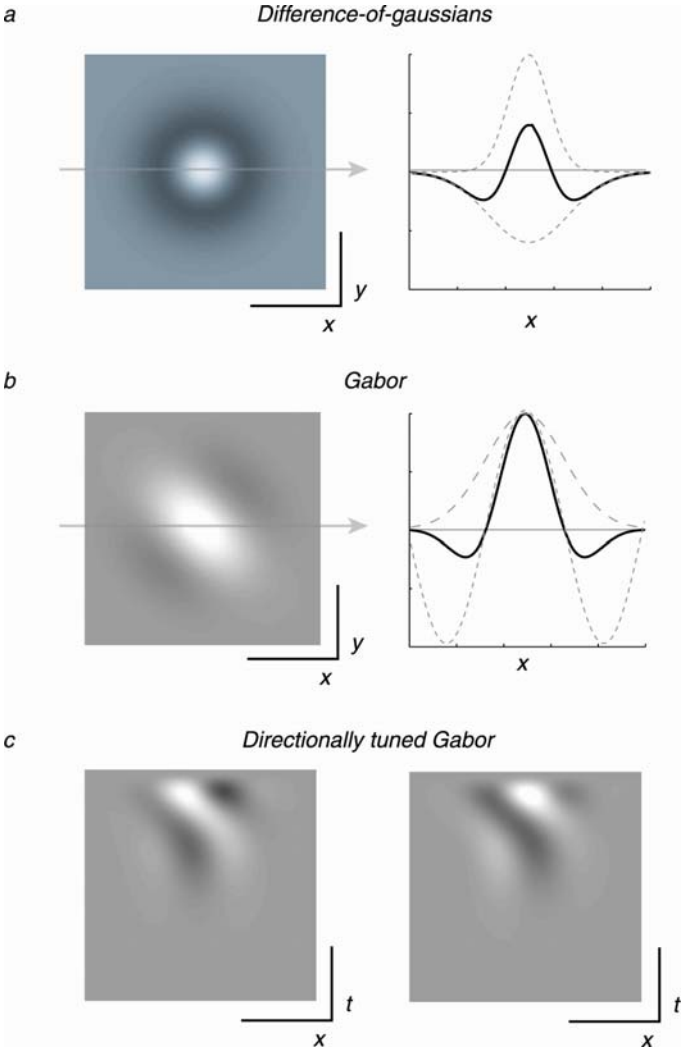

Figure 1-2: Linear filters used to describe receptive fields of neurons in early vision. a) The difference-of-gaussians model used to describe the center-surround organization of retinal ganglion cells shown in x-y spatial coordinates (left). Right: a slice taken across the x spatial dimension (black). Also shown are the two Gaussian components of the model (dashed lines). b) The Gabor model used to describe the receptive field of V1 simple cells shown in x-y spatial coordinates (left). Tilt in this model confers tuning for orientation. Also shown is a slice across the x spatial dimension (black) as well as the Gaussian and sinusoidal grating components of the Gabor (dashed). c) The space-time tilted receptive fields used to describe directionally tuned V1 simple cells. The two filters are 90 degrees out of phase (quadrature pairs) and together form the subunits of an energy model complex cell.

displacements) to produce a receptive field elongated along one axis. Tests of whether simple cells could be described by their line weighting functions were first performed by Movshon et al (1978b). They constructed histograms of these neuron's responses to light and dark bars and demonstrated that Fourier transforms of these line weighting functions predicted the spatial frequency tuning of these cells to drifting sinusoidal gratings.

The functional model used to describe the spatial profile of simple cell receptive fields was first proposed by Gabor for general optimization of spatio-temporal localized signal transmission (Gabor 1946). This function, now referred to as a "Gabor" was first suggested for the description of simple cells by Marcelja (1980) and its parameters further specified by Daugman (1985). A Gabor consists of a sinusoid multiplied by a gaussian window (figure 1-2b). It's preferred orientation is determined by the tilt of the sinusoid and its aligned elongated 2-D gaussian window; the phase of the grating determines the location of its "on" and "off" subregions. To test the applicability of the Gabor to simple cell 2-D spatial profiles, Jones and Palmer (Jones & Palmer 1987a, Jones & Palmer 1987b, Jones et al 1987) mapped simple cell receptive fields using automated techniques similar to those used by Movshon et al (1976): receptive field maps were constructed by computing the mean response to small light and dark bars and taking the difference between the light and dark maps. They arrived at a nine parameter model for a Gabor, including terms to adjust the orientation, phase, and spatial frequency of the grating, and the aspect ratio of

the 2-D Gaussian window. The linear maps for simple cells were in fact well described by these functions and the spatial and temporal frequency preferences of simple cells were well predicted by a Fourier transform of their best fitting Gabor function.

The tilt of the grating in XY spatial coordinates determines the orientation tuning of this function; direction tuning can be included in the model by extending the model to include a time axis (Adelson & Bergen 1985, van Santen & Sperling 1985, Watson & Ahumada 1985). The receptive field can be envisioned as a three-dimensional volume; direction tuning in the model is conferred by shifting the phase of the grating rightward or leftward at different points in time (corresponding to rightward or leftward motion). This concept is most easily visualized by examining slices through the Gabor perpendicular to the long axis plotted against time (figure 1-2c). A neuron with an un-tilted spatiotemporal receptive field will respond equally well to opposite directions of motion whereas receptive fields that are tilted in space-time produce directionally tuned responses. Tests of directionally tuned simple cells using bars or counterphase gratings reveal that most directionally tuned simple cells in V1 produce space-time tilted maps (DeAngelis et al 1993, McLean & Palmer 1989, Movshon et al 1978, Reid et al 1987, but see Murthy et al 1998). Although these linear maps correctly predict the preferred direction of these neurons, they systematically underestimate direction selectivity by overestimating the responses in the non-preferred direction, suggesting that a

nonlinearity plays a role in shaping the direction selectivity of these neuron's responses.

Simple cell nonlinearities:

The linear map of a simple cell predicts response increments and decrements to stimuli, yet these cells have low baseline firing rates and spiking responses cannot be negative. This discrepancy is remedied by simply setting the negative responses of a filter to zero (half-wave rectification). Evidence exists to suggest that this half-rectification stage should include either an exponent (e.g. squaring) or equivalently an over-rectified process coupled with a gain parameter (Heeger 1992a). As described above, the linear estimate of directional neurons has systematically underestimated the direction bias for a neuron due to an overestimation of the response to stimuli moving in the direction opposite the preferred (Albrecht & Geisler 1991, Reid et al 1987, Tolhurst & Dean 1991). However, intracellular measurements of the receptive field are well aligned with linear estimates (Jagadeesh et al 1997). The discrepancy between the actual versus predicted direction tuning of a cell are reconciled by the addition of an expansive nonlinearity (e.g. squaring or over-rectification) which reduces the responses to the null direction relative to the preferred. (Albrecht & Geisler 1991, Heeger 1993). Similarly, conversion of a simple cell's spatial frequency tuning curve into a spatial profile (via an inverse Fourier transform) has predicted maps with more side lobes than observed in the

actual responses of many V1 neurons (Dean & Tolhurst 1983, DeAngelis et al 1993, Tadmor & Tolhurst 1989). For these cells, an expansive nonlinearity applied to the filter output has been shown to align the predicted and actual responses. Biophysically, the expansive nonlinearity is believed to arise in simple cells via intracellular noise. When the membrane voltage of a neuron is near threshold, random events can raise the intracellular potential above threshold to produce a spike. As a result, the relationship between filter output (membrane potential) and firing rate takes on the form of a power law when averaged over trials (Anderson et al 2000, Miller & Troyer 2002)

In addition to squaring, a number of apparent nonlinear phenomena in simple cells have been tied to a single nonlinear mechanism, contrast gain control. At high contrasts, simple cell responses saturate. The responses of simple cells to an excitatory stimulus are also reduced by simultaneous presentation of an ineffective stimulus, a phenomenon known as masking. Both nonlinear phenomena can be mathematically described by divisive normalization (Carandini & Heeger 1994, Heeger 1992b). Under this formulation, a simple cell's excitatory response is described by its spatiotemporal linear weighting function. The neuron is simultaneously suppressed by a signal that is proportional to the total contrast of the stimulus via a divisive process. In the case of a contrast response function, excitation exists in the numerator and contrast in the denominator and the normalization produces sigmoidal responses. In the case of a masking stimulus, the mask

suppresses the excitatory response in proportion to its total contrast. Biophysically, the source of the suppressive signal is debated. Contrast dependant suppression has been proposed to arise from inhibitory inputs from neighboring neurons (Heeger 1992b), synaptic depression (Carandini et al 2002, Chance et al 1998, Kayser et al 2001), variations in the levels of balanced excitatory and inhibitory activity (Chance et al 2002), and other sources (Kayser et al 2001).

Nonlinear processing in complex cells:

Movshon et al (1978a) introduced a two-bar interaction paradigm to characterize the responses of complex cells. They found that a bar of a given polarity (bright or dark) would evoke responses at all positions across the receptive field, indicative of phase-insensitivity. To obtain a second-order estimate of the neuron's response, they presented a bar of one polarity at a fixed position within the receptive field while varying the spatial position of a second bar of the same polarity. A two-line interaction profile was computed as the difference in the histograms with and without the stationary bar present and was shown to be in agreement with the spatial frequency tuning of the cell.

Adelson and Bergen (1985) proposed an energy model for the construction of phase-invariant responses based upon linear filter subunits. This model summed the squared outputs of two Gabor filters whose sinusoidal components were 90 degrees out of phase (figure 1-2c). Stimulus selectivity is

conferred in the model by passing the stimulus through the two linear filters at the first stage. Phase-insensitivity is conferred in two ways. First, the squaring operation applied to the output of each filter results in large responses from a filter when presented with stimuli that resemble the filter or its inverse. Second, signals from the two filters are combined to produce perfect phase invariance. Automated extensions of the techniques used by Movshon et al (1978a) have been used to further characterize the spatiotemporal fields of V1 complex cells; the results derived by those techniques are consistent with energy model predictions (Emerson et al 1992, Emerson et al 1987, Gaska et al 1994, Livingstone & Conway 2003, Szulborski & Palmer 1990).

Historical overview:

Early work in the retina introduced an elegant philosophical concept to sensory neuroscience. Researchers, including Rodieck, Enroth-Cugell, Robson, and others, derived methods for determining the linear weighting functions to describe the transformation of light intensity into the spiking responses of retinal ganglion cells. Furthermore, they demonstrated the utility of deriving parametric models whose parameters can be adjusted to describe large classes of neurons at a given stage of processing. In other words, these researchers set the criteria for "understanding" a class of neurons as a derivation of a functional model that is a good descriptor of a neuron's response to any arbitrary stimulus

(albeit improvements to functional models of retinal ganglion cells are still being made e.g. Keat et al 2001).

In primary visual cortex, numerous researchers have applied similar mapping techniques to simple cells and demonstrated that these maps are well described by Gabor functions. However, simple cells display multiple nonlinearities which are not included in the Gabor model. Furthermore, functional models exist for complex cells but fitting the parameters of these models has proven difficult, due to the strong nonlinearities in these neurons. Despite decades of work in V1, we still have not constructed and tested complete functional descriptions of these neurons. Recent advances in nonlinear systems analysis make this problem more tractable, as described below. Application of modern analysis techniques to arrive at complete functional models for V1 neurons (both simple and complex) is the focus of chapter 3.

Advances in linear and nonlinear systems analysis:

Computation of higher order Wiener kernels to account for the nonlinear behaviors of half-rectified neurons (like simple cells) has proven unsuccessful. Proper description of threshold nonlinearities requires many higher order Wiener terms but terms beyond the second-order require more data than are accessible by current experimental techniques. More recently, recovery of the functional model in two stages has been proposed as a means of determining a full model for quasilinear receptive fields like the simple cell (Anzai et al 1999,

Brenner et al 2000, Chichilnisky 2001, Chichilnisky & Baylor 1999, Emerson et al 1992, Sakai 1992). In such a characterization, a cell is stimulated with a dense, random, Gaussian noise stimulus and the linear component of its receptive field is estimated by computing the mean stimulus before the response, the spike-triggered average (equivalent to the first order Wiener kernel). Given a model of a simple cell as a single linear filter followed by an instantaneous rectifying nonlinear function and Poisson spike generation, the spike-triggered average is an unbiased estimate of the linear filter (Chichilnisky 2001, Paninski 2003). The nonlinear function that relates the output of this filter and firing rate can then be reconstructed as the ratio between histograms of the raw (all stimuli) and spiking stimulus distributions projected onto the filter, thus completing a full functional model for the neuron.

To fully account for a complex cell's behavior using the two-bar interaction technique, one must calculate all possible second-order correlations between spatiotemporal dimensions (equivalent to the second-order Wiener kernel or the spike-triggered covariance matrix). This unwieldy data structure is difficult to interpret. Furthermore, this technique only classifies spatiotemporal interactions up to the second order and thus provides an incomplete model for a neuron. A clever solution to these problems was introduced to create functional models of the fly visual system, spike-triggered covariance (Brenner et al 2000, de Ruyter van Steveninck & Bialek 1988). In this analysis, a matrix similar to the second order Wiener kernel is calculated and resolved into a set of filters by

applying a principal components analysis (PCA). Statistical methods are then used to identify filters that have a significant impact upon a neuron's spiking. STC recovers filters for which the response variance changes relative to chance occurrences between stimuli and spikes (e.g. filters whose outputs are squared), and is capable of identifying filters that have an excitatory and/or a suppressive impact on spiking. Once a set of excitatory and suppressive linear filters are recovered, the nonlinear function that describes the combination of the signals arising from each filter can be reconstructed to complete the model. In this procedure the nonlinear function is estimated empirically, thus bypassing the problems associated with classical Wiener kernel approaches. STC has proven successful at recovering the subunits of complex cells in cat visual cortex (Touryan et al 2002). Characterization of full models of both simple and complex cells using STC is the focus of chapter 3.

1.2.2 Computation of motion direction in V1 and MT

We know from psychophysics that a suppressive signal must be involved in our computation of the direction of moving stimuli. After prolonged presentation of a moving stimulus in one direction, a static stimulus will appear to move in the opposite direction – a phenomena known as the motion after-effect. The perception of movement of the static stimulus is believed to arise from adaptation of neurons that normally suppress neurons tuned for the opposite direction. Most computational models of motion processing include an opponent

(subtractive) computation between signals with opposite direction preferences (Adelson & Bergen 1985, Simoncelli & Heeger 1998, van Santen & Sperling 1984). Details regarding the neural mechanisms that confer the source of this suppressive signal remain unresolved.

While a subpopulation of neurons in V1 are tuned for the direction of moving stimuli, neurons in the LGN are not, suggesting that the computation for motion direction occurs at the first stage of visual cortical processing. Neurons in V1 can have a array of direction selectivities, ranging from neurons that respond similarly to motion in both directions to neurons that respond vigorously to motion in one direction but have little or no response to motion in the direction opposite (Hubel & Wiesel 1962). In area MT, most cells produce a response to a stimulus moving in the direction opposite its preferred that is suppressed below baseline firing rate (Felleman & Kaas 1984, Maunsell & Van Essen 1983, Rodman & Albright 1987).

As described in section 1.2.1, most directionally tuned simple cells in V1 produce space-time tilted maps (DeAngelis et al 1993, McLean & Palmer 1989, Movshon et al 1978b, Murthy et al 1998, Reid et al 1987), indicative of a role for a linear process in shaping direction tuning. Space-time tilt is likely conferred by appropriately arranged non-directional inputs that are time lagged relative to one another (Adelson & Bergen 1985). The required time lags could be produced by the convergence of magnocellular and parvocellular inputs from the LGN (De Valois et al 2000, Saul & Humphrey 1990). Alternatively, time

lags could be conferred through cortical networks (Maex & Orban 1996, Suarez et al 1995), synaptic depression (Chance et al 1998), or through delays due to dendritic propagation (Livingstone 1998, but see Anderson et al 1999).

Whether an additional directionally tuned signal exists to suppress responses to stimuli moving in the direction opposite the preferred in V1 is unclear. Most comparisons of direction tuning before and after the administration of the GABA antagonist bicuculline have found a decrease in direction selectivity upon blocking of inhibition (Murthy & Humphrey 1999, Sato et al 1995, Sillito 1975, Sillito et al 1980, Sillito et al 1985, Sillito & Versiani 1977). However, cooling of the cortex to eliminate cortical processing is reported to have no effect on direction selectivity, even though projections from the LGN to V1 are exclusively excitatory (Ferster et al 1996).

In addition to a putative directionally tuned subtractive signal in V1, an untuned divisive suppressive signal exists in visual cortical neurons (Carandini & Heeger 1994, Heeger 1992b). As described in section 1.2.1, divisive normalization has been studied most thoroughly in orientation tuned simple cells where it has been introduced to simultaneously describe response saturation to signals at high contrasts and cross-orientation inhibition or masking (Carandini et al 1997). The question of whether a divisive signal can account for direction tuning in V1 has not systematically been explored.

Projections from V1 to MT arise primarily from layer 4B spiny stellate neurons (Shipp & Zeki 1989). Layer 4B spiny stellate cells receive their inputs

primarily from layer 4Cα neurons which in turn receive projections from magnocellular layers of the LGN (Yabuta et al 2001). A second direct projection from V1 to MT arises from the large Meynert V1 neurons located at the layer 5/6 border (Lund et al 1976, Spatz 1977, Tigges et al 1981). Although a range of direction selectivities are found across V1 neurons, directionally tuned V1 signals appear to form the majority of the direct input from V1 to MT (Movshon & Newsome 1996). V1 signals can also reach MT indirectly through projections that first pass through areas V2 and V3.

Within MT, most neurons are strongly directionally selective and are suppressed below baseline firing rate by a non-preferred moving stimulus (Felleman & Kaas 1984, Maunsell & Van Essen 1983, Rodman & Albright 1987). Potentially, a suppressive computation could occur in MT to sharpen directional responses. Alternatively, direction selectivity could be inherited exclusively form the responses of V1 neurons. An examination of the role of suppression in shaping directional responses in V1 and MT is the subject of chapter 4.

2 Reliability of developing visual cortical neurons

The spatial vision of infant primates is poor; in particular, infant monkeys and humans are 5-10 times less sensitive to contrast than adults (Banks and Salapatek, 1981; Boothe et al., 1988). The visually evoked responses of cortical neurons in infant monkeys are relatively weak, and during development firing rates increase, receptive fields become smaller, and temporal resolution improves (Blakemore, 1990; Boothe et al., 1988; Chino et al., 1997; Wiesel and Hubel, 1974). It is commonly believed that the postnatal increase in visual sensitivity reflects postnatal maturation of visual cortical response properties.

However, it is not only the absolute firing rate that determines how accurately a neuron can signal the presence or character of a particular stimulus. Information in a neuronal response is limited not only by firing rate but also by variability. Presented with the same stimulus on repeated trials, a neuron responds with a variable number of spikes. If there were a constant relationship between variability and firing rate throughout development, the low firing rates of infant neurons would imply that the information they can transmit increases with age. However, if the variability of responses in infant neurons were lower, this might compensate for their lower spike rates and permit them to transmit more information than their sluggish responses might suggest.

We wanted to determine whether changes in firing rate and tuning properties observed during development are associated with an increase in the information content of the visual signals carried by cortical neurons. To quantify the efficiency with which neurons signaled information during different stages of development, we calculated two measures: a ratio of the variance to mean spike count, and an information theory-based measure that relates the amount of information in a response to the number of spikes used to convey that information. Both measures suggested that the responses of infant neurons were more reliable than those of adult neurons, and that the increase in responsiveness during development is paralleled by a decrease in reliability. Therefore, the information that infant cortical neurons transmit need not, by itself, limit the contrast sensitivity of infant vision.

2.1 Methods

We made single unit recordings from the primary visual cortex of 11 anaesthetized, paralyzed pigtail macaques (*M. nemestrina)* between 1 and 99 weeks of age, using conventional methods that are detailed in the Appendix.

After isolating each recorded neuron, we tested the more effective eye, and optimized the orientation, spatial frequency, temporal frequency, and area of drifting achromatic sinusoidal gratings of 0.5 contrast presented on a gray background. The time- and space-average luminance of the display was 33

cd/m^2. We then measured each neuron's response to gratings at six contrasts ranging from 0 to 0.5. Stimuli drifted across the screen at a rate chosen so that an integer (1-8) number of cycles occurred in a 640 msec period (1.6 to 12.5 Hz). For the neurons reported here, 10 or more 640 msec trials were collected for each contrast; stimuli were interleaved and presented in pseudorandom order. The f1/f0 ratio of the response to drifting gratings was used to classify cells as simple or complex (Skottun et al., 1991). A few simple cells with a high spontaneous rate were excluded from the analysis because spike-count based techniques do not correctly capture the information these neurons transmit.

A direct method was used to calculate the information about contrast (see e.g. Cover and Thomas, 1991 for a review on Information Theory), as the difference between the total entropy across all contrasts and the mean noise entropy at each contrast:

$$I = -\sum_{r} P(r)\log_2 P(r) + \sum_{s} P(s)\sum_{r} P(r \mid s)\log_2 P(r \mid s)$$

where *r* is the number of spikes in a 640 msec trial and *s* is the contrast level of the grating. This equation was used to calculate both the full mutual information (about six contrast levels) and pairwise information (about two contrast levels). To compensate for overestimation of information caused by the limited number of available trials (mean N = 28.5 cycles), we applied an analytical correction. When the number of trials was less than four times the peak spike count, the responses were quantized into R bins (Panzeri and Treves, 1996) with R chosen

such that convergence to the large-N asymptote was observed over the entire data set (this resulted in R=0.4N in the case of contrast pairs). The effect of this strategy is to trade some small degree of under-estimation due to quantization loss against over-estimation due to sampling bias in order to obtain the most accurate results over the entire data set. The analysis was also performed with fixed bin size, and qualitatively identical results obtained.

We devised a novel metric to compute reliability by relating the pairwise information available in stimulus-evoked responses to differences in spike rates; we will refer to this metric as *information density*. To calculate information density, mutual information was calculated about all possible pairs of contrasts (6 contrasts; 15 pairs) from spike counts in 640 msec bins. Because information was calculated about pairs of contrasts, information could be plotted against the difference in firing rates, which should be related to information, rather than a potentially less correlated measure such as the mean rate. The relation between mutual information and the difference in spike count was fit with the curve:

$$I = [1 - (1 - \alpha)^{(\Delta n)^{\beta}}] \log_2 S$$

where I is the mutual information, Δn is the difference in spike count, S is the number of stimuli (2), and α and β are free parameters. This curve asymptotes at the theoretical limit of I=1 bit for large values of Δn. For β=1, the curve corresponds to an exponential saturation model in which the information provided by each spike has a random overlap with that provided by any other; in

this case, α measures the extent of that overlap (Gawne and Richmond, 1993; Rolls et al., 1997a). For β=2, the curve corresponds to the rate at which information grows as the firing rate distributions for two stimuli are separated, if those distributions were Gaussian. Allowing β to vary allows the function to account for a variety of firing rate distributions; the value of β for our sample varied between 1 and 4. The maximum slope of this function represents the peak rate of information growth with difference in spike count; we term this quantity *information density* to distinguish it from other measures of information. The values of information density obtained by fitting other empirically-chosen functions were very similar to those obtained using Equation 2. Neurons were excluded from this and other analyses if the correlation between pairwise mutual information and spike count did not achieve significance on an F-test ($P < 0.05$). The number of neurons so excluded was small (1 week, 1 of 48 neurons; 4 week, 7 of 60 neurons; 16 week, 2 of 68 neurons; adult, 6 of 72 neurons).

We wanted to know whether the choice of test contrasts had an effect on the full (all stimuli) mutual information values we computed. In particular, if contrast values were placed too high or too low, most responses would be either small or large, skewing the distribution of responses and reducing the amount of information transmitted. We calculated full mutual information for a Poisson neuron with a conventional contrast-response function and deliberately skewing the chosen contrast values. The full mutual information measure proved quite insensitive to this skewing within the range of skews in our data set, and we

used the simulations to estimate the amount by which our full mutual information calculations would have been in error for real neurons. The effect of skewing was modest (<10% underestimate of information for almost all cases), and there was no difference in our estimated errors across the four age groups.

We also measured responses to high contrast gratings of optimal orientation and spatial frequency drifting at frequencies between 0.4 and 25 Hz, and fit the data with a suitable descriptive function; we took temporal resolution as the frequency at which the response fell to one-tenth of maximum (Foster et al., 1985; Saul and Humphrey, 1992). Response latency also provides a measure of integration time (Gonzalez et al., 2001; Maunsell and Gibson, 1992). We measured response latency by plotting response histograms (in 5 msec bins) over multiple data sets and estimating latency as the first bin in which the response was greater than the mean spontaneous rate measured in response to a gray screen. For simple cells, a latency was recorded only if cycle triggered averages indicated that at least one stimulus started in the cell's excitatory phase. For a few cells (19 of 232), we could not determine latency reliably and those cells were omitted from the latency analysis.

2.2 Results

Consider the two cells whose data are shown in figure 2-1. Figure 2-1, a and b, shows the mean responses of an infant and an adult neuron, respectively, to an

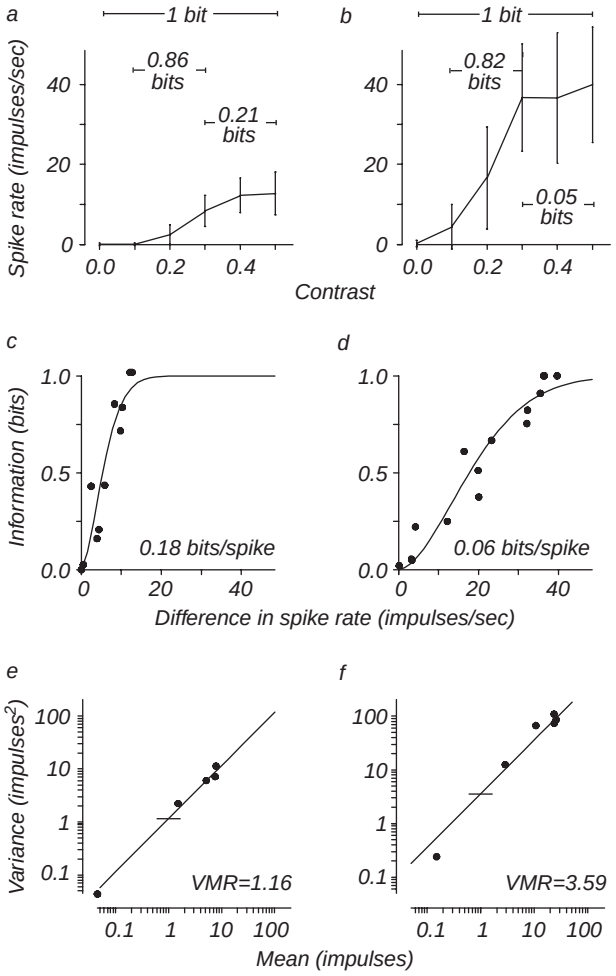
a
1 bit
0.86 bits
0.21 bits
Spike rate (impulses/sec)
40
20
0
0.0
0.2
0.4
b
1 bit
0.82 bits
0.05 bits
Contrast
c
Information (bits)
1.0
0.5
0.0
0.18 bits/spike
0
20
40
d
0.06 bits/spike
Difference in spike rate (impulses/sec)
e
Variance (impulses^2)
100
10
1
0.1
VMR=1.16
0.1
1
10
100
f
VMR=3.59
Mean (impulses)

Figure 2-1. Calculation of information density and variance to mean ratio for two cells. ***a, b*** Mean and standard deviation of the responses of a neuron from a 4-week old infant (a) and from an adult (b) to an optimized, drifting sinusoidal grating stimulus at six different contrasts evenly spaced between 0 and 0.5. The mutual information about selected contrast pairs is indicated. ***c, d*** Mutual information about every possible pair of the six contrasts (15 pairs) in a and b is plotted against the difference in the mean firing rate between each pair of contrasts. These data are fit with a function whose maximal slope is a measure of information density (see *Methods*). Information density has units of bits per spike, and the computed information densities for each cell are indicated. This measure, unlike total mutual information, does not depend on the specific contrasts tested, which differed somewhat from cell to cell. ***e, f*** Spike count variance at each contrast is plotted against mean spike count for the example cells in a and b. The variance to mean ratio (VMR) is taken from the best fitting line with slope = 1; horizontal ticks mark the ratios for each cell. The counting window was 640 msec and contained an integer number of temporal cycles of the drifting stimulus.

otherwise optimal grating stimulus at six different contrasts; error bars indicate the standard deviation of the firing rate distributions. As is typical of visual cortical neurons, firing rate grew with contrast and saturated at high contrasts for both cells. To discriminate two stimuli perfectly, a neuron with high trial-to-trial variability like the adult cell must signal two different stimuli with very different mean firing rates. Conversely, a neuron with low variability like the infant cell can convey the same amount of information with a smaller dynamic range.

We used Shannon's mutual information to measure how accurately different stimuli can be distinguished, based upon the number of spikes elicited from a neuron during repeated stimulus presentations (Rolls et al., 1997b; Tolhurst, 1989; Werner and Mountcastle, 1965). The information is related to the distance between the two firing rate distributions, and is similar to the *d'* measure used in signal detection theory (Parker and Newsome, 1998). To illustrate the relationship between the firing rate and information, figure 2-1, a and b, also shows the information transmitted by each neuron about selected pairs of contrasts. Note that both the infant and adult neuron were capable of perfectly discriminating a zero contrast stimulus (mean gray background) from the highest stimulus contrast, yielding 1 bit of information. However, the infant neuron signaled this information with fewer spikes.

To quantify the relationship between information and the number of spikes needed to convey that information, we plotted the information conveyed

by a neuron about each of the 15 different contrast pairs against the mean firing rate difference between the members of each pair (Figure 2-1c, d). Information about a contrast pair cannot exceed one bit, representing perfect discrimination, and we therefore fit these points with a curve whose form accounts for this saturation. The maximum slope of this function captures the shape of the relation between information and spikes; we term the maximum slope of this curve the *information density* (see *Methods)*, with units of bits per spike. This measure differs from the more usual full mutual information in that it depends only on pair comparisons and not on the total number of stimuli used (Rolls et al., 1997b; Tolhurst, 1989). Neurons with larger values of information density use fewer spikes to convey information (Figure 2-1c). Neurons with smaller values require a larger dynamic range to discriminate contrast pairs (Figure 2-1d).

Another way to capture the change in firing patterns is to analyze the relationship between response mean and variance for the example cells. The variance of cortical neuron spike counts increases in proportion to their mean (Tolhurst et al., 1983; Tolhurst et al., 1981) and the ratio of the two is inversely related to the amount of information transmitted by cortical cells (de Ruyter van Steveninck et al., 1997). Figure 2-1, e and f, shows the relation between response variance and mean for the two example cells. As indicated by the reference lines at a spike count of 1, the infant cell had a lower variance to mean

ratio than the adult cell, as would be expected from its higher information density.

We calculated information density for populations of V1 cells recorded from macaques in four age groups: 1, 4, and 16 weeks, and adult (31-99 weeks). Surprisingly, we found that V1 neurons in the youngest animals had the highest information density: mean information density decreased two-fold during development (Figure 2-2a). We also calculated the variance to mean ratio for the same populations; as expected from the information density calculation, the variance to mean ratio of cortical cells increased during development (Figure 2-2b). Adult cells tended to have higher variance to mean ratios than infant cells even when cells with similar dynamic range were selected, implying that this developmental difference cannot be attributed to the subpopulation of adult cells with high firing rates (data not shown). It is also interesting to note that simple cells had higher information densities for each age group (mean information densities for simple cells from the 1-week, 4-week, 16-week, and adult animals were 0.33, 0.25, 0.20, and 0.12; for complex cells the values were 0.19, 0.15, 0.11, and 0.09, respectively); simple cells had correspondingly lower variance-to-mean ratios than complex cells. A multiple linear regression analysis suggests that these differences cannot be accounted for by differences in spontaneous rate or dynamic range.

Together, these two measures suggest that the coding properties of neurons change during development. How are they related? Figure 2-2c shows

that information density and the variance to mean ratio were inversely but imperfectly correlated. This is because the variance to mean ratio measures the average variability of the response to a single stimulus, while the mutual information quantifies the fraction of the total variability that is attributable to the difference between responses. These two measures are comparable in that each indicates the reliability of neuronal firing, and the regular relationship shown in Figure 2-2c suggests that during development there was a decrease in the reliability of visual signaling by cortical neurons.

Despite the decrease in reliability during development, total information transmission could be maintained if the range between the lowest and highest firing rates (the dynamic range) also increased. The mean dynamic range did indeed increase two-fold during development, and a plot of the mean information density versus the geometric mean evoked firing rate for each age reveals the reciprocal relationship between these two measures (Figure 2-3a). In the youngest infants, information density was high and firing rate was low, whereas in the adults information density was low and firing rate was high.

The mutual information about all of the 6 contrasts presented in an experiment (which we term “full” mutual information to avoid confusion with the pairwise measure) quantifies the ability of these neurons to distinguish stimuli, and depends on both information density and dynamic range. However, unlike information density, full mutual information depends on both the number and distribution of the contrasts tested. We did not always use the same test

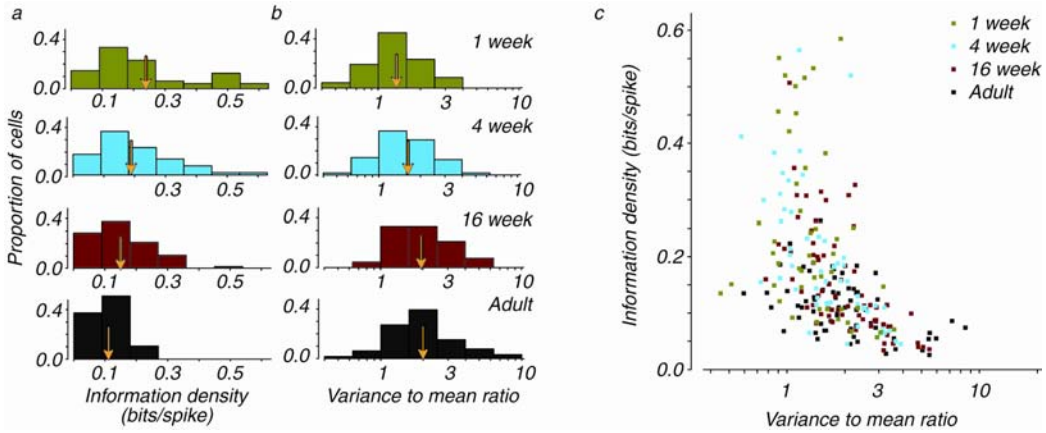

Figure 2-2. Changes in information density and the variance to mean ratio during development. *a* Distributions of information density for neurons from monkeys in the four age groups (see *Methods* for calculation). Arrows indicate the means. *b* Distributions of the variance to mean ratio for each age group (see *Methods* for calculation). Arrows indicate the geometric means. *c* Scatter plot of the data displayed in a and b, for 232 neurons from animals in the four age groups: 1 week (47), 4 week (53), 16 week (66), adult (66).

contrasts because we tried to place the contrasts so that they spanned the response range of each cell, but we verified that the chosen contrasts did not have an important effect on the full mutual information measure for our population (see *Methods*). Mean full information values for the four age groups are given next to each point in Figure 2-3a. The modest and inconsistent change in the full mutual information values is due to the opposing effects of increasing firing rate and decreasing information density as development progresses. In other words, infant neurons may fire few spikes, but each infant spike carries more information. As a result, 1week infant neurons can transmit 80% of the total information adult neurons transmit.

2.3 Discussion

Our results suggest that lower firing rates in infant neurons are partially compensated for by lower variability and that infant neurons, therefore, are more efficient at transmitting information about contrast than adult neurons. This leads to an interesting quandary. If infant neurons are capable of signaling 80% of the information that adult neurons signal, why is it that contrast sensitivity in infant primates is 5-10 fold lower than in adults (Boothe et al., 1988)? One possibility is that infant neurons have higher contrast thresholds than adult neurons (compare responses in Figure 2-1a,b). Our results might have been different had we tested infant neurons with very low contrast targets, but we did not explore systematically the contrast range below 0.1. A second possibility is

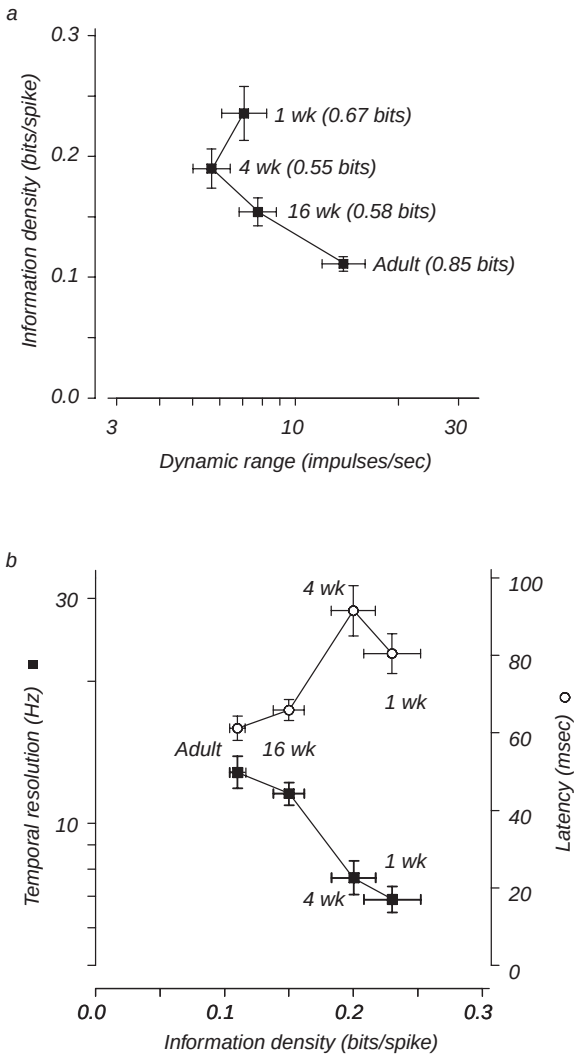

Figure 2-3. The relationship between information density, dynamic range, and temporal parameters during development. *a*, Mean information density and geometric mean dynamic range are plotted for each age group. Dynamic range is taken as the largest mean response to a grating target minus the mean baseline response. The mean transmitted full mutual information for all 6 contrasts is indicated beside each point. *b*, Mean information density, geometric mean temporal resolution (solid squares), and geometric mean latency (open circles) are plotted for each age group. For each cell, temporal resolution was taken as the drift rate at which the cell's response fell to one-tenth of its peak. Latency was taken as the time after stimulus onset at which the firing rate first deviated from baseline. Standard errors are plotted for all axes.

that the limits to infant contrast sensitivity are not set by V1 neurons and, instead, lie in downstream structures (Kiorpes and Movshon, 2003). The low spike rates of infant neurons might contribute to this by driving downstream neurons less effectively, even if their responses are reliable.

How might the reciprocal relationship between information density and firing rate arise? Many aspects of the visual system change during development, including improvements in the optics of the eye (Jacobs and Blakemore, 1988; Williams and Boothe, 1981) migration of cones in the fovea (Packer et al., 1990), increases in spatial resolution and decreases in receptive field size (Blakemore, 1990; Chino et al., 1997; Movshon and Kiorpes, 1993; Movshon et al., 2000). Our first thought was that developmental decreases in receptive field size might underlie our observations, but we have shown that these changes are almost entirely attributable to changes in retinal optical magnification and cone distribution (Movshon et al., 2000; Wilson, 1993), and do not reflect neural changes in receptive field organization. However, there are marked changes in the temporal fidelity of responses during development that may drive the change in information density. Figure 2-3b plots mean information densities for the neurons from each of the four age groups against two temporal measures: the latency of response after stimulus onset and the highest temporal frequency of drift that elicited a response (temporal resolution). A relationship between information density and each of these temporal parameters is clear. The decrease in latency and increase in temporal resolution with age suggest that

infant neurons integrate their inputs over longer times than adult neurons. A neuron with a longer integration time would average over more synaptic input events and thus reduce variability associated with rapid fluctuations in those inputs; such a neuron would carry more information with each spike by sacrificing temporal bandwidth. To improve their resolution of fine temporal structure, developing V1 neurons decrease their integration times, which would increase the variability of spiking. Such an increase would increase variance-to-mean ratios and have a deleterious effect on information transmission, but these effects could be overcome by increasing dynamic range (Figure 2-3a).

Developmental changes in temporal integration might arise from changes in either neuronal properties or synaptic properties. Interestingly, in the gerbil lateral superior olive and rat cortex, EPSPs are of longer duration in infant than adult neurons (Burgard and Hablitz, 1993; Sanes, 1993); this change may be due to changes in patterns of glutamate receptor expression (Krukowski and Miller, 2001). Whatever the biological basis, a shift in coding strategy from high information density, low bandwidth, and low firing rate to low information density, high bandwidth, and high firing rate would ensure that information transmission is not sacrificed as temporal resolution grows to adult levels.

3 Spike-triggered covariance reveals unexpected substructure in V1 simple and complex cells

To understand the processing that occurs at each stage of a sensory system, it is useful to develop models that describe the transformation of stimuli into neuronal firing patterns. The primary purpose of this type of "functional" model is not to provide a biophysical description, but rather to provide a compact yet precise description of the computation that transforms a stimulus into neuronal response. Although the neurons in primary visual cortex (area V1) have been studied extensively, functional descriptions of these neurons remain incomplete.

V1 neurons are commonly classified in two categories, based on their responses to light and dark bars (Hubel & Wiesel 1962). Simple cells have distinct 'on' and 'off' regions which are excited by bars of the same polarity and suppressed by bars of the opposite polarity. This behavior is captured in the typical model of a simple cell, which is based on a single linear filter with alternating positive and negative lobes (Hubel & Wiesel 1962, Movshon et al 1978b). Complex cells respond to light and dark bars in a manner that is independent of the polarity of the stimulus, suggestive of overlapping 'on' and 'off' regions (Hubel & Wiesel 1962, Movshon et al 1978a). These cells are commonly described using an "energy model", in which the outputs of two phase-shifted linear filters are squared and summed (Adelson & Bergen 1985). Although these models capture the main characteristics of V1 neuronal response

properties, they fail to capture other behaviors observed in these neurons. For instance, it is common to encounter neurons whose phase sensitivity lies in between the two extremes predicted by the standard models of simple and complex cells. Furthermore, suppressive influences such as those suggested by cross-orientation suppression, contrast gain control and other forms of adaptation are not included in these models.

In order to better account for the full range of behaviors found in these cells, we have developed a generalization of the standard V1 models (figure 3-1). In this model, the stimulus is analyzed using a small set of linear filters whose outputs are combined nonlinearly in order to determine the firing rate. The number of filters is allowed to vary, and the nonlinear combination is unconstrained, which allows (for example) the influence of individual filters to be either excitatory or suppressive.

This type of model can be fit and validated with experimental data through a spike-triggered analysis (Simoncelli et al 2004). In a spike-triggered characterization, neural responses to a sequence of random stimuli are recorded, and the ensemble of stimulus blocks that precede spikes analyzed to determine both the linear filters and the nonlinear rule by which their outputs are combined. The most well known method of this kind, known as reverse correlation, has been widely used to characterize simple cells (DeAngelis et al 1993, Jones & Palmer 1987). Assuming a model based on a single linear filter followed by a rectifying nonlinearity, an unbiased estimate of the linear filter can be recovered

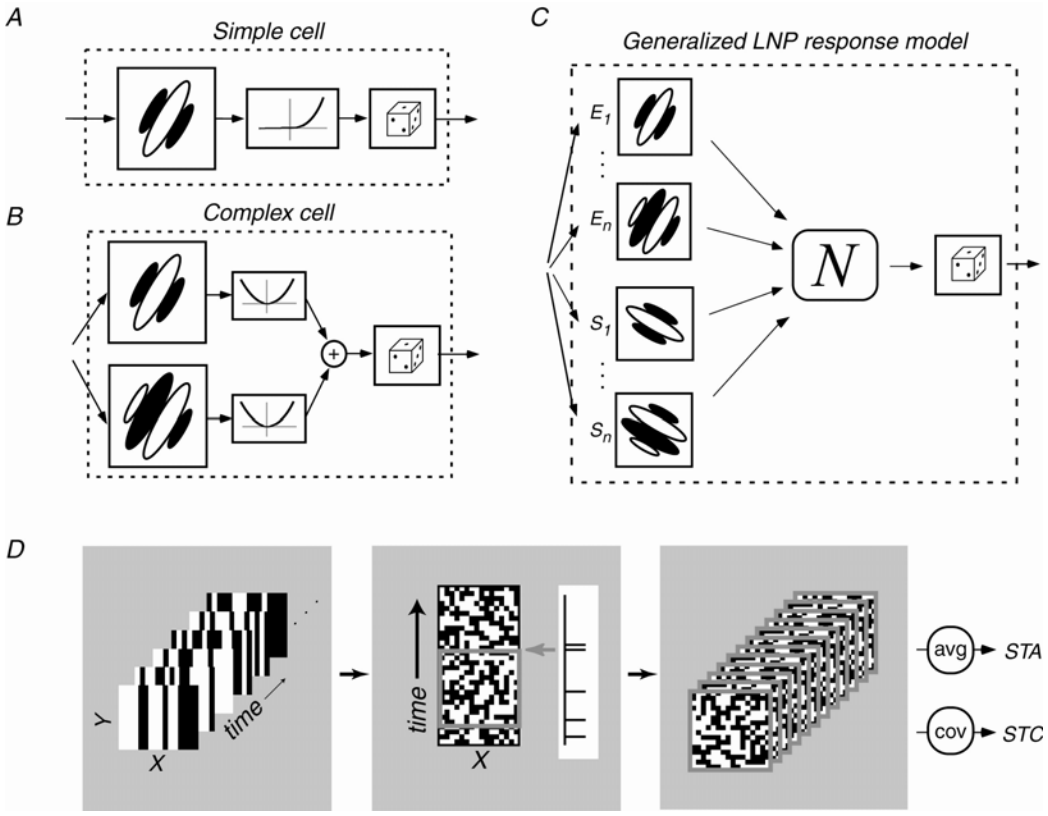
A
Simple cell
B
Complex cell
C
Generalized LNP response model
E_1
E_n
S_1
S_n
N
D
y
x
time
time
X
avg
STA
cov
STC

Figure 3-1. LNP functional models for V1 neurons, and their characterization using spike-triggered analyses. A) A standard simple cell model, based on a single space-time oriented filter. The stimulus is convolved with the filter and the output is passed through a halfwave-rectifying and squaring nonlinearity. This signal determines the instantaneous rate of a Poisson spike generator. B) The "energy model" of a complex cell, based on a pair of space-time oriented filters with a quadrature (90 degree) phase relationship (Adelson & Bergen 1985). Each filter is convolved with the stimulus, and the responses are squared and summed. The resulting signal is used to drive a Poisson spike generator. C) The generalized linear-nonlinear-Poisson (LNP) response model used in this paper. The cell is described by a set of *n* linear filters (L), which can be excitatory (E) or suppressive (S). The model response is computed by first convolving each of the filters with the stimulus. An instantaneous nonlinearity (N) governs the combination of excitatory and suppressive signals. Finally, spikes are produced via a Poisson spike generator (P). D) Spike-triggered analysis. Left panel: a random binary bar stimulus used to drive V1 neurons. The bars were aligned with the neuron's preferred orientation axis and the stimulus array spanned the classical receptive field. Middle panel: An X-T slice of the stimulus sequence – each pixel represents the intensity of a bar at a particular location in one frame. The collection of stimulus blocks during the 16 frames (160 msec) before each spike (example in gray box) form the spike-triggered stimulus distribution. Right panel: The STA is a block of pixels, each corresponding to the average of the corresponding pixel values over the distribution. The STC is a matrix whose entries contain the average product of pair of pixel values (after the mean has been subtracted). See methods for details.

by taking the average of the stimuli preceding spikes (Chichilnisky 2001, Paninski 2003) The nonlinear function that describes the transformation of the filter output into a firing rate can also be reconstructed, thus completing a quantitative model that predicts the neuron's firing rate in response to any time-varying stimulus.

The nonlinear behavior of complex cells precludes their characterization by spike-triggered averaging. Specifically, consider the energy model (figure 3-1B). Due to the squaring operation applied to each filter's output, every stimulus that excites the cell a will be matched with a stimulus of opposite polarity that equally excites the cell, and the resulting STA will be flat. An extension to the spike-triggered averaging concept, known as spike-triggered covariance (Brenner et al 2000, de Ruyter van Steveninck & Bialek 1988), has been introduced as a means of resolving filters that have this type of symmetric nonlinear influence on a neuron's response. Once a set of excitatory and suppressive linear filters are recovered, the model is completed by estimating the nonlinear function that describes how the signals arising from each filter are combined to determine the firing rate. STC has proven successful in revealing the excitatory subunits of cat V1 complex cells (Touryan et al 2002) as well suppressive influences in the retina (Schwartz et al 2002).

3.1 Methods

We recorded from isolated single units in primary visual cortex (V1) of adult macaque male monkeys (*Macaca fascicularis* and *Macaca Nemestrina*) using methods that are described in the Appendix.

Stimuli were generated with a Silicon Graphics XX workstation and presented on a gamma-corrected monitor with a refresh rate of 100 Hz and a mean luminance of 33 cd/m^2. The monitor was directed toward the monkey via a front surface mirror; total length between the eye and the monitor was 165-180 cm.

Stimuli were presented monocularly. Upon encountering a cell, the initial characterization involved a determination of the best direction, spatial frequency, and temporal frequency of drifting grating stimuli. The size of the classical receptive field was defined as the size at which an optimized full contrast sinusoidal grating saturated the response without impinging upon the suppressive surround. Stimuli used in the spike-triggered characterization were extended temporal sequences in which each frame contained a set of parallel non-overlapping black and white bars with randomly assigned intensity. The orientation of the bars was aligned with the cell's preferred orientation and the stimulus array was confined to the classical receptive field. The number of bars (8-32) was chosen to match cell's preferred spatial frequency such that 4-16 bars fell within each spatial period. A new frame was displayed every 10 msec.

Spike-triggered analysis (recovery of the linear filters):

We give a brief description of our spike-triggered analysis (Simoncelli et al 2004). We define a spike-triggered stimulus block, $S_n(x,t)$, as the set of bar intensities in the 16 frames preceding the *n*th spike (figure 3-1D, middle panel). Each of these spike-triggered stimuli can be envisioned as a point in a *D*-dimensional space (*D* is 16 times the number of bars presented in each frame, and ranges from 64 to 512 in our experiments, with the component along each axis representing the intensity of the corresponding bar, relative to the mean stimulus intensity, at position x and time t before the spike).

In the conventional procedure of reverse correlation, one averages the spike-triggered stimulus blocks to obtain the spike-triggered average (STA):

$$STA(x,t) = \frac{1}{N}\sum_n S_n(x,t),$$

where *N* indicates the number of spikes. More specifically, if one assumes that the neural response is generated by convolution with a single linear filter followed by an instantaneous asymmetric (e.g. half-squaring) nonlinearity and Poisson spike generation, the STA provides an unbiased estimate of the linear filter (Chichilnisky 2001, Paninski 2003).

Despite its widespread use in estimating linear receptive fields, the STA is known to fail under several commonly occurring circumstances. For example, if the neural response is symmetric, as is commonly assumed in models of

complex cells, the STA will be zero. If the neuron depends upon more than a single axis within the stimulus space, the STA will be some weighted average of these axes, but will not provide any indication of the full set. In either case, the STA produces a misleading description of the receptive field properties of the neuron. In order to handle these situations, one must examine higher-order statistical properties of the spike-triggered stimulus distribution. A number of authors have examined the second-order statistics of spike-triggered stimuli by computing the spike-triggered covariance (Brenner et al 2000, de Ruyter van Steveninck & Bialek 1988, Paninski 2003, Schwartz et al 2002, Simoncelli et al 2004). This procedure amounts to examining the variance of the cloud of spike-triggered stimuli and identifying those axes in the stimulus space along which the variance differs significantly from that expected due to chance correlations between the stimulus and spikes. These axes correspond to a set of linear filters underlying the neuron's response.

In our analysis, we first compute the STA and remove this component from the stimulus distribution. Specifically, we compute the normalized (unit vector) STA, the *nSTA*, and define:

$$S_n'(x,t) = S_n(x,t) - \left[\sum_{x,t} S_n(x,t) \cdot nSTA(x,t) \right] \cdot nSTA(x,t).$$

This differs from a traditional covariance calculation (in which the STA would be subtracted from each $S_n(x,t)$), but ensures that the axes obtained in the STC

analysis will be orthogonal to the STA and helps to avoid unwanted interactions between the STA and STC analyses.

We then compute the *DxD* spike-triggered covariance:

$$COV(x_1, x_2, t_1, t_2) = \frac{1}{N_s - 1} \sum_n S'_n(x_1, t_1) S'_n(x_2, t_2).$$

This matrix, with the parameter pairs $\{x_1,t_1\}$ and $\{x_2,t_2\}$ specifying the row and column indices, fully represents the variance of the spike-triggered stimulus ensemble in all possible directions within the stimulus space. Geometrically, the surface swept out by a vector whose length is equal to the variance along its direction is a hyper-ellipse, and the principal axes of this hyper-ellipse, along with the variance along each axis, may be recovered using principal components analysis (PCA). More concretely, principal axes of this ellipse correspond to the eigenvectors of the covariance matrix, and the variance along each of these axes is equal to the corresponding eigenvalue.

In the absence of any relationship between the stimulus and the spikes (and in the limit of infinite data), the spike-triggered ensemble would just be a randomly selected subset of all stimuli, and the variance of this subset in any direction would be identical to that of the full stimulus set. In an experimental setting, the finiteness of the spike-triggered ensemble produces random fluctuations in the variances in different directions. We are interested in recovering those axes of the stimulus space along which the neuron's response leads to an

increase or decrease in the variance of the spike-triggered ensemble that is greater than what is expected from this random fluctuation due to finite sampling.

We tested a nested sequence of hypotheses to determine the number and identity of axes corresponding to significant increases or decreases in variance. We began by assuming that there were no such axes (i.e., that the neuron's response was independent of the stimulus). We used Monte Carlo simulation to compute the distribution of minimal and maximal variances under this hypothesis. Specifically, we randomly time-shifted the spike train relative to the stimulus sequence, performed our STA/STC analysis on the resulting spike-triggered stimulus ensemble, and extracted the minimum and maximum eigenvalues. Based on 500 such calculations, we estimated the 99% confidence interval for both the largest and smallest eigenvalues. We then asked whether the eigenvalues obtained from the true spike-triggered ensemble lay within this interval. If so, we concluded that the hypothesis was correct. Otherwise, we assumed the largest outlier (either the smallest or largest eigenvalue) had a corresponding axis (eigenvector) with a significant influence on neural response. We added this axis to a list of significant axes, and proceeded to test the hypothesis that all remaining axes were insignificant.

STC approaches are intended for use with Gaussian-distributed stimuli due to the circular symmetric properties of such distributions (Paninski 2003). Unfortunately, the low contrast of Gaussian stimuli leads to low firing rates in

V1 neurons and the STC characterization requires a large number of spikes. We were thus forced to use higher contrast binary stimuli. We noticed that in neurons with many excitatory subunits, we occasionally found suppressive filters that differed from the others in that they contained only a small number of highly correlated bars. Simulations confirm that these filters are artifactual, and are due to the use of binary distributed stimuli (figure 3-2). To provide a conservative estimate of the number of significant filters for V1 neurons, we applied an additional criteria to the filters deemed significant by the hypothesis test described above. For each filter, we computed a histogram of the values obtained from its Fourier amplitude spectra. Filters with diffuse frequency spectra produced compact distributions, whereas filters with regions of highly correlated energy produce distributions with long tails (figure 3-2). We differentiated these filters by comparing the fourth moment of these distributions with a threshold value. The threshold was determined by generating artifactual filters from repeated simulations of model LNP neurons and estimating the 95% confidence interval of the distribution of fourth moments expected due to artifacts produced by random stimuli.

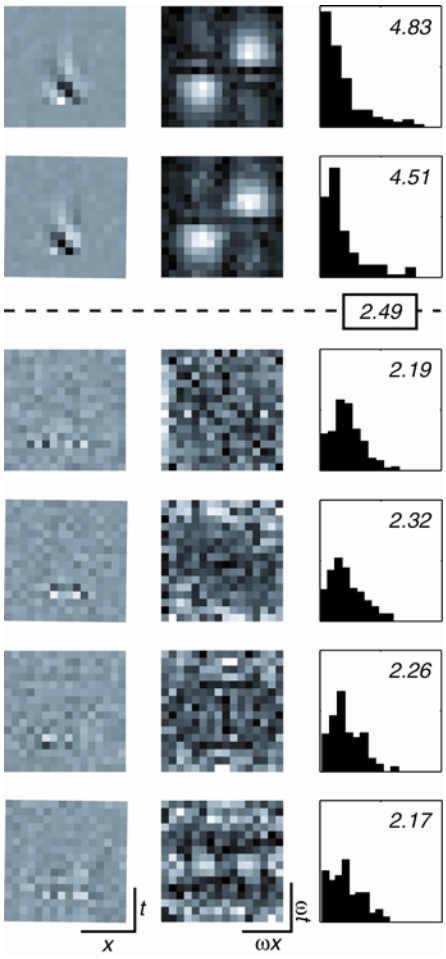

Figure 3-2: Artifactual suppressive filters produced by binary (non-gaussian) stimuli. Shown are two real and four artifactual filters revealed in a simulated complex cell along with their amplitude spectra. The artifactual filters can be identified by the absence of clear spatiotemporal structure and small number of correlated pixels, or equivalently their diffuse amplitude spectra. Histograms of the values taken from the amplitude spectra are shown; diffuse spectra have compact distributions whereas spectra with clear spatiotemporal tuning have distributions with long tails. Artifactual filters were identified by computing the fourth moment of these distributions, labeled for each filter. A large number of artifactual filters were collected from simulations of model neurons and the fourth moment of their spectra computed. In our data, we only considered filters with values greater than the 95% confidence interval of this distribution (2.49) to be significant.

Estimating the nonlinearity:

The firing rate of the neuron is a nonlinear function of the outputs of the set of linear filters recovered from the spike-triggered analysis. It is possible to estimate this function directly by binning the filter outputs and estimating firing rate for each bin (Brenner et al 2000, Chichilnisky 2001). For example, we can examine the structure of the nonlinearity as a function of the output of any single filter by taking the quotient of the number of spikes and the number of stimuli for each (binned) filter output value (figure 3-7A-C, marginals). Similarly, firing rate as a function of the responses of two filters may be examined by taking the quotient of the joint (two-dimensional) counts of the number of spikes and the number of stimuli (figure 3-7A-C). Unfortunately, the data required for such a direct estimate grows exponentially with the number of filters. For example, to estimate this function for a neuron with 10 significant filters over a set of 15 bins along each axis would require collecting multiple samples in 15^10 bins. Thus, for the neurons in our study, we needed to somehow reduce the dimensionality of the problem in order to estimate the nonlinearity.

We found that contours of constant firing rate associated with the nonlinearity for any pair of excitatory or suppressive filters were well fit by ellipses, with vertical/horizontal principal axes. Based on this observation, we decided to define the firing rate nonlinearity in two stages (figure 3-7D). In the first stage, the excitatory and suppressive filter outputs were combined in two

separate pools via weighted sums of squares (with the STA half-squared). The weights for each pool were obtained by maximizing the mutual information between the weighted sum-squared output of the joint excitatory and suppressive pools and the spikes. This approach makes no assumptions regarding the mathematical form of the interaction between excitation and suppression.

The second stage of the nonlinearity computed firing rate as a function of the joint output of the pooled excitatory and suppressive responses, and was estimated by taking the quotient of the two-dimensional binned counts of the number of spikes and the total number of stimuli. Because the data were not uniformly distributed across this two-dimensional space, individual bin widths were adjusted to maintain a uniform distribution of data points across the marginals. For neurons with no suppressive axes, the second stage nonlinearity described firing rate as a function of the output of the pooled excitatory response alone.

Parametric models for the interaction of excitatory and suppressive:

We fit several simple parametric models to the binned second-stage nonlinear firing rate function. The variable bin sizes used to evenly distribute the data across these surfaces resulted in an unequal distribution of data within each bin. For fitting purposes, we assumed a value of excitation (E) and suppression (S) for each bin equal to the center of mass of the data in that bin. We determined

the best fitting model to be one with a sigmoidal excitatory function containing a scalar (*a*), a baseline (β), an exponent (*p*), and weight (*c*) for the normalization term. The suppression was allowed to have both a subtractive (b) and divisive (d) influence on the response:

$$R = \beta + (aE^{p} - bS^{p})/(cE^{p} + dS^{p} + 1)$$

The model was fit using the STEPIT algorithm (Chandler 1969) to minimize the mean-squared error between the actual responses and model predictions.

Classification of V1 neuronal types:

Cells were classified as simple or complex based upon their response to full contrast drifting sinusoidal gratings optimized for direction, spatial frequency, temporal frequency, position and size. Gratings were presented for an integer number of cycles with the first cycle removed to eliminate effects from the onset transient. The relative modulation index (F1/DC) was quantified as the ratio of the vector average response at the grating temporal frequency and the baseline subtracted mean response. Baseline was defined as the response to a blank (mean gray) screen. Post-stimulus time histograms to one cycle of a drifting grating (cycle-triggered averages) were constructed by binning time (relative to the stimulus cycle) with a resolution of 10 msec, the duration of one frame.

Cells were classified as directional or non-directional based upon a comparison of their responses to grating stimuli drifting in the direction that produced the largest response versus the direction opposite. Specifically, a directional index was calculated as: 1 – (nonpreferred response / preferred response) where both responses were baseline subtracted. Cells with an index greater than 0.8 were considered directional and those with an index less than 0.8 non-directional.

Predicting the responses to arbitrary stimuli:

Determination of the predicted response of the STA-based model began by convolving the stimulus with the STA. The nonlinear function that relates the output of the STA to firing rate was reconstructed as described above and fit with the exponentiated, threshold-rectified function:

$$R = k * \lfloor x - T \rfloor^n$$

to correct for discretization by binning and to better estimate poorly sampled bins. The predicted F1, DC, and blank responses were calculated directly from the firing rate functions produced by the model.

Determination of the predicted response of the energy model containing the first two filters revealed by STC began by convolving the grating stimulus with the two filters and combining the signals via a weighted sum of squares.

The relationship between firing rate and the output of these filters were was fit with an appropriate function to correct for discretization and the resulting function was used to transform the pooled excitatory signal into a firing rate prediction. The predicted response modulation was calculated in the same manner as for STA-based model prediction.

Determination of the predicted response of the full model containing the STA and filters revealed from the STC began by convolving the grating stimulus with each of excitatory and suppressive filters recovered for a neuron. The excitatory and suppressive signals were combined separately, each via a weighted sum of squares. Firing rate was computed from the pooled *E* and *S* signals by using the combination model fit to the data (see above). Upon obtaining firing rate predictions, the predicted response modulation was calculated as described for the STA.

While sinusoidal gratings were of a similar contrast to the bar stimulus used to characterize the neuron, stimuli presented at lower contrasts (e.g. figure 3-10D) required a gain parameter to adjust the contrast sensitivity of the cell. To determine the gain adjustment for a neuron, we fit a single scalar to the pooled excitatory and suppressive signals before these signals were converted into firing rates.

3.2 Results

Recovery of the linear filters:

We stimulated each neuron with a dense, random binary bar stimulus aligned with its preferred orientation axis (figure 3-1D, left). As described above, we assume a functional model consisting of a set of linear filters (L), followed by an instantaneous nonlinear function (N) that combines their outputs to obtain a rate and a Poisson spike generator (P; figure 3-1C). The model posits that the generation of spikes is based on the stimulus contained in an interval preceding each spike. We chose a duration of 16 stimulus frames (160 msec) for this interval (figure 3-1D, middle). The ensemble of stimulus blocks preceding each spike define the spike-triggered stimulus distribution (figure 3-1D, right). The linear filters for each neuron were recovered from the statistics of this distribution. For every cell, we began by calculating the first-order statistic (the mean) of this distribution, the spike-triggered average (STA). We then recovered additional filters by calculating the second-order statistics of the STSD, the spike-triggered covariance matrix, and resolving this matrix into a small set of filters through the application of principal components analysis (STC). Specifically, components associated with significant increases or decreases in the variance of the spike-triggered stimulus ensemble (relative to the variance of the raw stimuli) provide estimates of the filters in the LNP model (see methods for details).

This procedure recovers a set of linear filters that form the front end of the LNP model, and thus define the fundamental stimulus selectivity of the cell (more formally, they determine the linear subspace of the stimulus space in which the cell's response is generated). Note that signals that are combined in a purely linear fashion, such as the excitatory and inhibitory signals arising from positive and negative subregions of a receptive field, will be resolved into a single filter by this analysis. Signals that are combined after a nonlinear operation, such as rectification or squaring, will be revealed as different filters. Each recovered filter can have an excitatory or a suppressive impact on spiking, depending on the way in which its output is incorporated into the nonlinear stage. The individual filters are unique only up to a linear transformation (i.e., one can form an equivalent model based on an alternative set of filters that are related by an invertible linear transformation), and thus cannot be taken too literally as an indication of underlying mechanisms. Nevertheless, the overall subspace they cover is uniquely determined, as is the full LNP response model.

The spatiotemporal filters for a representative simple cell, along with their Fourier spectra, are shown in figure 3-3A. This cell produced an STA with clear spatiotemporal structure. The space-time tilt of this filter, or equivalently the localization of spatiotemporal energy in opposite quadrants of the Fourier domain, indicates a preference for the direction of a moving stimulus. If this simple cell were adequately described by a single linear filter (as suggested by the standard model of figure 3-1A), no additional filters would be revealed by

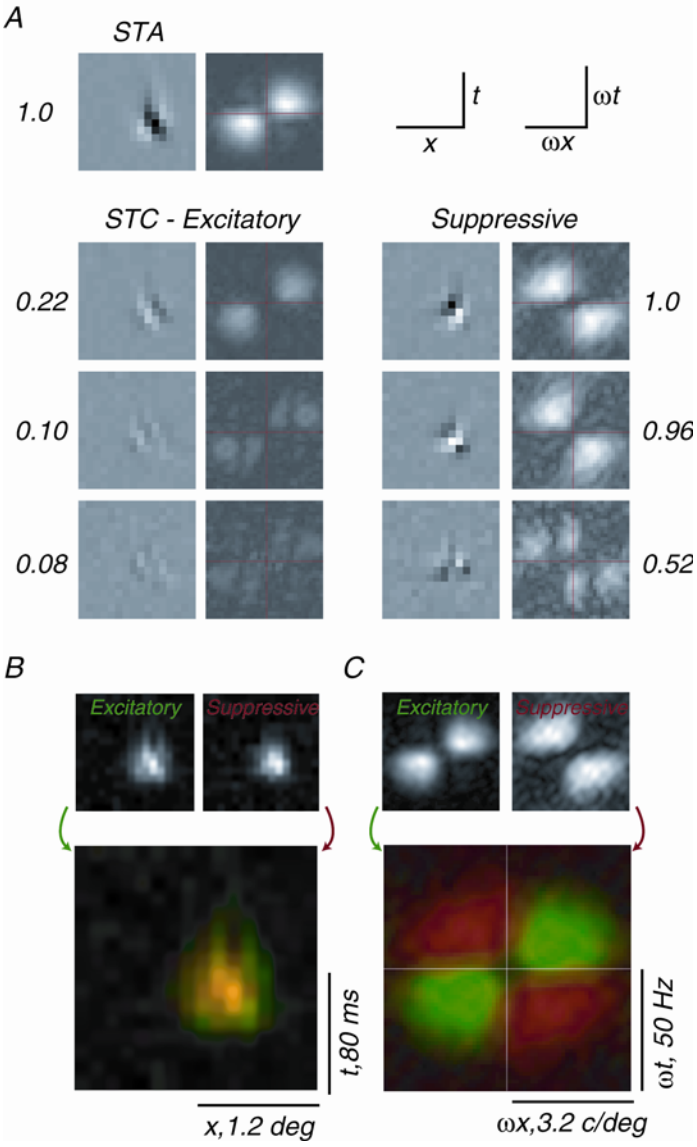
A
STA
1.0
t
x
ωt
ωx
STC - Excitatory
Suppressive
0.22
0.10
0.08
1.0
0.96
0.52
B
C
Excitatory
Suppressive
Excitatory
Suppressive
t,80 ms
x,1.2 deg
ωt, 50 Hz
ωx,3.2 c/deg

Figure 3-3. Model filters recovered for an example cell classified as simple by its response modulation to an optimized drifting sinusoidal grating (F1/DC = 1.51). A) The STA, three excitatory, and three suppressive filters recovered from the STC analysis shown in X-T coordinates. Each filter is scaled by the square root of its recovered weight (value indicated next to each filter. Weights were independently normalized for the excitatory and suppressive pools, with the largest in each case set to a value of 1. Also shown are the Fourier amplitude spectra in spatial and temporal frequency coordinates, similarly scaled by the square root of their weights. B) Pooled excitatory (green) and suppressive (red) filter spatiotemporal envelopes computed as the L^2-norm (square root of the weighted sum of squares) of the filter values for each X-T pixel. Regions of overlap are indicated by yellow. C) Pooled excitatory (green) and suppressive (red) frequency spectra as a weighted-sum of the amplitude spectra for each filter. As in B, regions of overlap are displayed in yellow.

STC. However, the STC analysis produced three additional excitatory filters, all with the same direction preference. In addition, three suppressive filters were recovered, all tuned for the direction opposite that of the excitatory filters.

It is also of interest to examine the net spatiotemporal and spectral extent of the excitatory and suppressive portions of the model. We computed separate spatiotemporal and spectral envelopes for the pooled excitatory and suppressive filters by summing the squared filters of their spectra. These results for are shown in figure 3-3B and 3-3C. For the example simple cell, the pooled excitatory and suppressive signals are almost completely overlapping in space and time. In the frequency domain, the excitatory and suppressive spectra are non-overlapping and tuned for opposite directions of motion.

Next consider a typical complex cell. The energy model (figure 3-1B) predicts a zero-valued STA and two STC filters with clear spatiotemporal structure. Data for an example complex cell (figure 3-4A) do show an essentially flat STA, but in addition to the two strongest excitatory filters, five additional excitatory and seven suppressive filters were revealed. As in the case of the simple cell, all of the excitatory filters had the same direction preference and most of the suppressive filters had the opposite direction preference. Note the weakest suppressive filter had the same direction preference as the excitation and was time-delayed relative to the excitatory filters. Unlike the simple cell, the excitatory and suppressive filters for the complex cell appeared in pairs, with each member of a pair appearing as phase-shifted copy of the other. The pooled

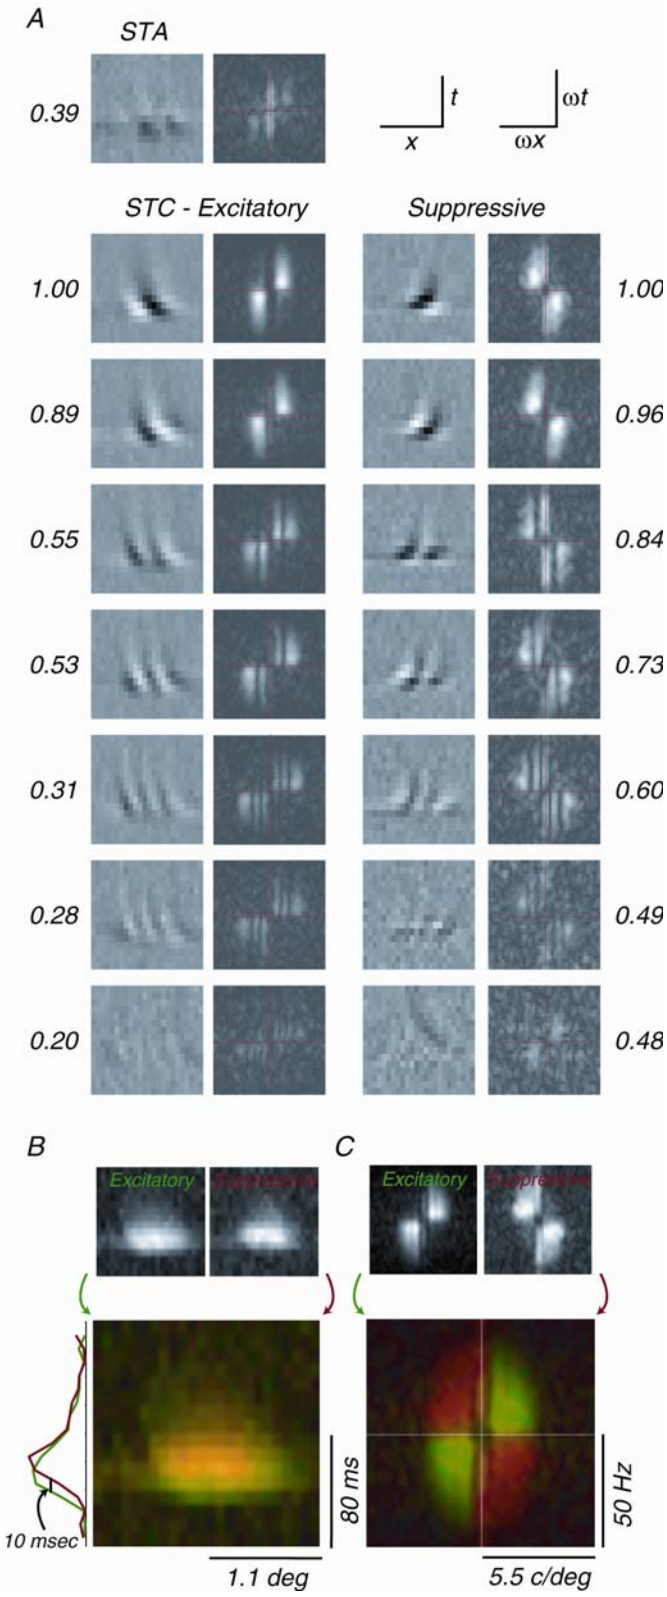
A
STA
0.39
t
x
ωt
ωx
STC - Excitatory
Suppressive
1.00
0.89
0.55
0.53
0.31
0.28
0.20
1.00
0.96
0.84
0.73
0.60
0.49
0.48
B
C
Excitatory
Excitatory
10 msec
80 ms
50 Hz
1.1 deg
5.5 c/deg

Figure 3-4: Model filters recovered for an example cell classified as complex by its lack of response modulation to an optimized drifting sinusoidal grating (F1/DC = 0.10). A) The STA, seven excitatory, and seven suppressive filters recovered from the STC shown with the same conventions as figure 3-3. The recovered weights for each filter are labeled. B) Pooled excitatory (green) and suppressive (red) filters, similar to figure 3-3B. Plotted along the y-axis is the time course of the pooled excitatory and suppressive signals for bar 12 to illustrate the delay of suppression relative to excitation. C) Pooled excitatory (green) and suppressive (red) frequency spectra (see figure 3-3C).

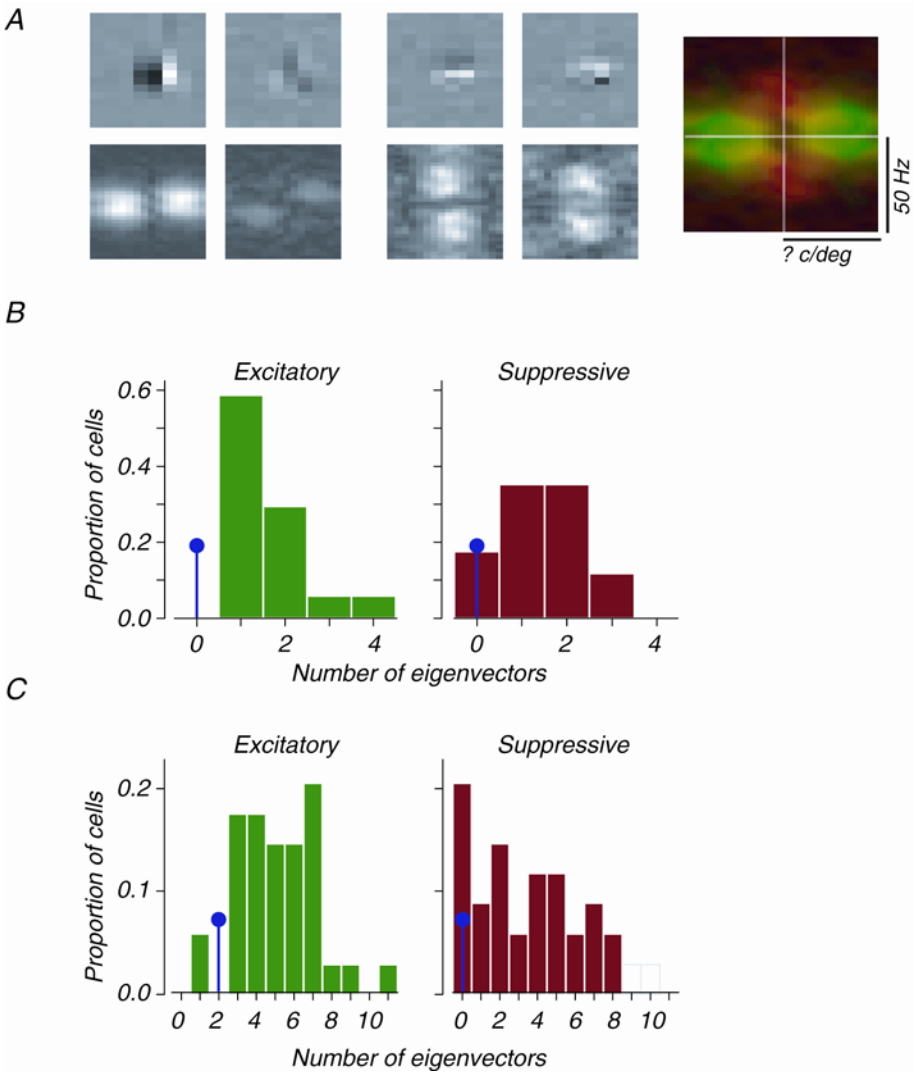

Figure 3-5. Characteristics of the population of V1 neurons. A) Filters recovered for a non-directionally tuned simple cell (F1/DC = 2.28). Excitatory and suppressive filters are shown along with their frequency spectra. Also shown are the pooled spatiotemporal and frequency spectra, computed for excitatory (green) and suppressive (red) filters with the same convention as panel 2C. (B-C) Number of filters revealed by STC (not including the STA) for simple (B, n=17) and complex (C, n=34) neurons. Numbers expected by standard models (figure 1A-B) are shown in blue. Only cells for which we gathered at least 50 spikes per spatiotemporal dimension were included in this analysis.

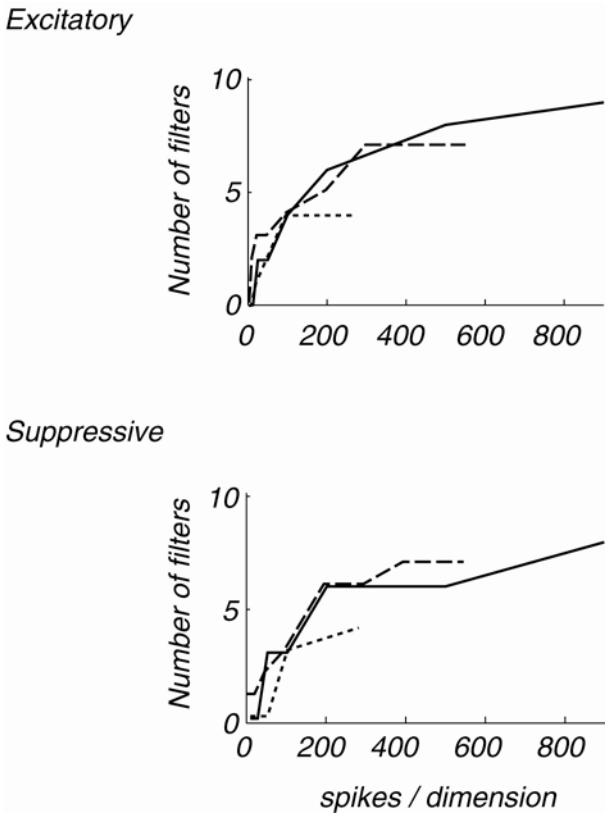

Figure 3-6: Dependency of the number of filters revealed by STC on the number of spikes included in the analysis. Shown are the number of excitatory and suppressive filters revealed as a function of the number of spikes collected per spatiotemporal dimension for three neurons. Dotted line: 512 dimensions; 143K spikes total. Dashed line: 384 dimensions, 212K spikes total. Solid line: 256 dimensions, 230K spikes total.

spatiotemporal excitatory and suppressive envelopes indicate that the suppression was time delayed relative to the excitation for this neuron. As seen from a time slice at a fixed spatial position (figure 3-4B, y-axis), this delay is approximately 10 msec. The excitatory and suppressive frequency spectra were almost completely non-overlapping (figure 3-4C), as in the case of the simple cell.

The two example cells shown thus far are representative of directionally-tuned simple and complex cells. The spatiotemporal relationship between the excitation and suppression for non-directional neurons had different characteristics. Figure 3-5A shows the strongest excitatory and suppressive filters for a non-directional simple cell that was tuned for the orientation of drifting gratings. The response of this cell was dominated by a sign-sensitive space-time non-oriented filter but also influenced by a weaker, directionally tuned filter, resulting in a pooled excitatory signal tuned broadly tuned for in spatial frequency. On the other hand, the suppressive axes for this cell were more narrowly tuned in spatial frequency but broadly tuned in temporal frequency.

Although the standard model of a simple cell (figure 3-1C) predicts a single linear filter for these neurons, across the population of simple cells we always recovered at least one (and as many as four) additional excitatory filters (figure 3-5B). The energy model of a complex cell predicts two excitatory filters, but in all complex cells but one, we recovered more than two excitatory filters (figure 3-5C). Suppressive filters were found for most simple and complex cells (figure 3-54B, C). While similar numbers of suppressive filters were re-

covered from directional and non-directional simple cells (means: 1.7 versus 1.3, respectively), directionally selective complex cells produced more suppressive filters than non-directional complex neurons (means: 5.9 and 1.8, respectively).

In general, the number of filters recovered depends on both the strength of their influence on the response, as well as on the number of spikes in the collected data (Chichilnisky 2001, Paninski 2003). We found this to be true experimentally (figure 3-6), and only included cells for which we collected at least 50 spikes per dimension in our population analysis (on average 255 spikes per dimension or 58,000 spikes total). Even so, we cannot guarantee that we have recovered all the excitatory filters for a cell and our results should thus be regarded as lower bounds on the true number of filters required to model these neurons.

Recovery of the nonlinearity:

After recovering a set of linear filters, the model is completed by recovering the nonlinear function (N) that combines the filter outputs to produce a firing rate. When the number of filters is very small (e.g., one or two), this can be done by computing the filter responses to the stimulus sequence, binning them, and constructing a table of the average number of spikes observed for each combination of filter responses. Unfortunately, we simply cannot collect enough data to fill

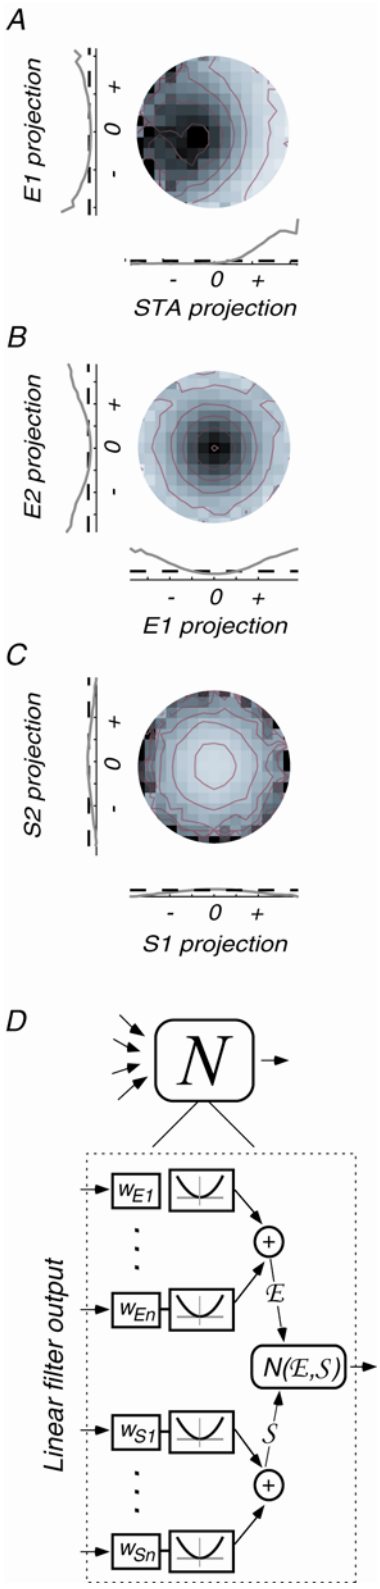
A
E1 projection
+
0
-
- 0 +
STA projection
B
E2 projection
+
0
-
- 0 +
E1 projection
C
S2 projection
+
0
-
- 0 +
S1 projection
D
N
Linear filter output
wE1
wEn
wS1
wSn
+
+
E
S
N(E,S)

Figure 3-7. The nonlinearity. (A-C) Firing rate as a function of the output of single filters and as the joint output of two filters. Firing rate is indicated by pixel intensity and red curves outline contours of constant firing rate. Black pixels indicate bins that were not estimated due to insufficient data. Also shown along the marginals are the 1-D firing rate functions for each filter individually. Dotted lines indicate the mean response to all stimuli. One and two dimensional firing rate functions are plotted for (A) The STA and strongest excitatory filter revealed by STC for the simple cell in figure 3-3, (B) The two strongest excitatory filters revealed by STC for the complex cell in figure 3-4, and (C) The two strongest suppressive filters revealed by STC for the complex cell in figure 3-4. D) The separability of the 2-D firing rate functions allows for a two-stage nonlinearity. In the first stage, the output of the excitatory and suppressive signals are pooled separately, each via a weighted sum of squares. In the second stage, a two-dimensional function governs the combination of the excitatory and suppressive signals to produce a firing rate.

in such a table to operate on the full set of filters we typically recover (4-15). However, we can examine the marginal responses for individual filters or pairs of filters. Figure 3-7A-C show firing rate as a function of the output of pairs of filters selected from the example simple and complex cells in figures 3-3 and 3-4. Also shown are firing rates as a function of the output of individual filters (along the vertical and horizontal axes).

Across our population of neurons, we found firing rate functions associated with the STA that were consistently halfwave-rectified (figure 3-7A, x-axis) whereas the firing rate functions associated with the STC filters were symmetric (figure 3-7A, y-axis; figure 3-7B-C, x- and y-axes). Excitatory STC filters (i.e., those recovered from the STC analysis corresponding to increased variance) had firing rate functions that increased monotonically with the absolute value of their outputs (figure 3-7B). Suppressive filters (those recovered from the STC analysis with decreased response variance) always produced firing rate functions that decreased monotonically with the absolute value of their outputs (figure 3-7C).

We examined the joint 2-D firing rate functions along different pairs of excitatory or suppressive axes and found that they took on a characteristic form: contours of constant firing rate along these pairs were well fit by ellipses and circles, suggesting that the firing rate can be expressed as a function of a weighted sum-of-squares of the filter responses (figure 3-7B-C). In the case of the STA and another excitatory dimension, the contours outlined a crescent

shape, consistent with a halfwave rectification followed by squaring (figure 3-7A). This regularity suggests that the dimensionality of the nonlinearity may be reduced by resolving it into two stages (figure 3-7D). In the first stage, the excitation and suppression are pooled separately, each by a weighted sum of squares. The STA is included in the excitatory pool, but is half-squared. In the second stage, the firing rate is computed as a function of the joint output of the excitatory and suppressive pools (labeled *E* and *S* in figures 3-7 and 3-8).

To recover the first stage of the nonlinearity, we obtained the weights for each filter by maximizing the mutual information between the weighted sum of squares of the joint excitatory and suppressive pools and the spikes; the recovered weights associated with each filter for the example cells are given in figures 3-3 and 3-4. The use of mutual information as an optimization criterion is advantageous, as it makes no assumptions regarding the form of interaction between the excitatory and suppressive signals. The second-stage of the nonlinearity is then recovered by binning the excitatory and suppressive signals and constructing a table of estimated firing rates for each pair of binned values. The table recovered for the example complex cell shown in figure 3-4 is shown in figure 3-8A.

Properties of the suppressive signal:

For those cells in which suppressive filters were revealed, we quantified the strength of the suppression as the decrement in the response of a strong excita-

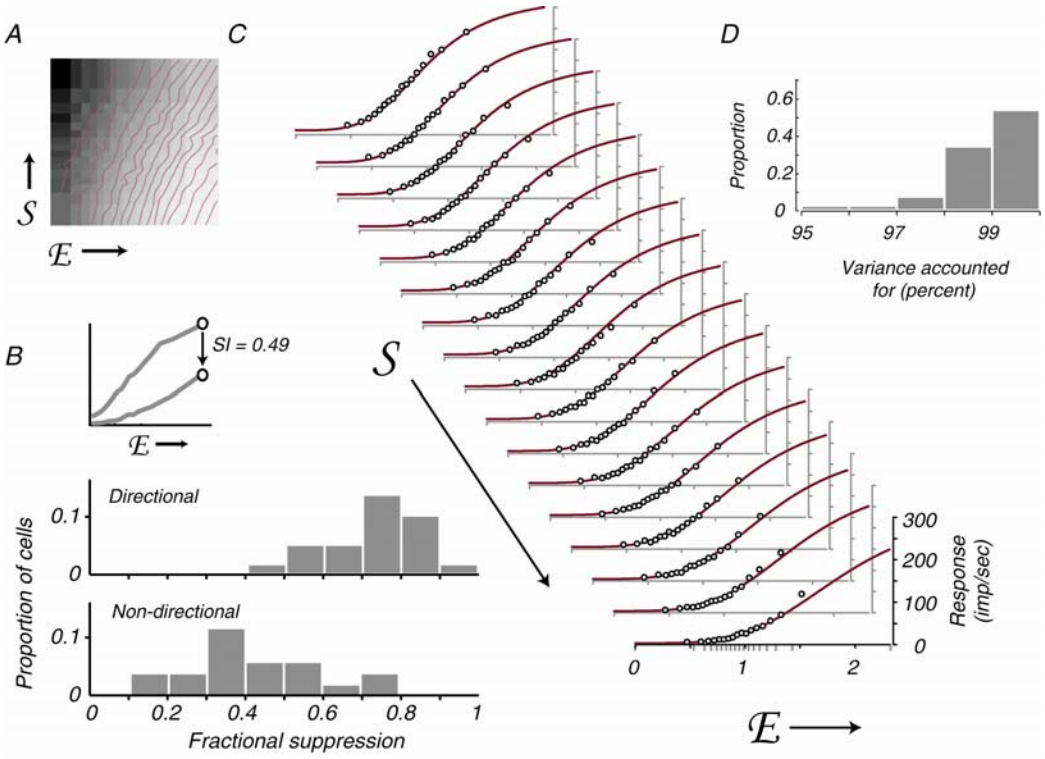

Figure 3-8. Characteristics of the suppressive signal. (A) Firing rate as a function of the joint output of the excitatory and suppressive pooled signals N(*E*,*S*) for the complex cell in figure 3-4. (B) Fractional suppression calculated as 1 - (maximal excitation with maximal suppression / maximal excitation with minimal suppression) where 0 indicates no suppression and 1 indicates complete suppression; the inset illustrates the two relevant data points taken from panel C. Shown are histograms of the fractional suppression computed for the 24 directional and 27 nondirectional cells with suppressive filters (means 0.73 and 0.42, respectively). (C) Slices through the same surface shown in A, taken through increasing excitation at constant suppression with minimal suppression at the top of the figure and maximal suppression at the bottom. Data points are shown. Also shown are fits for a model containing both divisive and subtractive influences (see Results). (D) Comparison of the variance accounted for by the model fits calculated as one minus the ratio of the mean-squared error of the fits and the variance in the data.

tory stimulus with and without suppression (figure 3-8B). Suppression in most cases was strong – for the example complex cell shown in figure 3-4, a suppressive stimulus was capable of reducing the response by 49%. On average, the fractional suppression was 58%. Suppression was stronger in directional (mean 73%) as compared to non-directional (mean 42%) neurons.

To determine the nature of the combination of excitation and suppression, we compared parametric models fit to the two-dimensional nonlinearity that describes firing rate as a combination the excitatory and suppressive pools (figure 8C). The data (289 points) were fit with a sigmoidal excitatory function in which suppression could enter through a subtractive and/or a divisive term. The model also contained weights for the excitation in the numerator and denominator, an exponent and a baseline offset parameter, resulting in 6 parameters. Figure 3-8C shows the data points and model fits plotted as slices of increasing excitation through different levels of suppression. Changes in the slope of the points at different levels of suppression are captured in the model by the divisive term whereas downward shifts are captured by the subtraction. As was typical of most cells, the divisive suppressive component described most of the suppression in the example cell but an additional subtractive component was also required to adequately describe the small downward shift of the neuron's response with increasing suppression. Figure 3-8C compares the variance of the data accounted for by the model for population of cells that produced significant suppressive axes. Across the population, the model provided an excellent

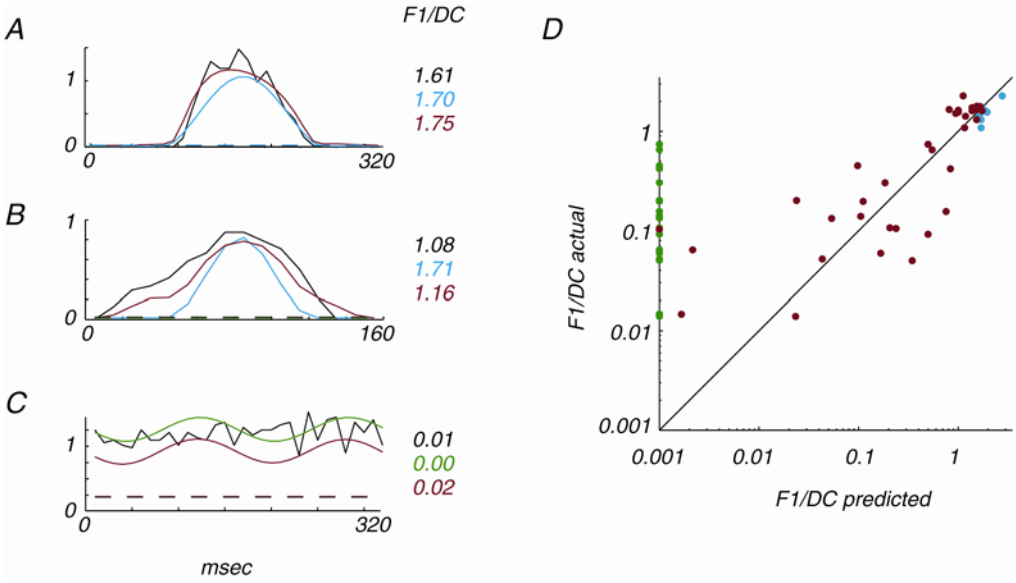

Figure 3-9. Predictions of response modulation to optimized drifting sinusoidal gratings. (A-C) The PSTH of the response to one cycle of a drifting sinusoidal grating for the actual response of the cell (black) the standard model prediction (for simple cells, the STA model, blue; for complex cells, the energy model, green) and the full STC model (red) for three V1 neurons. Dashed lines indicate actual and predicted baseline responses (responses to a mean gray screen). The actual and predicted response modulation indices (vector average of the first harmonic of the response / baseline subtracted mean response, F1/DC) are labeled. Weights of the excitatory filters for each cell: (A) STA: 1 STC: 0.04 (B) STA: 1 STC: 0.40, 0.18 (C) STA: 0.40, STC: 1.0, 0.99, 0.69, 0.61, 0.49, 0.44, 0.35 (D) STA model (blue), Energy model (green) and full STA+STC model predictions (red) of response modulation versus actual response modulation for a population of 35 V1 neurons. The Energy model predictions of response modulation were less than 0.001 and are shown at the edge of the plot (green).

account of the data.

Model predictions of responses to drifting sinusoidal gratings:

Given a complete mathematical description of a neuron's responses, we were interested in determining how well the model predicted the responses to another stimulus. Drifting sinusoidal gratings are commonly used to characterize these neurons and thus serve as a good test case. V1 neurons are often labeled as "simple" or "complex" based upon the modulation of their response to optimized drifting sinusoidal gratings (Skottun et al 1991), quantified as the ratio of the first harmonic of the response to the mean response (F1/DC). Cells with a modulation index greater than one are classified as simple while cells with a modulation index less than one are classified as complex. Figure 3-9A-C show an average response histograms for one cycle of a drifting grating for three V1 neurons. The top panel (figure 3-9A) shows data from a prototypical simple cell that responded to half of the cycle of grating drift. The bottom panel (figure 3-9C) shows data from a prototypical complex cell that responded in a phase insensitive manner to the entire cycle. The middle panel (figure 3-9B) corresponds to a cell with an intermediate behavior. This cell was classified as simple by the F1/DC criterion but its response lasted more than half a cycle. Colored traces indicate the predictions for two models: the standard model of these neurons and a full model including both the STA and filters recovered from STC

(green). For the simple cells (figures 3-9A and B), the standard model included the STA (blue). For the complex cell (figure 3-10), the standard model included the first two STC filters (the Energy model; green). In the case of the prototypical simple cell, the STA predicted the neuron's response well; adding the additional filters obtained by STC had little effect on the prediction due to the low weights recovered for these filters (figure 3-9A). For the "intermediate" cell, the STA was a fair predictor of the neuron's response; adding the filters obtained from STC improved the prediction (figure 3-9B). For the complex cell, both the energy model and full STC model were good predictors of response modulation (figure 3-9C).

A plot of the actual modulation index versus the predicted index for the three models illustrates that these three cells are representative of the population (figure 3-9D). The model including the STA alone predicted highly modulated responses from simple cells, and erred on the side over over-estimating the response modulation (blue points). Addition of the STC filters to the model reduced the response modulation prediction. Of note are the population of simple cells that live on the simple/complex cell border (like the cell in figure 3-9B). These cells had a strong STA that predicted a robust response to a drifting grating. However, the sign-insensitive filters revealed by STC were required to accurately predict the phase sensitivity of these cells, thus validating the existence of the additional excitatory filters in these neurons. For complex cells, the energy model predicted zero response modulation (green), despite the modulation

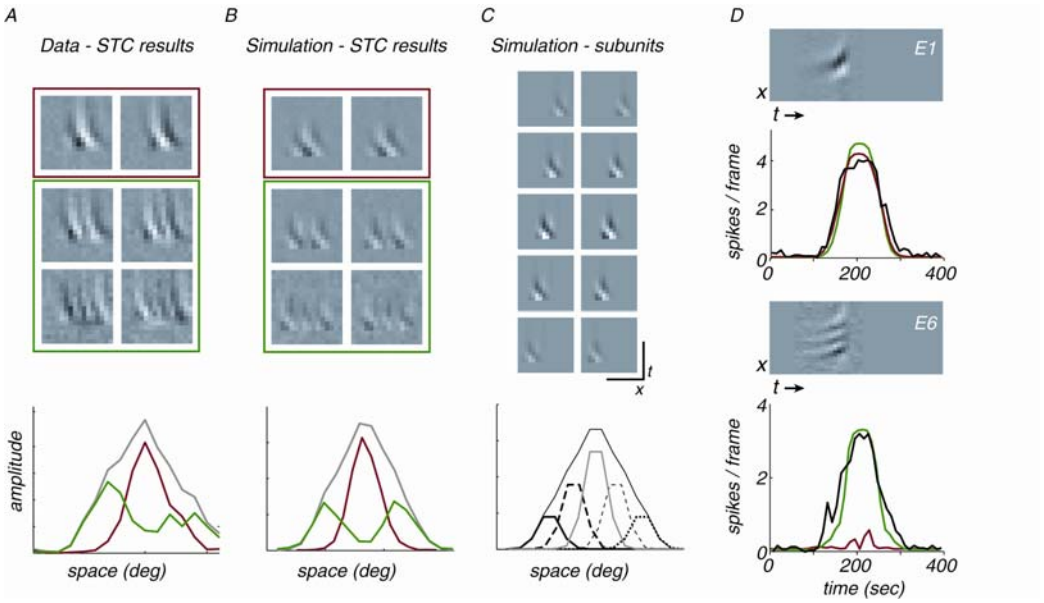
A
Data - STC results
B
Simulation - STC results
C
Simulation - subunits
D
E1
x
t
spikes / frame
4
2
0
0
200
400
E6
amplitude
space (deg)
space (deg)
space (deg)
time (sec)

Figure 3-10. Complex cell subunits. A) Top: the six excitatory filters revealed for an example complex cell (weights: 1, 0.98, 0.73, 0.49, 0.68, 0.47). Note the similarity with the cell in figure 3-4. Bottom: a cross section of the spatial profile of the receptive field, computed by taking the square root of the weighted summed squared combination of the pixel values for each of the filters. The cross section was taken at the peak temporal offset, t = 65 msec before a spike. The spatial profile is shown for all the filters (gray), the two strongest filters (red) and the remaining four filters (green). B) A set of filters and spatial envelope structure similar to the results to A, obtained by performing an STC analysis on the data produced by a simulation of the 10 subunits shown in C. C) The five spatially shifted quadrature pairs used in the simulation. Simulated data were generated by convolving each of the subunits with the random binary bar stimulus used in the experiment and combining the signals via a weighted sum of squares with weights of 0.33, 0.66, 1, 0.66, and 0.33 applied to the filter pairs as shown top to bottom. Bottom: the spatiotemporal profile of the pooled filters and individual filter pairs. D) Experimental validation of the unexpected filters for the cell shown in figure 3-4. Shown are the stimuli presented to the cell, the response of the cell (black), the predicted response of the energy model (the STA and first two filters, and nonlinearity reconstructed based on the pooled output of these three filters, red) and the full model including all excitatory and suppressive filters (green). Top: The stimulus presented was a movie of the strongest excitatory filter preceded and followed by periods of gray screen. Both models predict a similar response to this stimulus and as a result it can be used to determine the multiplicative factor required to predict the gain adjustment to low contrast stimuli. The gain adjustment was simulated as the best fitting scalar applied to the pooled excitatory and suppressive signals before they were converted into firing rates. For this stimulus, the scalar was determined to be 5. Bottom: Presentation of the sixth excitatory filter as a stimulus predicts a different response from the standard energy model and the STC model. The STC model better predicts the actual response of the cell.

observed in some of these neurons. Due to the inclusion of an asymmetric STA, the full model better predicts modulation in those complex cells with intermediate modulation values (red). For neurons at all points along the simple to complex continuum the recovered functional models produced respectable predictions of response modulation (red points, figure 3-9D).

Complex cell subunits:

In complex cells, the STC analysis consistently recovered more than just the filter pair predicted from the standard energy model. Figure 3-10A shows the filters recovered from STC for an example complex cell with similar characteristics to the cell in figure 3-4 to demonstrate the consistency with which we observed these results. In complex cells of this type we observed a specific structural relationship between the filter pairs: the temporal envelopes for all the filters revealed by STC were similar, but the spatial envelopes differed across filter pairs. Figure 3-10A compares the spatiotemporal envelope for the first two excitatory filters with the spatiotemporal envelope of the last four. Shown is a slice across the spatial envelope at the peak temporal offset. While the spatial structure of the two strongest STC filters was confined to the center of the receptive field, the spatial structure of the additional filters flattened in the middle but was robust at the receptive field edge. Mindful that the actual subunits of a cell can be linear combinations of the filters revealed by STC, we wondered what types

of subunits could produce such results. Simulations confirmed that a model of a complex cell that included multiple (6-8) spatially shifted subunits produces responses that are strikingly similar to filters we revealed from many complex cells (figure 3-10B).

To validate the existence of the unexpected filters in these cells, for a subpopulation of neurons we presented stimuli designed to discriminate between the standard (energy) model and the full model revealed by STC. For many cells, presentation of the unexpected filters as movies could be used as discriminating stimuli due to a combination of the orthogonality between the filters and the concentration of stimulus energy at the receptive field fringe. Figure 3-10D (top) illustrates the responses of the neuron shown in figure 3-4 to a stimulus that is predicted to produce a similar response from the standard and full STC model. Shown below is the response of the neuron to a stimulus that produces a drastically different prediction in the two models. While the standard model predicts only a small response to this discriminating stimulus, a robust response was evoked from the neuron as predicted by the full STC model.

3.3 Discussion

Conventional reconstructions of spatiotemporal receptive fields for V1 cells have been specific with regard to cell type; e.g. simple (DeAngelis et al 1993, Jones & Palmer 1987, McLean & Palmer 1989, Movshon et al 1978b) or com-

plex (Emerson et al 1987, Lau et al 2002, Livingstone & Conway 2003, Movshon et al 1978a, Touryan et al 2002) and focused primarily on linear or quasi-linear descriptions of neural response. In this paper we have presented a generalized model for V1 neurons that can be applied to any neuron regardless of cell type. This model includes a linear processing stage based on responses of a small set of filters (adjustable in number), followed by a nonlinear function that combines the filter outputs in order to generate a firing rate. The resulting quantitative model can be fit to data using spike-triggered techniques, and used to predict responses to arbitrary stimuli.

In the process of fitting these models to extracellular data, we found that most cells required substantially more filters than predicted by standard models of V1 neurons. For every simple cell we tested, we recovered more than the single linear filter predicted by the standard model of these cells (figure 3-1A). For nearly every complex cell tested, we recovered more than the two filters predicted by the energy model. In addition to mapping the excitatory influences in these cells, we also uncovered the spatiotemporal tuning of strong suppressive influences in V1.

Initially, we were concerned that the unexpected filters could be an artifact resulting from the small involuntary eye movements that are known to exist in the anesthetized, paralyzed macaque preparation (Forte et al 2002). Although we cannot rule out this possibility completely, several pieces of evidence indicate that eye movements alone cannot explain the discrepancies between our

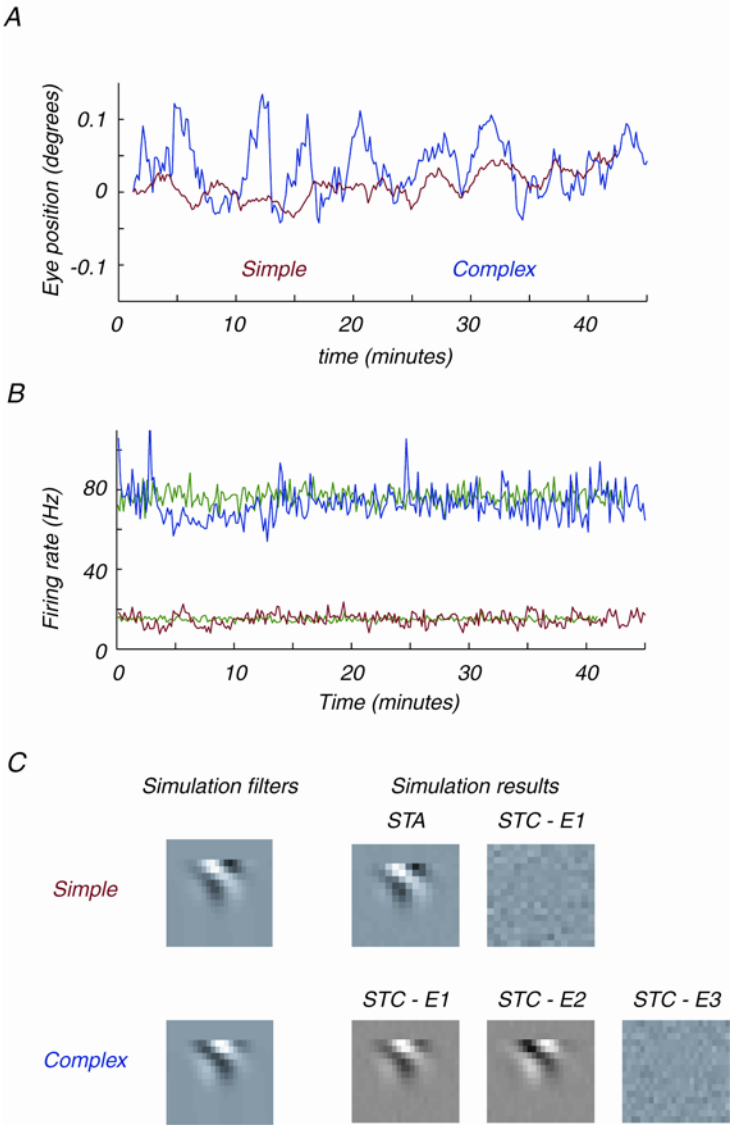

Figure 3-11: Eye movement analysis. A) Estimation of eye position during data collection for the example simple cell shown in figure 3-3 and the example complex cell shown in figure 3-4 (located in different animals). For both cells, eye position was estimated from the data collected in 2.5 minute windows; successive eye position estimates were obtained by shifting the window forward in 10 second increments. For the simple cell (red), eye position was estimated

from the STA computed for each window. A Gabor (sine multiplied by a Gaussian) was fit to the time slice at the peak offset (t = 65 msec before a spike) and a position parameter (the center of the Gaussian) was used as an estimate of eye position. The estimated eye position deviated over 0.09 degrees, approximately half the width of one bar (0.2 degrees). For the complex cell (blue), a STC analysis was computed from the windowed data and the spatial envelope was calculated by taking the L^2-norm (square root of the sum of squares) of the two strongest excitatory filters. A position parameter was extracted by fitting a Gaussian to the data at the peak offset (t = 55 msec before a spike). The eyes moved over an absolute deviation of 0.17 degrees, approximately 2 bar widths (bar width, 0.09 degrees). The magnitude of the estimated eye movements are within the range reported by direct tracking of the eyes under similar experimental conditions as are the oscillations shown in both traces with a periodicity of 3-8 minutes (Forte et al 2002). To examine the effects of eye movements of this magnitude on the STC analysis, we simulated a standard model simple (figure 3-1A) and complex (figure 3-1B) cell. The top filter was included in model simple cell; both filters were included in the complex cell model. Eye movements were simulated by shifting the filters by the magnitude suggested by the traces in A and taking the dot product of the resulting filters and a binary bar stimulus every 10 msec. For both simulations, the size of the receptive field, number of bars used in the experiment, firing rate, and experiment duration (total number of spikes collected) were matched to the experiment. B) Actual firing rates over the course of data collection for the simple (red) and complex cell (blue). Also shown are the firing rates over the course of the two simulations (green). C) In simulation, the eye movements shown in A fail to produce artifactual filters not included in the models. For the model simple cell, only an STA was recovered. Also shown is the strongest (nonsignificant) excitatory filter revealed by STC, which had no spatio-temporal structure. For the model complex cell, only the two expected excitatory filters were revealed by STC. Also shown is the third strongest (nonsignificant) filter, which had no spatio-temporal structure. Additional simulations reveal that larger movements of the eyes can produce unexpected filters. Simple cells appear to be particularly prone to artifactual filters, due to the residual variance remaining after the STA is projected out of the spike-triggered stimulus distribution in preparation for STC (e.g. the red eye movement trace in A magnified four-fold produced an artifactual filter in simulation). In both simple and complex cells, large deviations of the eyes result in shifts of the receptive field away from the stimulus array and consequently decreases in firing rate during these episodes. We have explored the parameter space and failed to find suitable conditions under which the firing rate remains constant throughout the simulated experiment (as shown in B) and a large number (>4) filters are revealed.

results and standard models. First, the predictions of response modulation to a drifting sinusoid grating are based on the F1/DC ratio, which is relatively unaffected by eye movements. For all but the most highly modulated simple cells, the additional filters revealed by STC are required to properly predict phase sensitivity. Second, we have estimated the extent and timecourse of the eye movements that occurred during data acquisition by analyzing short segments of the data. These eye movement traces fail to produce artifactual filters in simulation (Figure 3-11). Furthermore, we have systematically explored the effects of the eye movements described by Forte et al (2002) on STC results in simulation and found that the number of artifactual filters produced by those eye movements are inconsistent with our data (not shown).

We also wondered whether the unexpected excitatory and suppressive filters were produced by deviations from the Poisson spiking assumed by the model (figure 3-1C). For example, the intracellular mechanisms associated with spike generation, such as the refractory period, can produce suppressive filters in an STC analysis that do not reflect true subunits (Aguera y Arcas & Fairhall 2003, Pillow et al 2004). Similarly, correlated excitatory events such as bursting could suggest artifactual excitatory subunits. Simulations confirm that the characteristics of the filters we are observing are inconsistent with the filters expected to arise from non-Poisson spiking. In the case of spatiotemporally inseparable (directionally tuned) excitation, the suppressive filters resulting from models containing a refractory period or integrate-and-fire dynamics appear as filters with

the same direction preference but time-delayed relative to the strongest excitatory filters (not shown). Such filters are not consistent with the suppressive filters that we are observing, which commonly have a direction preference opposite the excitation. Similarly, the artifactual filters produced by a model of a bursting neuron are time-delayed relative to the strongest excitatory filters whereas the excitatory filters we observed varied in the spatial as opposed to their temporal profiles.

The results presented here suggest that V1 neurons are constructed of more subunits than predicted by the standard models of these cells. These extra subunits, in turn, provide an explanation of the phase-sensitivity of the cells, which varied inversely with the number of excitatory filters. Simple cells with the highest degree of phase sensitivity were well described by a single half-rectified excitatory filter (an STA), although evidence for a weak additional excitatory subunit was consistent in these neurons. Cells with intermediate phase sensitivity had weaker STAs and stronger excitatory filters revealed by STC, indicative of a combination of a number of half-rectified (and potentially full-rectified) subunits. Complex cells with the highest degree of phase-invariance appeared to be comprised of 5-6 smaller full-rectified subunits that converged to form their spatial profiles.

What biophysical mechanisms produce the filters that we are uncovering? Our model attempts to describe the responses of cells as a function of the input stimulus, and thus includes all processing preceding the V1 neuron in question

as well as any time-delayed inputs (e.g. feedback or lateral connections) that these neurons receive. Even at the earliest stages of visual processing, multiple linear filters are resolved from retinal ganglion cells using this technique (J.W. Pillow, E. P. Simoncelli, and E. J. Chichilnisky (2003). Soc. for Neurosci. abstracts). The filters we are recovering in V1 may reflect nonlinear processing (e.g. rectification) of signals in the retina and the LGN. A second possible source of multiple filters is the convergence of rectified signals within V1. Complex cells had on average twice as many subunits as simple cells, suggestive of convergence within this area. Alternatively, nonlinear intracellular mechanisms could be the source of multiple subunits. We rarely observed evidence for time-delayed excitatory influences (e.g. feedback), which would have been revealed as excitatory filters with different temporal profiles (although see figure 3-4A).

Parametric model fits to the data revealed that the suppressive signal had a combination of divisive and subtractive influences on the excitation. Note that this analysis resolves signals into different filters only following a nonlinear operation. In the push-pull model of a simple cell, the excitation and inhibition are combined linearly and together would produce a single linear filter (corresponding to the STA). Similarly, in the STC analysis the labeling of a filter as "excitatory" or "suppressive" is dependent upon changes in the second-order statistics of spiking stimuli. If excitation and suppression coincide along a single axis (where an axis is defined by a filter and its inverse), STC will resolve

the signals as a single axis or multiple axes, depending on the form of the combination. Hence our results may underestimate the subtractive influences in these neurons.

In a related STC characterization of cat V1 neurons, fewer excitatory filters (on average 2) were recovered from complex cells, although as many as five filters were recovered from some neurons (Touryan et al 2002). Suppressive filters weren't reported in this study, but an examination of the non-significant stimulus dimensions resulting from the PCA suggested that they were better described as having divisive than subtractive influence on V1 neuron's responses. The additional excitatory and suppressive filters we are recovering from monkey V1 neurons as compared to the cat may be explained by the large number of spikes (on average 58,000 per cell) we collected. Alternatively, a difference in processing between the two species may exist.

Here we have presented the most profound deviations from standard models of V1 neurons observed when constructing functional models of neurons in V1 using spike-triggered techniques. Many other interesting properties of these neurons could be revealed by extending the analysis to cover the second spatial dimension as well as other stimulus attributes like color, binocularity, and surround suppression. These extensions would be required to build a model of a neuron that could predict the response to *any* arbitrary stimulus. Even when confined to examining a limited number of stimulus attributes, spike-triggered

techniques have proven themselves to be sensitive tools that can be used to uncover the subtleties of neuronal computation.

4 The role of suppression in shaping direction selectivity in visual areas V1 and MT

The computation of motion direction begins in primary visual cortex (V1). Neurons in the lateral geniculate nucleus (LGN) are untuned for the orientation and direction of stimuli. These input signals are transformed in V1 into responses that live on a continuum from oriented (equally responsive to both directions of motion along one axis) to directionally selective (responsive to motion in one direction with little or no response to motion in the opposite direction). A subpopulation of V1 neurons tuned for direction project to the next stage of motion processing, visual area MT (Movshon & Newsome 1996).

The majority of neurons within MT are strongly tuned for motion direction. Most neurons in MT respond vigorously to stimuli moving in a preferred direction and are suppressed below spontaneous firing by motion in the opposite (null) direction (often referred to as "motion opponency"). Models of direction computation typically include a multi-stage process in which directional bias is formed, and these directional signals are then further sharpened by a suppressive signal tuned to motion in the opposite direction (Adelson & Bergen 1985, Simoncelli & Heeger 1998). A number of contradictions exist in the literature with regard to the source of this suppressive signal.

If neurons in MT receive inhibitory input from other MT neurons with opposite direction preferences, the null suppressive signal would be expected to

act globally across the large MT receptive field. Experimental evidence suggests that this is not the case (Qian & Andersen 1994). The local nature of the suppressive signal suggests that it must act before spatial pooling within MT neurons.

Where does the null direction suppressive signal act? The local nature of this computation suggests that its locus is unlikely to be MT. Long-range feedforward intracortical connections (from V1, V2, and V3) are believed to be exclusively excitatory (White 1989). Thus local inhibition in MT would require the existence of a population of MT inhibitory interneurons with small receptive fields that receive excitatory projections from V1 and in turn produce inhibitory projections that combine locally with excitatory signals. These hypothetical neurons are unlikely; MT receptive fields are all 6-10 times the diameter (30-100 times the area) of V1 receptive fields at any given eccentricity (Van Essen et al 1981). However, investigations of V1 have suggested that the null suppressive signal is unlikely to act there (Qian & Andersen 1994, Qian & Andersen 1995). We were interested in revisiting questions related to the nature and source of the null suppressive signal in MT through stimuli specifically designed to isolate components of this signal.

4.1 Methods

See the *Appendix* for details regarding experimental preparation. Stimuli were presented on a gamma-corrected monitor with a refresh rate of 100 Hz and a

mean luminance of 33 cd/m^2 and generated with a Silicon Graphics XX workstation. The monitor was directed toward the monkey via a front surface mirror; total length between the eye and the monitor was 80 - 180 cm.

We recorded from isolated single units stimulated monocularly. Upon encountering a cell, the initial characterization involved an optimization for the best direction, spatial frequency, temporal frequency, and size of drifting grating stimuli. Beyond this stimulus optimization, stimuli were presented at optimal spatial and temporal frequency and confined to the classical receptive field. Only strongly directional MT (n=32) and V1 (n=18) cells were recorded.

Counterphase family stimuli

The first set of experiments are described in the Results (figure 4-1); the stimulus used for the remaining experiments are described here. The basic stimulus set is diagramed in figure 4-2a. Individual stimuli were weighted combinations of a preferred (P) and null (N; direction opposite preferred) drifting grating presented in blocks of constant total contrast (T) such that:

$$w_pP + w_nN = T$$

where w_p and w_n correspond to the contrast of the preferred and null drifting grating, respectively. Stimuli were constructed from 11 combinations of preferred and null weights including all preferred energy (w_p=1, w_n=0), all null energy (w_p=0, w_n=1), and nine points in between. In the case of equal contribu-

tions of preferred and null grating (w_p=w_n=0.5) the stimulus is a stationary contrast inverting (counterphase) grating.

Curves were collected at 5 different total contrasts (T). For most cells, T = 6.25%, 12.5%, 25%, 50%, 100%. The total contrast range was adjusted for some cells to account for the neuron's sensitivity. Stimuli at the five contrast levels were presented concurrently for 320 msec in 1.5 min blocks of constant total contrast. An additional block contained preferred stimuli at different contrast levels to obtain a traditional contrast response function for each cell. The final stimulus block contained a 15 second blank (mean gray) stimulus to measure the steady-state baseline response. Within each block, stimuli were randomly interleaved and blocks were interleaved across trials. Stimulus variations included:

1. As a control, responses to the basic stimulus set with blocks of constant contrast were compared to response to all stimuli interleaved (figure 4-3c).
2. For a subpopulation of cells, contrast energy rather than total contrast was held constant according to:

   $$(w_pP)^2 + (w_nN)^2 = T$$

   In these experiments, a limited range of preferred contrasts were explored due to physical limitations imposed by the upper limit of 100% contrast (figure 4-3d).

3. To explore the degree to which the influence of the null drifting grating depended upon direction, the null drifting grating was replaced with a grating drifting in another non-excitatory direction (e.g. 90 degrees from preferred; figure 4-3b, figure 4-11).
4. To assess changes in baseline response during different levels of total contrast presentation, in a subpopulation of cells, 320 msec blanks were randomly inserted into each total contrast block (figure 4-5, figure 4-10 a,b).

Post-stimulus time histograms (PSTHs) were constructed with 10 msec bins and the latency for each cell was determined by eye as the first time bin after stimulus onset that exceeded background firing rate for the 100% contrast stimuli. Mean firing rates were determined by aligning stimuli and latency adjusted spike trains and determining mean firing rates in 270 msec of the 320 msec stimulus presentation. PSTHs were examined for uniformity across the 270 msec to ensure that the spikes and stimuli were properly aligned.

Models were fit using the STEPIT algorithm (Chandler 1969) to minimize the chi-squared error between the actual responses and model predictions. We assessed goodness of fit by calculating the variance accounted for by the model as the ratio between the mean squared error of the fits and the variance in the data. The models accounted for a minimum of 92% of the variance (97% on average).

4.2 Results

4.2.1 Spatial extent of null suppression in MT

To explore the location specificity of null suppressive interactions between MT receptive field subregions, we stimulated the receptive field with small grating patches. In one patch, we collected a contrast response curve (the test patch). We presented this stimulus alone and in the presence of a second patch containing a 50% contrast sinusoidal grating stimulus drifting in the non-preferred direction (the null pedestal); the null pedestal was located either at the same location as the test or at a spatially distinct but approximately equally responsive location of the receptive field (figure 4-1a). In the separated configuration, the patches were presented symmetrically about the center of the classical receptive field and perpendicular to the cell's direction axis. In all cells, a preferred drifting grating presented at the separated location evoked a vigorous response from the neuron. We set a criterion that the preferred pedestal patch evoke at least 60% the response of the test patch at 50% contrast (mean 90%, mode 80%).

Figure 4-1a shows the responses of a typical cell. Shown are the fits to the contrast response functions when collected alone and in the presence of a null pedestal at the same location (figure 4-1a, top). This neuron had a considerable baseline firing rate (~31 spikes/sec) and a response that increased with contrast (dark grey). Presentation of the null pedestal alone decreased the response of the neuron relative to its spontaneous rate (to ~18 spikes/sec). Presen-

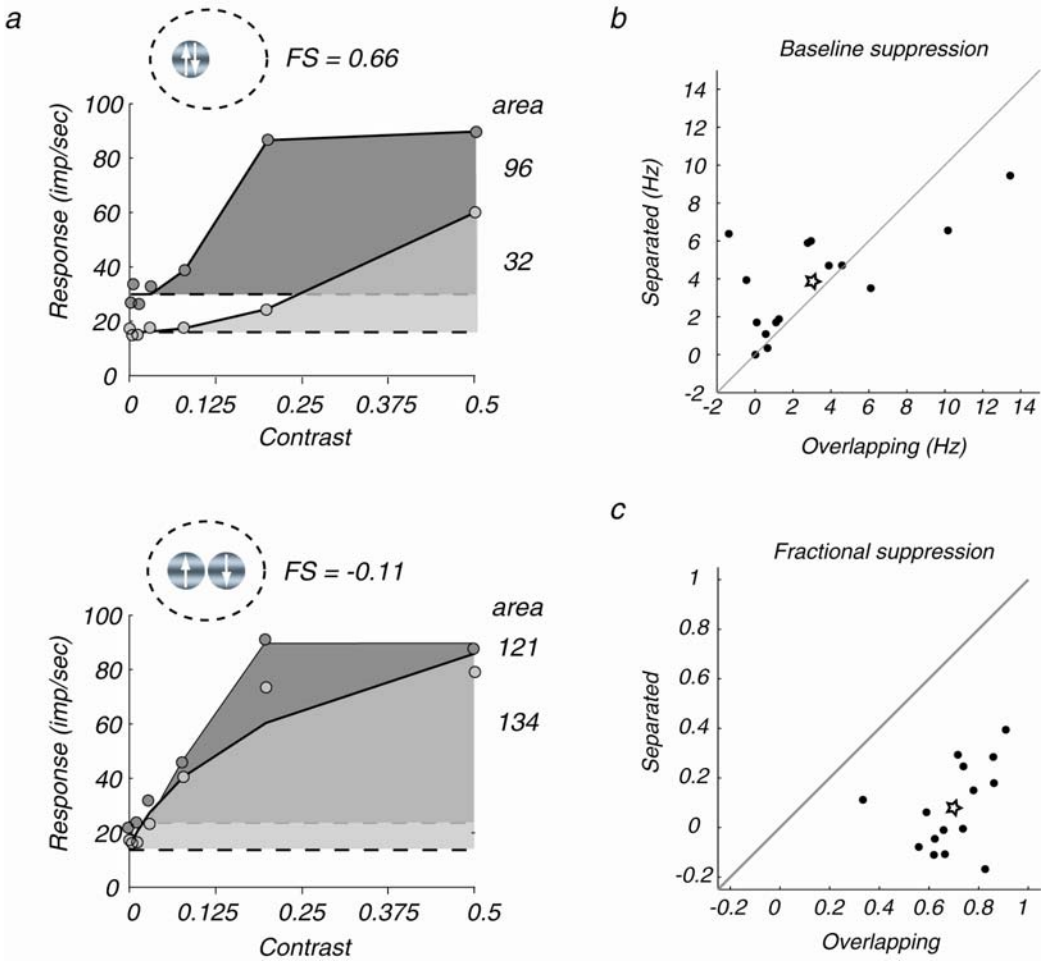
a
FS = 0.66
area
96
32
Response (imp/sec)
Contrast
FS = -0.11
area
121
134
Response (imp/sec)
Contrast
b
Baseline suppression
Separated (Hz)
Overlapping (Hz)
c
Fractional suppression
Separated
Overlapping

Figure 4-1: Spatial extent of the null suppressive signal in MT. Stimuli presented in these experiments were small patches of sinusoidal gratings drifting in the preferred or null direction. On each trial, two gratings were presented. The "test" grating always moved in the excitatory direction, was always presented at the same location (left patch), and was used to measure the contrast response of the neuron. The "pedestal" contained a 50% contrast grating drifting in the null direction and was either presented at the same (top panel) or a separate (bottom panel) location. The data collected for each condition were fit with an appropriate function and the area under each curve calculated (labeled). The contrast response function collected alone is indicated by dark-gray; the contrast response function collected in the presence of a null pedestal is indicated by light-gray. Dashed lines indicate the response to a gray screen (dark gray) and the response to the null pedestal alone (light gray). b) The decrement in the response from baseline when the null grating was presented in the same location as the test patch or at a separate location. Population means are indicated by the star. c) The fractional suppression of the test patch by the null pedestal, calculated as the ratio of the area under the contrast response function in the presence and absence of the null pedestal, subtracted from 1. A fractional suppression of 1 indicates complete suppression, 0 indicates no suppression, and values less than zero indicate an excitatory response that was stronger in the presence of the null pedestal than in its absence. Population means for the overlapping and separated conditions are indicated by the star.

tation of the null pedestal simultaneously with excitatory stimuli greatly reduced the responses to all excitatory contrasts (light grey). The null pedestal placed at a spatially distinct location within the receptive field produced a similar decrement to the baseline firing rate (figure 4-1a, bottom). When the contrast response function was collected in the presence of this non-localized null pedestal, it appeared similar to the contrast response in the absence of a pedestal but shifted downward.

This cell was representative of most neurons we recorded: in the overlapping configuration, a null drifting pedestal had a dramatic impact upon the neuron's response whereas in the separated configuration the null drifting pedestal primarily shifted the contrast response downward. To summarize this behavior for a population of cells, we considered two components of these responses. Figure 4-1b plots the decrement in baseline in the overlapping versus separated configuration for 15 cells. On average, the baseline decrement was similar under the two conditions. To assess the effect the null pedestal had on the test patch after this baseline decrement was considered, we calculated the fractional suppression as the ratio of area under the contrast response function with and without the null pedestal subtracted from 1 (labeled FS). A fractional suppression of 1 corresponds to complete suppression, 0 to no suppression, and negative values correspond to an excitatory response that was larger in the presence of the null pedestal than when presented alone. For the example cell, the effect of an overlapping null pedestal on the contrast response was strong, reducing the

response by 66%. In the separated configuration, the null drifting pedestal slightly facilitated the response (FS = -11%; figure 4-1a). Across the population of MT neurons we recorded, the effect of a null pedestal upon an excitatory response was dramatic when the two stimuli were co-localized but weak when the stimuli were presented in separate locations (means 70% and 8%, respectively). For most cells, a null drifting patch presented at a spatially distinct location had minimal effect beyond a reduction in the baseline firing rate.

These results suggest two components to the null suppressive signal in MT. First, null suppression has a dramatic effect upon co-localized excitation beyond that expected in an additive model by the decrement of a null stimulus to its baseline firing rate. This behavior is indicative of a strong suppressive signal that is masked by rectification at low firing rates. The combination of excitation and suppression followed by rectification appears to occur locally before spatial pooling in MT neurons, as evidenced by the inability of a null stimulus to affect spatially distinct excitation beyond its effect upon baseline (see also Qian and Andersen, 1994). However, if the suppression were rectified before reaching MT, a null stimulus would fail to suppress the response below baseline. Thus a second component of null suppression must exist to account for the downward shift of the contrast response when the test and pedestal are separated.

4.2.2 Counterphase family

Upon confirming previous reports that a component of the null suppressive signal is local in MT, we were interested in further characterizing this signal in MT and the strongly directionally V1 cells that are known to form a large component of the MT input. To do so, we used a stimulus that we will refer to as the "counterphase family". In each instantiation, the stimulus presented was a combination of a preferred drifting and null drifting grating, each presented at the same location and at the full size of the classical receptive field. While the total (summed) contrast of the two gratings was held constant, the contrast of the preferred versus null grating varied in 10% increments ranging from a preferred motion exclusively, through equal preferred and null contrasts (counterphase flicker), to null motion exclusively. We presented 11 preferred-null ratios at 5 different total contrasts (figure 4-2a). We were interested in probing directionally tuned suppressive signals while minimizing the effects of untuned suppressive influences (like contrast normalization). Thus stimuli were presented in blocks of constant total contrast in an attempt to hold the gain state of the cell constant while assessing the effect of varying the preferred and null components.

The responses of a representative MT cell to the counterphase family stimuli are presented in figure 4-2. Figure 4-2b illustrates the strong suppression caused by the addition of a null drifting grating with a contrast that increases as the preferred contrast decreases. For this cell, stimuli dominated by null drifting

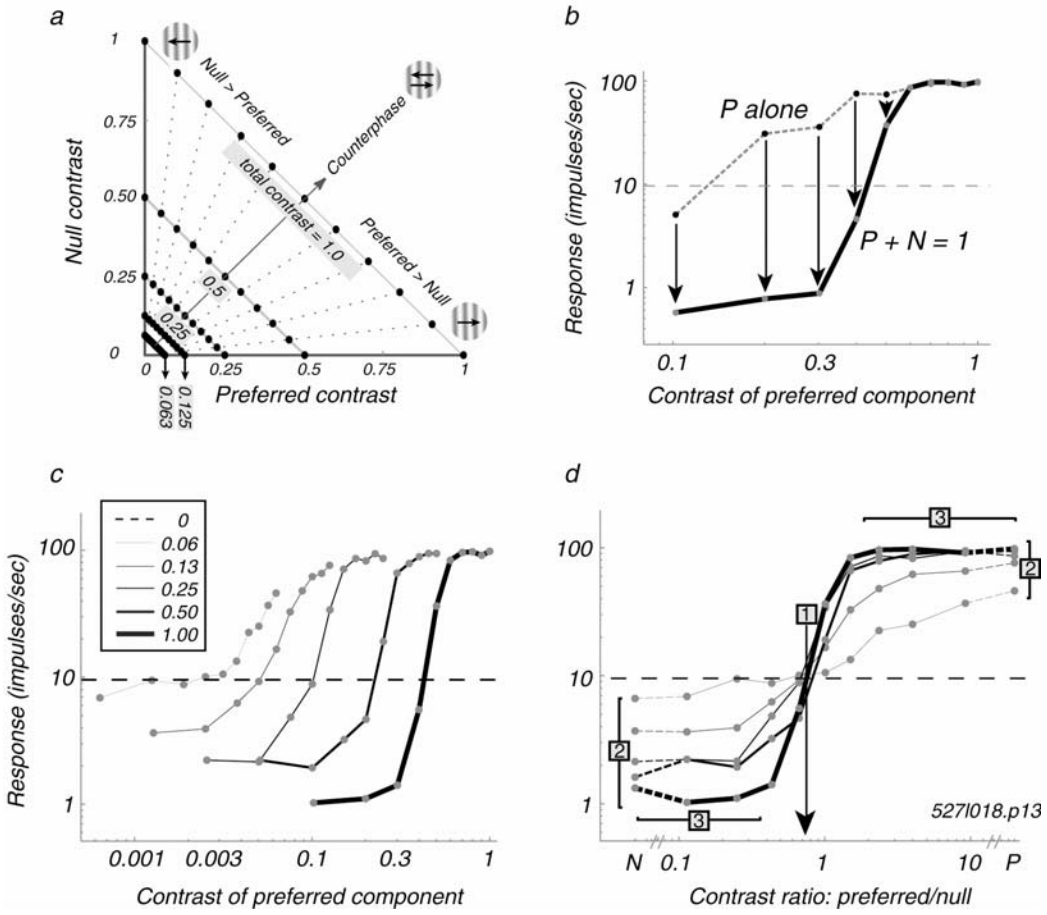
a
Null contrast
Preferred contrast
Null > Preferred
Counterphase
Preferred > Null
total contrast = 1.0
0.5
0.25
0.063
0.125
b
Response (impulses/sec)
P alone
P + N = 1
Contrast of preferred component
c
Response (impulses/sec)
0
0.06
0.13
0.25
0.50
1.00
Contrast of preferred component
d
Contrast ratio: preferred/null
527l018.p13

Figure 4-2: Representative response of an MT neuron to the counterphase family stimuli. a) Diagram of the stimulus design. Stimuli were weighted combinations of gratings drifting in the preferred and null (opposite) direction such that the total contrast was held fixed while the relative contribution by the preferred and null grating varied; equal contribution of preferred and null stimuli produces a contrast inverting stationary (counterphase) grating. Eleven preferred-null grating combinations were displayed at 5 different total contrasts; black dots correspond to the 55 stimuli presented. b) Response to a preferred grating of increasing contrast as compared to a stimulus containing the same preferred stimulus in addition to a null grating such that preferred contrast (P) + null contrast (N) = 1; arrows indicate the suppression caused by the null grating. Dashed line indicates the steady-state baseline response to a gray screen. c) Response to the full counterphase family plotted against the contrast of the preferred grating, disregarding the null grating contribution. Stimuli were presented in blocks of constant total contrast; line thickness corresponds the total contrast as indicated in the legend. d) The same responses shown in c but plotted against the ratio of preferred and null contrasts. "P" indicates the preferred stimulus alone (infinity on this axis), "N" the null stimulus (zero on this axis). Salient characteristics found across MT neurons and are labeled: 1) the curves collected at different absolute contrasts cross at a single point and this crossing point corresponded to a particular ratio, the "crossing-ratio" 2) curves collected at different total contrasts fan about the crossing point and the amount of this spread increases with total contrast 3) responses saturate for high contrast stimuli in both the preferred and null direction.

gratings were effective at decreasing the response to preferred gratings over 40-fold, well below baseline firing rates.

The responses of this cell to the full counterphase family are plotted as a function of the contrast of the preferred stimulus (figure 4-2c) and as function of the ratio of preferred and null contrasts (figure 4-2d). Across the population of MT neurons, responses to these stimuli displayed a characteristic structure illustrated by plotting response against the ratio between preferred and null contrasts on log-log axes (figure 4-2d). One of the most salient characteristics was the crossing of the curves collected at different total contrasts at a single point (labeled "1" in figure 4-2d). About this crossing point, the curves collected atdifferent total contrasts tended to fan out with a spread that increased with total contrast (indicated by "2" in figure 4-2d). The different total contrast curves often saturated at high preferred and null contrasts (labeled "3" in figure 4-2d).

The regularity of responses across MT provides clues into the mechanisms by which these receptive fields are constructed; the interpretation of each of these salient points are addressed below.

The crossing point

The crossing point occurred near steady-state baseline in all cells. This crossing point can be associated with a particular preferred:null contrast ratio; this "crossing ratio" represents the contrast ratio at which excitation overcomes suppres-

sion to cause an increase in firing. Cells with a crossing ratio of one have perfectly balanced excitation and suppression, a value less than one represents a cell dominated by excitatory influences, and a value greater than one represents a cell that is dominated by suppression. Both the cells in figures 4-2d and 4-3a have crossing ratios of ~0.7-0.8, indicating slightly stronger excitatory than suppressive influences (see figure 4-12a for a population comparison).

Because the crossing ratio reflects the strength of the null suppressive signal, we can use it to assess how the signal changes under different manipulations such as changing the direction of the non-preferred stimulus. Most cells in MT display at least partial "pattern" selectivity (Gizzi et al 1990, Movshon et al 1985) and will respond in a nonlinear manner to the intersection of two superimposed gratings drifting in different directions. This nonlinearity precludes probing the direction tuning of the suppressive signal. However, a subset of MT neurons are known to respond approximately linearly to the components of superimposed gratings (Gizzi et al 1990, Movshon et al 1985) and in these cells we compared the suppressive effect of the null drifting grating to that of an orthogonal (perpendicular) drifting grating. For the cell illustrated in figure 4-3, the counterphase family stimuli produced a crossing ratio of 0.8. Replacing the null grating with an orthogonal grating produced responses with a clear crossing ratio but one that was decreased to a value of 0.38, indicative of weaker suppression in the orthogonal versus null direction. This behavior suggests that the null suppressive signal is tuned for direction in MT.

The counterphase family experiments were run in blocks of constant total contrast in an attempt to keep the cell in a constant gain state. To determine whether the crossing point was dependent on long time-scale adaptation, we tested a subpopulation of neurons with all the stimuli interleaved (figure 4-3c). We observed clear crossing points in all cases.

Fanning and saturation

Saturation of the responses to stimuli at high preferred and null contrasts was common in MT. Saturation of signals at high contrasts have been described as contrast gain adjustments that occur through a process of divisive normalization in V1 and MT (Britten & Heuer 1999, Carandini et al 1997, Heeger 1992).

If the signal adjusting the gain of the MT cell were linear and untuned for direction, such saturation would not exist; the expected response would be linear rather than sigmoidal. Models of contrast gain control postulate a gain signal that is proportional to the square of the contrast (Heeger 1992). As a test of this normalization model, we held the energy of the stimulus constant (preferred contrast^2 + null contrast^2 = constant; figure 4-3d). If the squared form of the normalization model were correct, we would expect non-saturating (linear) responses to the constant energy stimulus. Although we couldn't test the full range of preferred and null contrasts while keeping the contrast energy constant

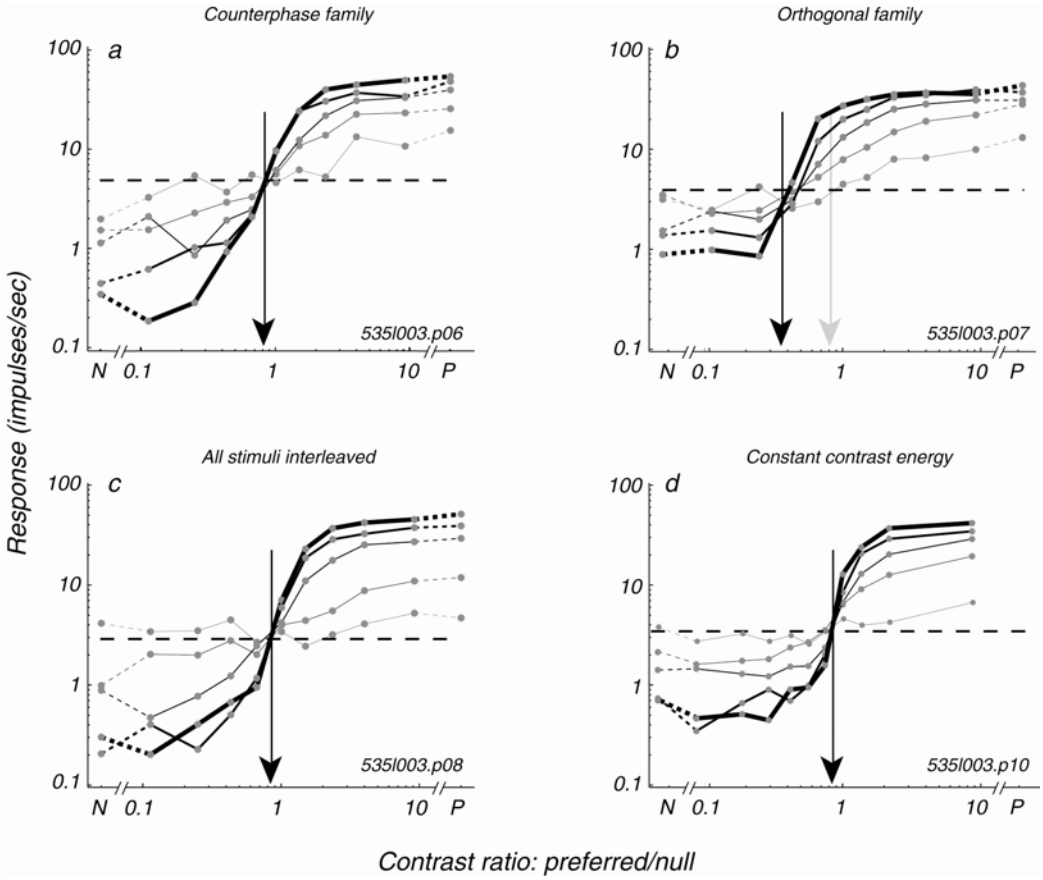

Figure 4-3: Counterphase family variants. In all figures, arrows indicate the crossing ratio for each cell; line thickness corresponds to total contrast with the same convention as figure 4-2. a) Counterphase family run in blocks of constant total contrast, similar to figure 4-2. b) Results from an experiment in which the null stimuli were replaced with orthogonal (90 degree from preferred) drifting gratings. Gray arrow indicates the crossing ratio from panel a for comparison. c) Counterphase family run with all 55 stimuli interleaved. d) Counterphase family with constant contrast energy as opposed to constant total contrast (see text).

(due to physical upper bound of 100% contrast), even within the range tested responses clearly saturated. These results are inconsistent a non-directional, squared gain control signal.

About the crossing point, the curves collect at different total contrasts displayed a consistent fanning structure. The amount of spread of the excitatory-dominated responses compared to the suppressive-dominated responses (those to the right and left of the crossing point, respectively) was often symmetric about the crossing point when plotted on a log response axis (figures 4-2d, 4-3a, 4-4c). In other cases, we observed a larger magnitude fanning above (figure 4-4b) and below (figure 4-4a) the crossing point. The amount of fanning was not reliably connected to the crossing ratio of the cell.

It is interesting to note that presentation of null suppressive stimuli at the highest contrasts never completely silenced MT neurons; we reliability recorded non-zero firing rate responses to null direction stimuli that decreased with the contrast of the null stimulus. The suppressive signal in many neurons is quite strong; for the cell displayed in figure 4-2b, a null grating reduced the responses to an excitatory stimulus by more than 25 impulses/sec and yet a null drifting grating presented alone failed to completely silence the 10 impulses/sec baseline firing rate. If the effect of the null suppressive were purely subtractive we would expect that a high contrast null drifting stimulus would hyperpolarize the cell below threshold, resulting in zero firing rates.

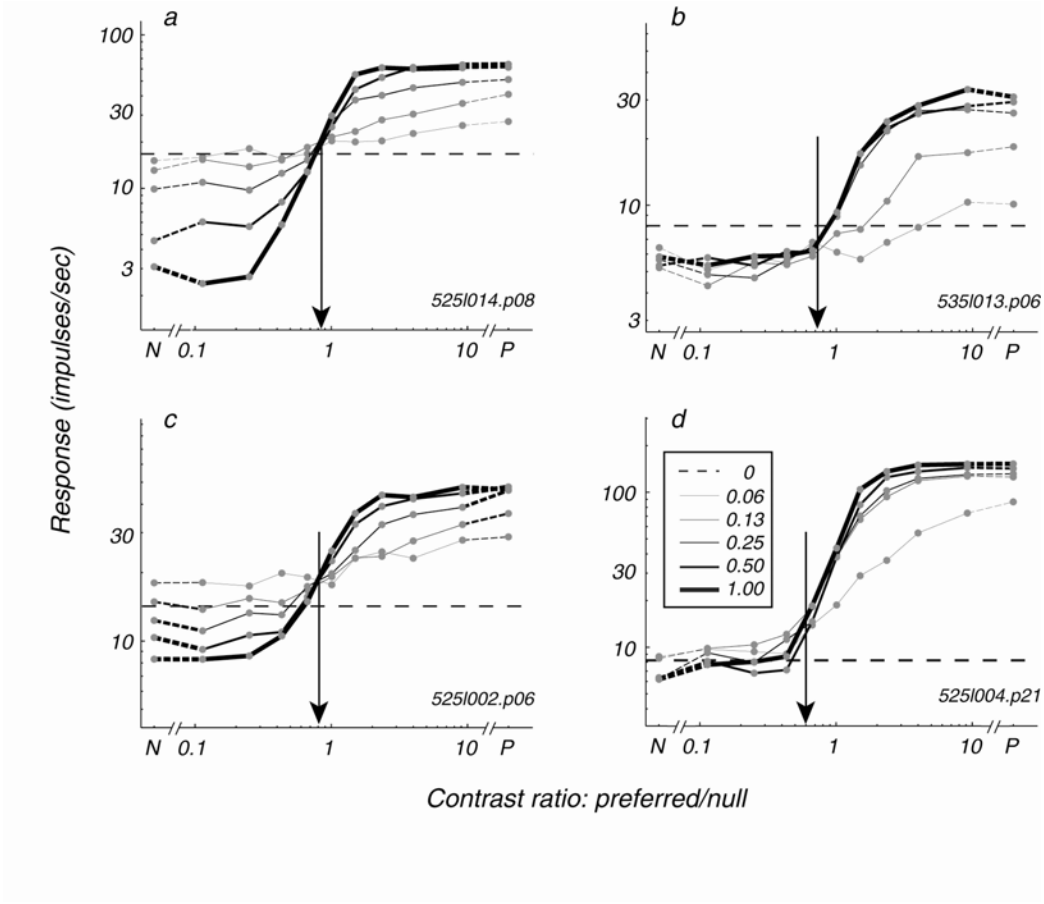

Figure 4-4: A sample of the range of responses observed in MT neurons to the counterphase family stimuli. In all panels, the gray arrows indicate the crossing ratio. As in figure 4-1, stimuli were presented in constant total contrast blocks and line thickness indicates the total contrast of the stimulus.

Baseline response

The steady-state baseline response was measured by recording the response to a gray screen over 15 seconds. In addition to recording a steady-state baseline, for some cells we sampled the baseline response during each total contrast presentation by interleaving short (320 msec) gray stimuli into each total contrast block. For many cells, the response during these "integrated" blanks was lowest at the highest total contrasts. This behavior is indicative of an adaptation signal that decreases the response of the cell in proportion to the total contrast of the stimulus. We found cells for which the suppression below steady-state baseline in response to a null grating could completely be explained by the reduction of the integrated baseline (figure 4-5a). For these cells, the suppression observed in the null direction is not due to a directionally tuned suppressive signal, rather, suppression is simply the absence of excitation and a baseline reduction at high contrasts. We also found cells that were clearly suppressed below both integrated and steady-state baseline by null drifting stimuli (figure 4-5b). For these cells, baseline adaptation may have recovered with a timecourse shorter than the 320 msec integrated blank period. Alternatively, an active null suppressive component may exist in these neurons.

To summarize, the counterphase family experiments revealed a number of interesting properties of MT neurons that constrain descriptions of their mechanism. First, curves collected at different total contrasts crossed at a single

point when plotted as response versus the ratio between preferred and null contrasts. The preferred:null ratio corresponding to this point (the crossing ratio), varied with the direction of the non-preferred grating, indicative of a directionally tuned suppressive signal. Second, strong saturation at high preferred and null contrasts was consistently observed in MT even when the contrast energy was held constant, suggestive of a nonlinearity that deviates from classical normalization model predictions. Third, the symmetry of "fanning" about the crossing point varied from cell to cell and was independent of the location of the crossing ratio. Fourth, the baseline firing rate tended to vary inversely with the total contrast of the stimulus, suggesting the existence of an adaptation to total contrast in MT. Finally, the decrement below steady-state baseline by a high-contrast null drifting stimulus was incapable of silencing MT neurons, despite the strong effect the null grating had on an excitatory stimulus. This suggests that the excitatory and suppressive signals are not pooled linearly.

Model fits

The salient characteristics of responses to the counterphase family stimuli constrain the class of models that can describe the data. We considered a number of models before arriving at one that provided a good description of all the cells we recorded.

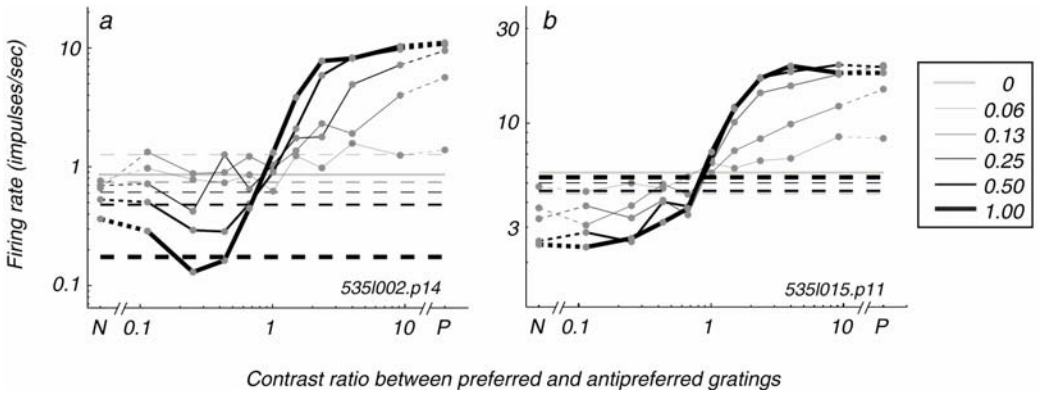

Figure 4-5: Suppression relative to baseline responses. Blank (gray) stimuli were integrated into each block of constant total contrast. Dashed line thickness indicates the total contrast of the stimulus block in which the blank was inserted; the solid gray line indicates the steady-state baseline response (collected in response to 15 seconds of a gray screen). a) A cell for which the integrated baseline recorded within each constant total contrast block varied inversely with total contrast and the suppression below steady-state baseline could be explained by this decrement. b) A cell for which null drifting gratings suppress the responses below both the absolute and integrated baseline response.

The crossing point is the most salient characteristic of MT responses to the counterphase family. As we have described above, the crossing point represents the balance of excitation and suppression in these neurons. Suppression can take on many mathematical forms; the existence of a crossing point is most easily described as a weighted difference between the contrast of the signal in the preferred direction (P) and the contrast of the signal in the null direction (N):

$$R_1 = \lfloor P - \alpha N \rfloor$$

The weight α determines the balance point of excitation and suppression and determines where the crossing point falls along the x-axis. The output of this difference operation is then rectified (indicated by the half brackets), a step that is crucial to account for the local nature of the null suppressive signal in MT (described below). To account for the saturation at high contrasts, the output of the rectified difference is self-normalized by a divisive process and a scalar β is applied:

$$R_{1norm} = \beta * R_1^{\varepsilon} / (R_1^{\varepsilon} + \sigma_N)$$

We fit a different σ_N to each total contrast curve to account for the adaptation observed during blocks of total contrast (Ohzawa et al 1982).

Motivated by the results shown in figure 4-5, we include a baseline response that is inversely proportional to the total contrast (C_{tot} = P + N). The decrement in baseline with increasing contrast is described as a divisive process

to account for the inability of a null drifting stimulus to completely silence the response of these cells:

$$R_2 = \delta / (\gamma * C_{tot} + 1)$$

The response of the MT neuron is the combination of the normalized signal R_{1norm} and the baseline R_2:

$$R = R_{1norm} + R_2$$

In the case of the regular counterphase family stimulus set, the model includes 10 parameters (α, β, ε, δ, γ, and $\sigma_{1\text{-}5}$) fit to 56 data points.

The model fit to an example cell (from figure 4-2) is shown in figure 4-6 as response plotted against the preferred component of the stimulus (figure 4-6a) and as a function of the ratio between preferred and null contrasts (figure 4-6b). For this cell, the weight of the subtractive suppressive term α was fit as 0.54. The difference between preferred and null signals produced negative firing rates at high contrast N; the differenced signal was then rectified to zero (R_1). Suppression of responses below baseline at high contrast N arises from adding a baseline parameter back into the model that is inversely proportional to contrast (R_2). The saturation of the response at low contrast ratios is thus produced by the rectification stage. Saturation at high contrast is produced through self-normalization (R_{1norm}), determined at each total contrast by the σ fit for that curve. Figure 4-6c plots σ as a function of the total contrast block; this cell was

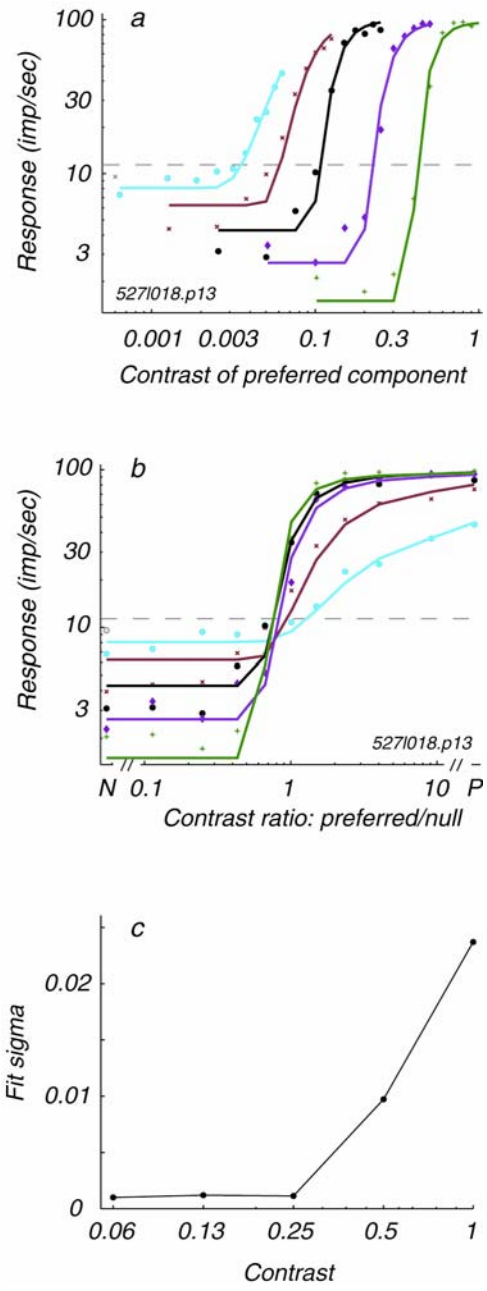
a
100
30
10
3
Response (imp/sec)
527l018.p13
0.001 0.003 0.1 0.3 1
Contrast of preferred component
b
100
30
10
3
Response (imp/sec)
527l018.p13
N 0.1 1 10 P
Contrast ratio: preferred/null
c
0.02
0.01
0
Fit sigma
0.06 0.13 0.25 0.5 1
Contrast

Figure 4-6: Model fits. A parametric model including rectified subtraction, self-normalization, and a baseline inversely proportional to firing rate fit to the data shown in figure 4-2 (see Results). a) Response plotted against the contrast of the preferred component of the stimulus, as in figure 4-2c. b) Response plotted against the ratio of preferred and null stimuli, as in figure 4-2d. The data and fits for each constant contrast block are indicated by a different color: 100% - green; 50% - pink; 25% - black; 12.5% - red; 6.25% - cyan. Parameters fit for the model were: α: 0.54; β: 95.8 ε: 2.64 δ: 11.47; γ: 6.66. c) The sigma fit to each total contrast curve as a function of total contrast.

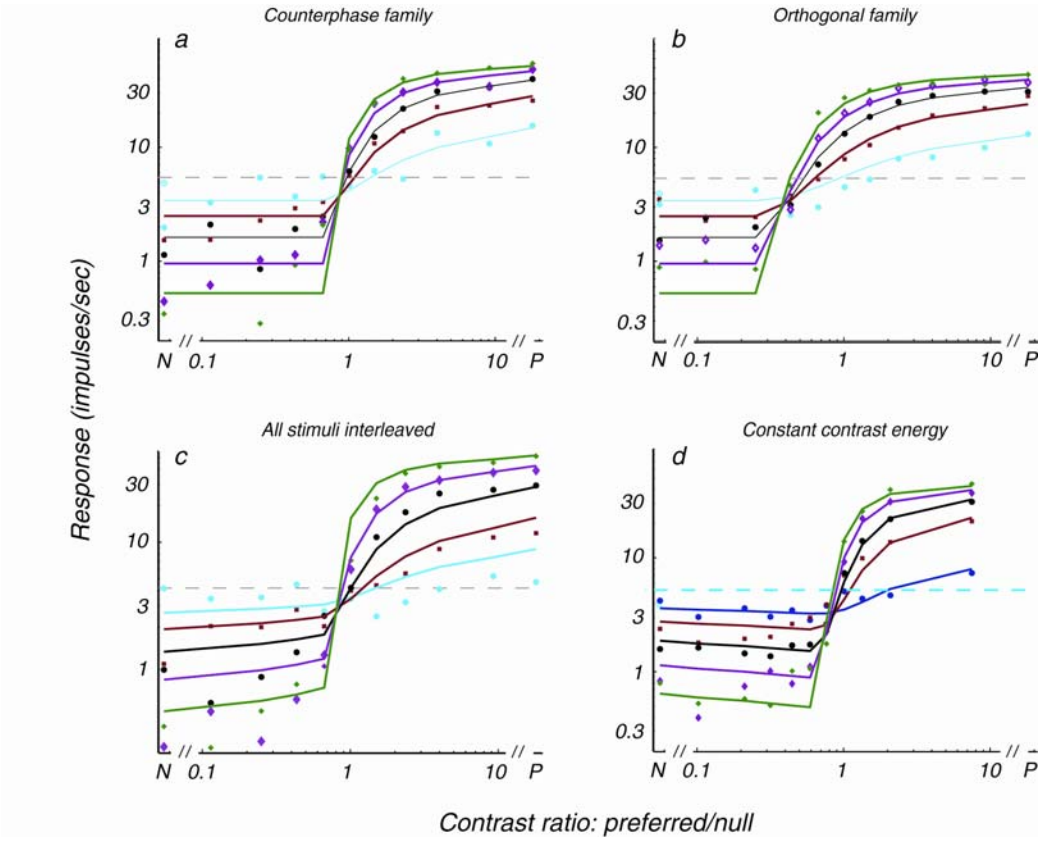

Figure 4-7: Model fits to the counterphase family variants shown in figure 4-3. Plots are shown with the same convention as figure 4-6. a) Counterphase family stimuli. Parameters fit for the model were: α: 0.7; β: 64.2 ε: 1.56 δ: 5.39; γ: 9.25, $\sigma_{1\text{-}5}$: 0.06 0.06 0.08 0.14 0.25. b) Responses of the same cell when the null drifting mask was replaced with an orthogonal mask. The model was fit with all parameters held to the values fit for panel a with the exception of α: 0.29 and β: 54.0. c) Responses of the same cell with all 55 stimuli interleaved. Data were fit with the same parameters as panel a, but with β: 54.8 and a single σ: 0.14. d) Responses of the cell when contrast energy was held constant. Parameters for the fit α: 0.62; β: 49.3 ε: 2.14 δ: 5.29; γ: 10.1, $\sigma_{1\text{-}5}$: 0.01 0.01 0.01 0.02 0.06.

typical of many neurons in that sigma increased with increasing contrast after a threshold (see figure 4-12b for a population summary).

Direction tuning of the suppressive signal is captured in the model by the weight of α. Figure 4-7 shows the fits for the cell illustrated in figure 4-3. In the case of a null drifting non-preferred stimulus, α was fit as 0.7 (figure 4-7a). When the model was fit to data collected with an orthogonal non-preferred stimulus, and all the parameters except the inhibitory weight α and the scalar β held constant (β was allowed to vary to account for slight variations in firing rate due non-interleaved stimulus conditions), the data were well fit and α decreased to 0.29 (figure 4-7b).

When presenting the stimuli in constant total contrast blocks, the cell is expected to adapt to the total contrast of the block, captured in our model by the different σ fit to each curve. When all the stimuli are interleaved, we might predict that a single σ would account for all the curves. Figure 4-7c shows an interleaved stimulus set fit in this manner; the data are well fit with a single sigma of 0.14. The data were also well fit in the case of constant contrast energy (using multiple σ; figure 4-7d).

The fanning of the excitatory-dominated response is controlled in the model by the scalar β whereas the fanning of the suppressive-dominated response is controlled by the baseline term R_2. The independence of these two

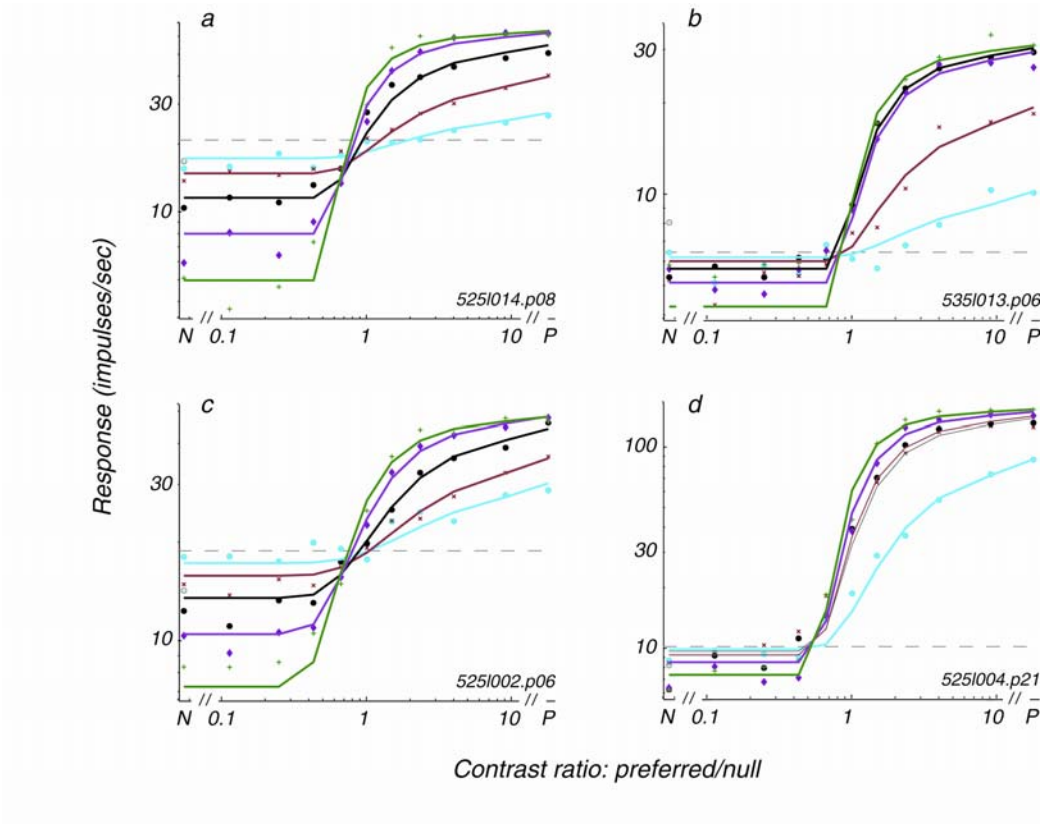

Figure 4-8: Model fits to the same cells shown in figure 4-4 with a range of asymmetries about their crossing points. Plots are shown with the same conventions as figure 4-6.

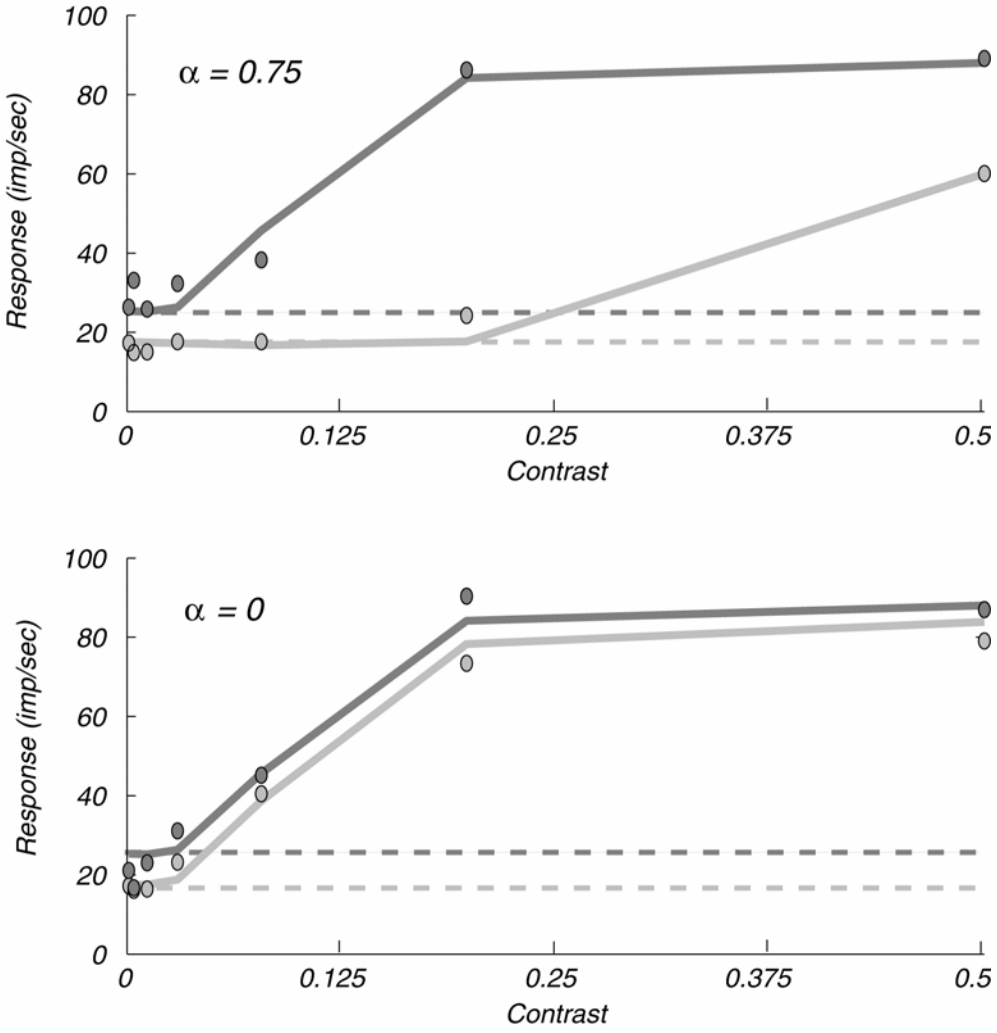

Figure 4-9: Fits of the model to the two patch experiment applied to the example cell shown in figure 4-1. Top: contrast response function collected alone (dark gray) and in the presence of a null pedestal at the same location (light gray). Bottom: data plotted with the same convention as in the top panel but for a non-localized null pedestal (see figure 4-1 for details). All the data points shown were fit simultaneously and the same parameters applied to both conditions with the exception of the subtractive weight α (fit α are labeled). Other parameters were fit as: β: 71.0 ε: 3.0 δ: 25.4; γ: 0.89, σ = 0.001.

terms allowed the model to fit data with a range of asymmetries about the crossing point (figure 4-8).

The model also provided a good account of the behavior observed in the two patch experiments (figure 4-9). We fit the four contrast response functions obtained from this experiment simultaneously (overlapping and separated configurations in the presence and absence of the null pedestal) by forcing the parameters in the overlapping and separated condition to be the same with the exception of the subtractive weight α. In the overlapping condition, α was typically set to a large value (e.g. in figure 4-9a $\alpha = 0.75$) whereas in the separated condition the α parameter was usually set to a much smaller value, often zero (e.g. figure 4-9b). Thus in the separated configuration, all that remained to account for the decrement in response below baseline to the null patch was the divisive effect that total contrast had on the baseline R_2 term. Rectification of the subtractive signal is crucial in accounting for the two-patch data; otherwise the subtractive process would act globally across the receptive field. Taken together, the two-patch and counterphase experiments provide strong evidence for two independent components of the null suppressive signal in MT.

4.2.3 V1

Both the two-patch and counterphase family experiments suggest a local, subtractive signal that is rectified before arriving at the soma of MT neurons. Such

behavior would be observed if the subtractive signal acted not in MT but in earlier visual areas such as V1. MT projecting V1 neurons are believed to be a subpopulation strongly tuned for direction (Movshon & Newsome 1996), thus we limited our investigation in V1 to strongly direction selective cells. In V1, we found neurons whose responses to the counterphase family stimuli were strikingly similar to those observed in MT. Figure 4-10a and 4-10b show the responses of two such neurons. Both have a clear crossing point corresponding to a preferred:null ratio of 0.6-0.7, fanning about this point, saturation, and suppression below steady-state baseline. For the cell in figure 4-10b, we observed a integrated baseline response that decreased with total contrast; this baseline decrement explained the majority of the suppression below steady state baseline in response to null-dominated stimuli (compare with figure 4-5). Within V1, we also found strongly direction-selective neurons that did not display a crossing ratio (figure 4-10c). As expected, neurons that failed to produce a clear crossing ratio also had small excitatory responses to null stimuli (figure 4-10d). Some neurons with very weak excitatory responses to null stimuli also produced clear crossing ratios (e.g. figure 4-10b).

To determine whether the null suppressive signal in V1 is tuned for direction, we substituted the null non-preferred stimulus with gratings drifting at different directions (similar to the experiment shown in figure 4-3b). Similar to the responses of MT neurons, the location of the crossing ratio varied with the direction of the non-preferred stimulus. For the example cell shown in figure 4-

11, an orthogonal non-preferred stimulus produced a crossing ratio of ~0.3. As the direction of the non-preferred stimulus approached the direction opposite the preferred, the crossing ratio increased to the point of approximately balanced excitation and suppression (crossing ratio 0.9). These results present strong evidence for the existence of a directionally tuned suppressive signal in the subpopulation of V1 neurons with crossing ratios.

Histograms of the crossing ratios across the population of MT and V1 neurons we recorded are displayed in figure 4-12a. The geometric mean crossing ratio across the MT population is 0.88, indicating that excitation tends to slightly outweigh suppression in MT. For the V1 neurons in which we observed crossing ratios, the geometric mean crossing ratio was slightly lower, 0.80.

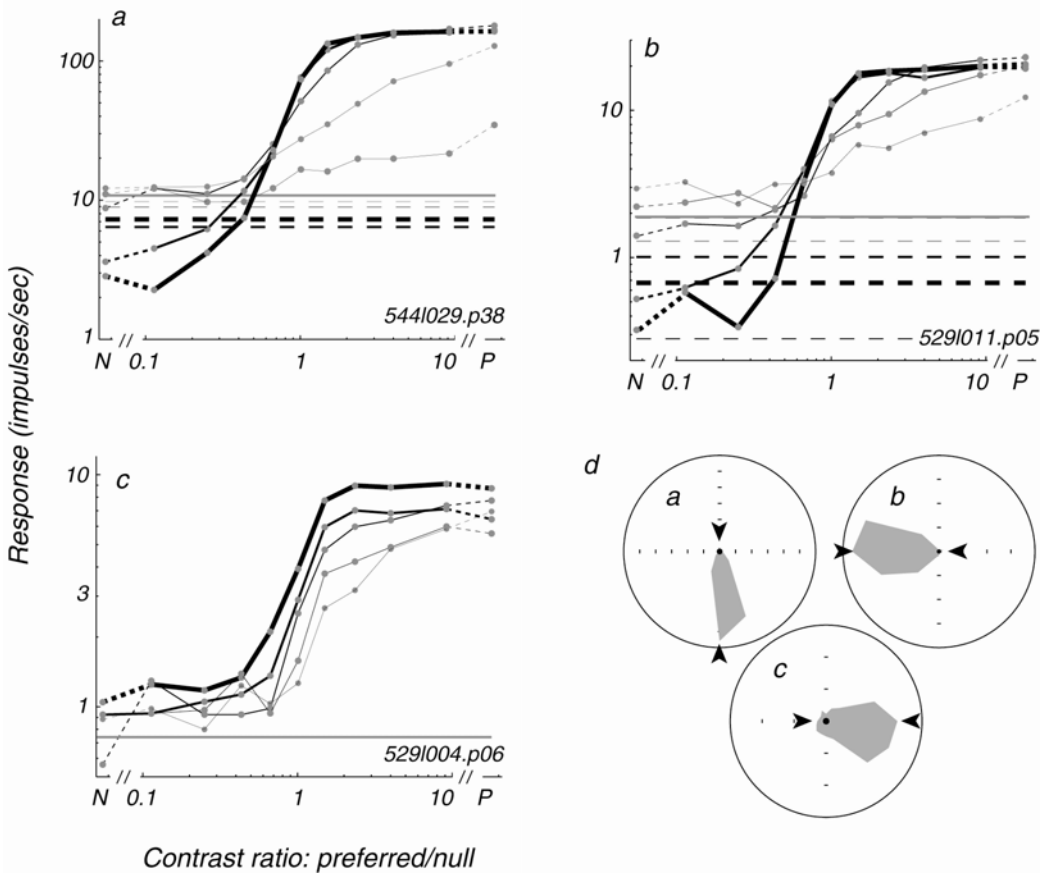

Figure 4-10: V1 responses to the counterphase family stimuli. a, b) Two cells representative of the strongly directional V1 neurons with behavior similar to that observed in MT. Line thickness corresponds to total contrast and dotted lines indicate integrated baseline firing rate with the same convention as figure 4-5. c) A cell representative of V1 neurons without a crossing ratio. d) Direction tuning curves for the cells shown in a-c. Arrowheads indicate the preferred and null directions used in the counterphase family experiments. Cells without crossing ratios displayed an excitatory response to stimuli drifting in the null direction (n=7). Cells with crossing ratios were either suppressed below baseline by null stimuli or displayed weak excitatory responses (n=11).

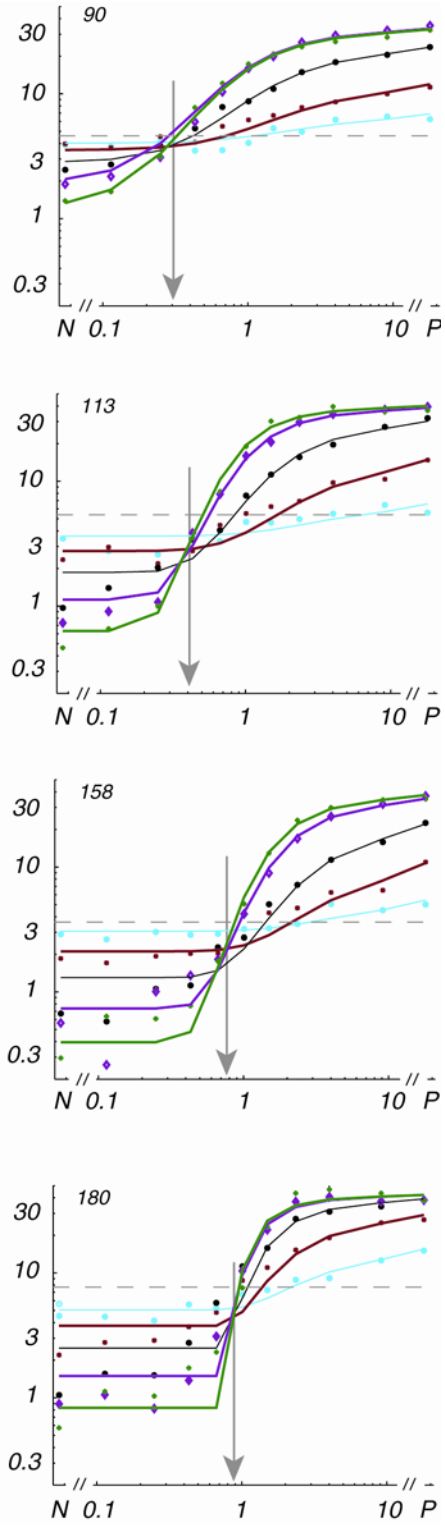
90
30
10
3
1
0.3
N 0.1 1 10 P
113
30
10
3
1
0.3
N 0.1 1 10 P
158
30
10
3
1
0.3
N 0.1 1 10 P
180
30
10
3
1
0.3
N 0.1 1 10 P

Figure 4-11: Direction tuning of the V1 null suppressive signal. The counter-phase family stimuli with non-preferred gratings drifting 90, 113, 158, and 180 degrees from the preferred grating. Gray arrows and labels indicate the crossing ratio for each cell. Data and model fits are shown with the same convention as figure 4-6. α values: 0, 0.13, 0.26, and 0.71 respectively.

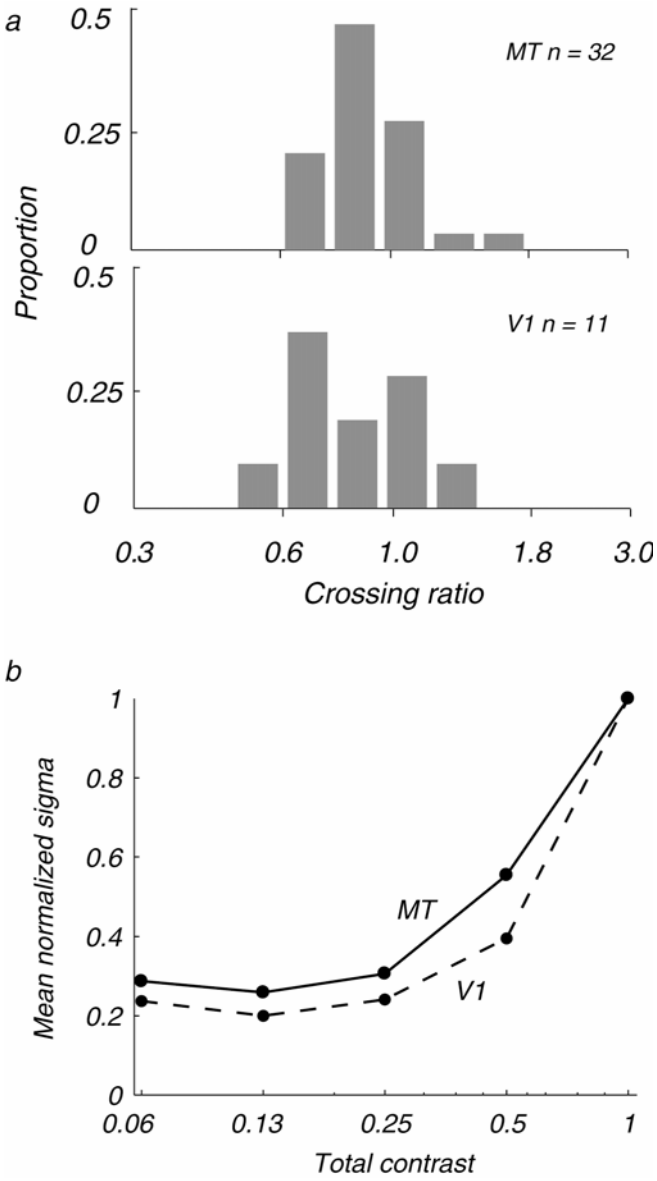

Figure 4-12: Population summary. a) Crossing ratios observed in MT and V1. Geometric means: 0.88 (MT) and 0.80 (V1). b) The sigma fit to each total contrast curve as a function of total contrast after normalizing the sigma at total contrast = 1 for each cell before taking the population averages.

4.3 Discussion

We present evidence in this paper for two components of the null suppressive signal in MT. The strongest component of this signal can only act upon co-localized excitation within large MT receptive fields. In our model, this signal produces the characteristic crossing point in the responses to the counterphase family stimuli and it is tuned for the null direction. We account for these behaviors in the model by representing the signal as a local, weighted subtraction between excitation and suppression followed by rectification. In the sense that the signal is tuned for the null direction and well represented by a subtractive process, it is most similar to the "motion opponent" signal postulated for these neurons.

Our results suggests that a second signal in MT produces the suppression below baseline in response to a null moving stimulus. This signal is not specifically tuned for the null direction; it is the product of adaptation to the total contrast of the stimulus and results in a reduction of the baseline firing rate. In our model, this signal describes the global suppressive effect observed in our two patch experiments. In V1, tonic hyperpolarization proportional to the total contrast of a preferred drifting sinusoidal grating has been directly recorded intracellularly in (Carandini & Ferster 1997).

To determine whether the local, subtractive component of the null suppression in MT reflected a computation performed in V1, we presented strongly

directional V1 cells with the counterphase family stimuli. We found that cells with the strongest directionality produced responses indistinguishable from those in MT, including a crossing ratio and baseline hyperpolarization proportional to total contrast. We also found that the strength of the suppressive signal, reflected by the crossing ratio, depended on the direction of the non-preferred stimulus and was strongest in the null direction. These results suggest that the "motion opponent" computation is in fact performed in some V1 neurons.

Is the motion opponent computation performed in V1 exclusively? As described in the introduction, it is unlikely that the motion opponent computation occurs in MT: the lack of feedforward inhibitory projections from V1 to MT require a subpopulation of inhibitory interneurons with receptive fields the size of V1 neurons and no such cells have ever been reported. If the motion opponent computation were to occur entirely in V1, the population of neurons projecting from V1 to MT would have to be comprised primarily of the highly direction-selective neurons that produce crossing ratios in their responses to the counterphase family stimuli. Movshon and Newsome (1996) identified MT projecting V1 neurons through antridromically activation resulting from electrical stimulation of MT. Most (60%) of the MT projecting V1 neurons they identified were inhibited below baseline by a null moving stimulus; 90% of MT projecting V1 neurons were strongly directional (direction index > 0.8). These characteristics are consistent with the types of cells that produced crossing ratios

to the counterphase family stimuli. Based on this evidence, we conclude that the motion opponent computation is likely exclusive to V1.

The Simoncelli-Heeger (SH) model of motion processing (1998) proposes that the direction computation occurs in two stages. First, contrast invariant direction tuning is conferred in a subset of V1 neurons through an untuned, divisive gain control signal. The direction tuning computed in V1 is then further sharpened in MT. MT neurons receive excitatory input from V1 neurons preferring one direction and inhibitory input from cells with the opposite direction preference (motion opponency). In MT, the gain of the cell is again adjusted through an untuned, divisive signal. A grating drifting in the null direction induces three types of suppression in the SH model: 1) divisive gain control in V1 2) motion opponency in MT and 3) divisive gain control in MT.

Our data are consistent with a multi-stage computation of direction but suggest modifications to this model. First, the motion opponent computation should be included in V1, not MT. If this modification were included in the model, the model would correctly predict stronger null suppression of co-localized versus non-localized excitation. Suppression below baseline could be conferred in the SH model if the gain control signal were allowed to act upon baseline firing rates. But the saturation of responses to constant contrast energy (figure 4-7d) suggests that the MT gain control signal should be modified to include a non-linear saturating process, either arising through a self-normalization or normalization via an alternative compressive nonlinear mechanism.

5 Discussion

The work presented in this thesis relates to a larger body of work whose goal is to "understand" neural processing in terms of the computations performed at each stage. In the work presented in chapter 2, I focused on signal transmission in primary visual cortex and found surprising evidence that the reliability of V1 neurons decreases during development. In chapters 3 and 4, I focused on signal representation and computation in V1 and MT using two different approaches. In chapter 3, I presented an extension to classical spike-triggered techniques to recover functional models of V1 neurons; application of this method revealed that V1 neurons have more subunits than are included in classical models. In chapter 4 I probed the role of suppression in sharpening direction selectivity by presenting stimuli comprised of preferred and null drifting sinusoidal gratings. There, I presented evidence that a directionally tuned suppressive (motion-opponent) signal likely operates in V1 and is merely reflected in MT responses. Below I develop the common themes of these chapters in terms of their experimental design and the physiological mechanisms they suggest.

5.1 Comparing responses mapped with gratings and stochastic stimuli

Simple and complex cell subunits:

Chapter 3 presents evidence for many more subunits in V1 simple and complex cells than standard models suggest. These neurons have been studied in great detail; why haven't these subunits been revealed before? This question can be viewed from the perspective of asking what stimuli and analysis techniques can distinguish the standard models from alternatives. Physiologists have been aware of the continuum from simple-to-complex for some time, based upon the responses to isolated bars and sinusoidal grating stimuli (Skottun et al 1991) and have held the intuition that cells on the simple/complex border are comprised of both asymmetric and symmetric subunits. The spike-triggered characterization presented in chapter 3 confirms this intuition. In contrast, sinusoidal grating stimuli provided no hints of the multiple spatially-shifted subunits we revealed from complex cells. Identification of subunits along an axis requires the presentation of multiple stimulus combinations along that axis and comparison of a neuron's differential response to those stimuli; a periodic grating stimulus such as a sinusoid fails to provide such variation. Two-bar interaction experiments (Emerson et al 1987, Gaska et al 1994, Livingstone & Conway 2003, Movshon et al 1978, Szulborski & Palmer 1990) satisfy the requirement of variation and effectively map the same covariance matrix as our dense bar

stimulus. The crucial difference between the two-bar interaction experiments and the spike-triggered covariance (STC) method we applied was in the application of the principal components analysis (PCA) to this data structure. PCA transformed the map of second-order response dependencies between stimulus dimensions (the covariance matrix) into a small number of linear axes that defined a subspace in the high-dimensional space of all stimuli. Standard models of complex cells predict that the neuron's response will reside within a two-dimensional subspace whereas our results suggest that this space is much higher dimensional.

Mapping the interaction of excitation and suppression:

Portions of chapters 3 and 4 examine the combination of excitatory and suppressive influences in V1. The counterphase family experiments presented in chapter 4 were designed to test the suppressive signal in strongly directionally selective neurons; below I compare the responses to this grating stimulus with the results obtained by the spike-triggered characterization of neurons with similar direction selectivities (chapter 3).

Our spike-triggered analysis identified the linear subspace of the higher dimensional stimulus space that impact a neuron's response, including both excitatory and suppressive influences. The joint distribution of raw stimuli across the excitatory and suppressive pooled signals is shown in figure 5-1 (top left) for the example cell from figure 3-4. For this cell, the excitatory and suppressive

signals were sampled with a nearly continuous distribution that peaked at moderate values of excitation and suppression. In contrast, the counterphase family stimulus was selected from 55 discrete combinations of a preferred-drifting and null-drifting grating (the same 11 ratios of preferred and null contrast at 5 contrasts; figure 5-1 top right).

In both chapters 3 and 4, I presented parametric models that were well fit to responses to the two stimulus classes. The model that best accounted for the spike-triggered data included a saturating excitatory function as well as subtractive and divisive suppressive terms. The model fit to the counterphase data included a rectified difference between excitation and suppression, self-normalization, and a baseline inversely proportional to contrast. We can use these models to fill-in the unsampled excitatory-suppressive combinations for each experiment; those surfaces are shown in figure 5-1 (second row). The axes for the spike-triggered covariance map are normalized and extended to include the prediction of the response to a 100% contrast sinusoidal grating at the same spatial and temporal frequencies used in the counterphase experiment (axes labeled E_c and S_c).

Despite the fact that the experiments were performed on the same neuron, the two surfaces are quite different. The surface determined for the spike-triggered characterization contains gentle response contours that radiate from the x-axis and change in slope (figure 5-1, second row, left). Along the surface determined for the counterphase family, contours of constant firing rate are

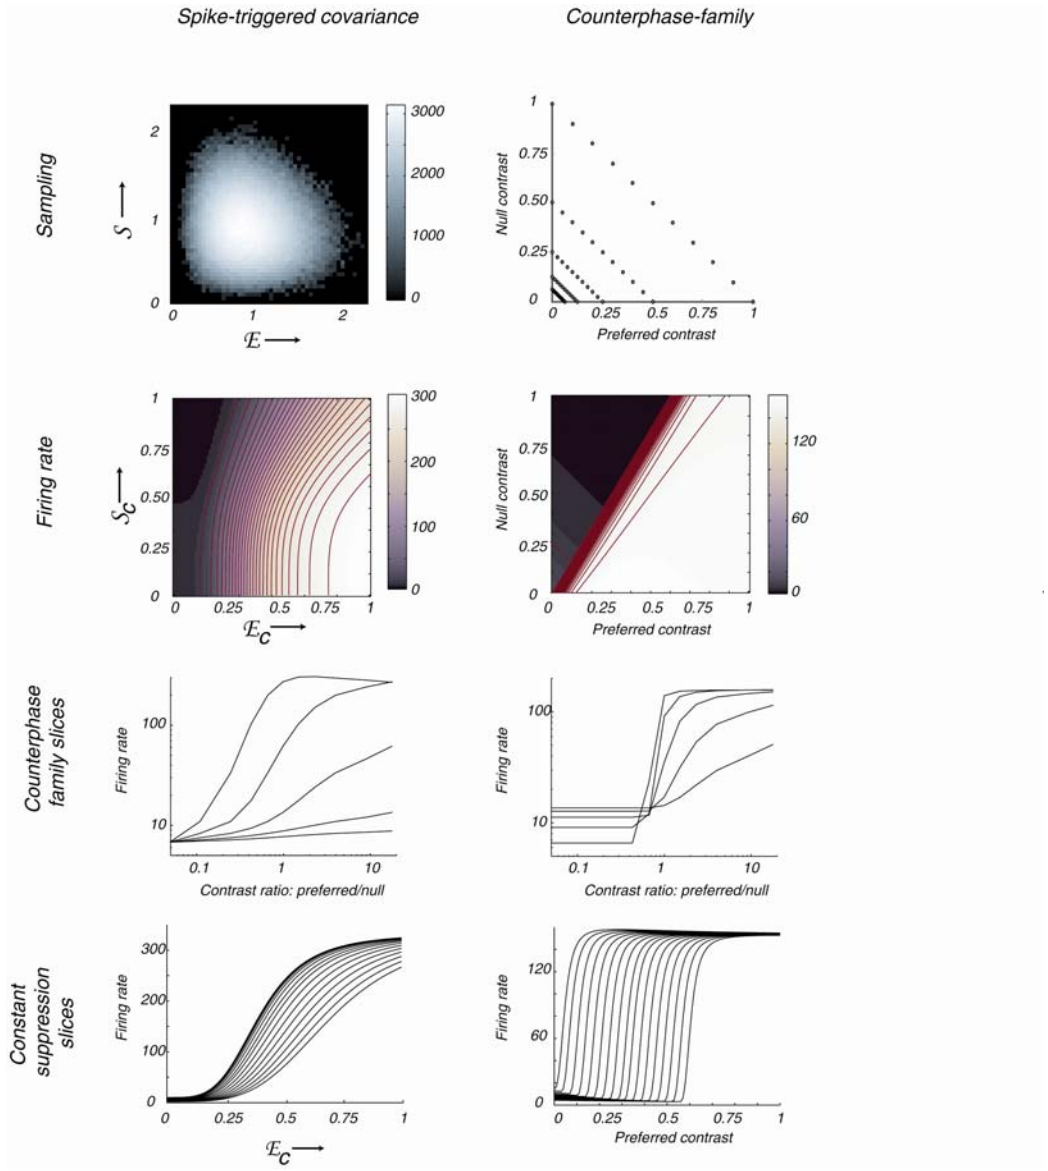
Spike-triggered covariance
Counterphase-family
Sampling
Firing rate
Counterphase family slices
Constant suppression slices
Null contrast
Preferred contrast
Firing rate
Contrast ratio: preferred/null

Figure 5-1: Comparison of the results from the spike-triggered characterization and the counterphase family experiments. Top row: Distribution of data collected in each experiment across the excitatory and suppressive axes. In the spike-triggered characterization, sampling density is shown proportional to intensity for a range of outputs of the excitatory and suppressive pools (left). In the counterphase experiments, discretely sampled preferred and null grating combinations are indicated by the dots (right). Second row: Firing rate as a function of the output the excitation and suppression determined by parametric models fit to each data set (see chapters 3 and 4 for details). The axes for the spike-triggered surface (labeled Ec and Sc) are normalized such that 1 corresponds to the predicted firing rate to a full contrast sinusoidal grating with the same parameters used in the counterphase experiments (E output 3.2; S output 1.8). Equally spaced contours of constant firing rate are indicated in red. Third row: the two surfaces sampled at the points used in the counterphase family experiments (top right). Bottom row: Horizontal slices across increasing excitation at constant levels of suppression.

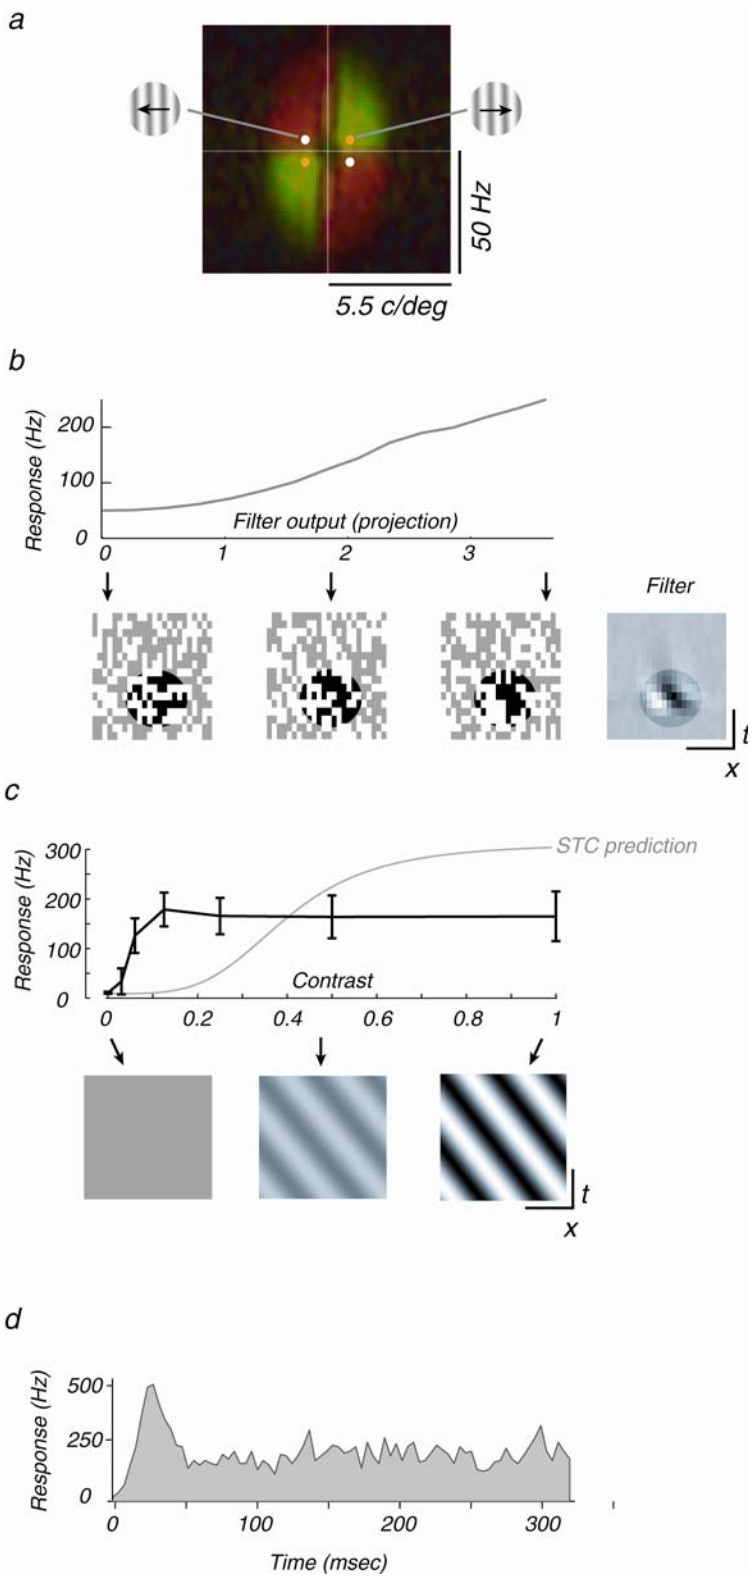
a
50 Hz
5.5 c/deg
b
Response (Hz)
200
100
0
Filter output (projection)
0
1
2
3
Filter
t
x
c
Response (Hz)
300
200
100
0
STC prediction
Contrast
0
0.2
0.4
0.6
0.8
1
t
x
d
Response (Hz)
500
250
0
0
100
200
300
Time (msec)

Figure 5-2: Comparison of the spike-triggered and counterphase family experiments II. a) Pooled excitatory (green) and suppressive (red) frequency spectra as a weighted-sum of the amplitude spectra for the set of filters revealed by STC for each class. The spatial and temporal frequency selected for the counterphase family experiments are shown for the preferred grating (yellow) and the null grating (white). b) Firing rate as a function of the output of the strongest excitatory filter (the projection of the filter and a stimulus) revealed for this cell by spike-triggered covariance. Instantiations of the binary bar stimulus that occurred at three points along this axis are shown. The strongest (non-zero) regions of the filter are highlighted as are the same regions of the bar stimuli for comparison. c) Response plotted as a function of contrast for an otherwise optimized drifting sinusoidal grating. Standard error bars are shown. Also shown are stimuli taken from three points along the contrast axis for comparison with b. The prediction of the response to the same contrasts by the spike-triggered model is shown in grey. d) A post-stimulus time histogram of the same neuron's response to 320 msec of a full contrast, optimized drifting sinusoidal grating to illustrate the onset transient. The first 40 msec have been removed to adjust for the latency of the cell.

clustered about a steeply sloped firing rate precipice that follows a constant ratio of excitation and suppression (figure 5-1, second row, right). For a closer look atthese responses, figure 5-1 (third row) shows the response surfaces sampled at the points used in the counterphase experiment (top row, right) plotted in the same manner as the data were presented in chapter 4 (firing rate as a function of the ratio between excitation and suppression). The data taken from the counterphase family experiment show the characteristic behaviors reported in chapter 4, including a point at which all the curves cross, steeply sloping functions that fan about this point, and saturation at high and low values of excitation and suppression. The data taken from the spike-triggered surface fail to cross at a single point and the functions are much shallower. Figure 5-1 (bottom row) shows slices of increasing excitation taken at different values of constant suppression, similar to figure 3-8. Again, the slices taken from the counterphase experiment are steeper and saturate more than those for the spike-triggered characterization. Note that these differences are not due to differences between the models fit to the two data sets, rather they reflect differences in the response properties of this neuron to the two stimuli (data not shown).

Why does the neuron respond so differently to the two stimulus sets? One difference between the two experiments is related to the nature of the excitatory and suppressive stimuli that were used in the characterization. Figure 5-2a compares the frequency spectra of the stimuli used in each experiment. Shown are the pooled frequency spectra for the excitatory and suppressive filters

in the spike-triggered characterization (green and red, respectively, taken from figure 3-4). The two signals cover a similar range of spatial and temporal frequencies but are tuned for opposite directions of motion. Also shown are the spatial and temporal frequencies of the gratings used in the counterphase family experiment performed on the same cell (preferred drifting grating: yellow; null drifting grating: white). The excitatory and suppressive frequencies chosen for the counterphase experiment were in fact deemed excitatory and suppressive by the spike-triggered characterization, although the location of the null-drifting grating at the edge of the pooled suppressive spectra may suggest that a more effective suppressive stimulus could have been chosen. Some models of contrast adaptation predict stronger adaptation to concentrated versus diffuse spatiotemporal excitatory energy (Carandini et al 2002), but the dramatic differences in the slope and saturation of these surfaces are unlikely to be explained solely by this effect.

More likely, the differences arise from the different means of manipulating excitability in the two experiments. To illustrate this difference, figure 5-2b shows firing rate as a function of the output of the strongest excitatory filter revealed by spike-triggered covariance. Below, the spatiotemporal structure of the filter is shown along with instantiations of the bar stimulus at three points along this axis. The subregion of the filter that is most relevant in determining its response is highlighted, as are the same regions of the bar stimuli for comparison. Stimuli that most resembled the filter produced the highest firing rates from the

neuron (right); stimuli that bore no resemblance to the filter produced the lowest responses from the cell (left). (In actuality, the full model of the neuron contained seven other excitatory filters and seven suppressive filters; the response to a stimulus depended on the output of all 15 filters). For comparison, the traditional sine-grating contrast response function for the neuron is shown in figure 5-2c (black) and stimuli from three points along this axis are indicated. Comparison of a 100% contrast grating stimulus with the highlighted region of the rightmost stimulus in 5-2b reveals that they are quite similar. However, comparison of the 0% contrast stimulus (figure 2c, left) with the bar stimulus bearing no resemblance to the excitatory filter (figure 5-2b, left) highlights the difference between these two experiments. In the counterphase experiment, excitation was titrated by manipulating the contrast of the stimulus; in the bar experiments, the excitation was manipulated (through random selection of stimuli) by introducing ineffective spatiotemporal structure while holding the stimulus intensity constant. Although the stimuli can be quite different in nature (e.g. in contrast), the model assumed by the spike-triggered characterization assigns the same firing rate to stimuli that produce the same output from the excitatory pool. Thus a 50% contrast stimulus with maximal spatiotemporal similarity with the pooled filters is assigned the same firing rate as a 100% contrast stimulus with half-the spatiotemporal similarity.

The contrast response function predicted by the full STC model is shown in figure 5-2c in gray. The cell is much more sensitive to low contrast gratings

than the spike-triggered model predicts, suggestive of a difference in the contrast adaptation state during the grating versus bar characterization. The difference in gain states may be attributed to masking, a well documented phenomenon in V1 neurons (Dean et al 1981, Morrone et al 1982). Masking refers to the ability of a non-excitatory stimulus (such as an orthogonal grating or spatiotemporal noise stimulus) to reduce the response to a simultaneously presented excitatory stimulus. Masking stimuli have been shown to have a divisive effect on a V1 neuron's response (Carandini & Heeger 1994, Heeger 1992). In the spike-triggered characterization, the titration of excitability through introduction of "noise" may similarly act as a mask to the neuron and reduce the response relative to an equivalent sinusoidal grating at low contrasts.

The responses of the neuron to highly excitatory instantiations of the bar stimulus were approximately two-fold larger than the responses to a high contrast, optimized sinusoidal grating (figure 5-2c). This may be attributed to the time course over which the two stimuli are characterized. The response to sinusoidal gratings represents the average response over a 320 msec presentation whereas the spike-triggered characterization was performed by presenting a new stimulus every 10 msec. Figure 5-2d shows a PSTH of the response to 320 msec of a preferred direction drifting grating. This neuron does respond with a large onset transient for 30 msec that is then reduced approximately two-fold and maintained for the remainder of the stimulus presentation. The spike-triggered model is likely to be more reflective of the onset transient as compared to the

steady state response. Potentially, some of the differences in the responses to the grating and bar stimuli can be ascribed to differences in the interaction between excitation and suppression during transient versus steady-state responses.

To summarize, the differences in the firing rate surfaces during the spike-triggered versus counterphase family characterizations are probably primarily described by differences in the gain states of the neuron during the two characterizations. During the spike-triggered characterization, the non-excitatory noise stimuli act as a mask, thus reducing the responses to preferred stimuli. Because the motivation of the spike-triggered analysis is in part to construct generalized models that can predict the responses to any stimulus, these differences highlight the need to incorporate a contrast gain adjustment into the functional models of these cells. Contrast gain control effects are known to act at many different time scales, making this a challenging undertaking. Differences in the results may also arise from the different time epochs analyzed in the two experiments. To the degree to which this is true, the stochastic stimulus better resembles the operating regime of the neuron during naturalistic viewing conditions.

5.2 Computation in area MT

Two computations are commonly thought to occur in MT: 1) sharpening of direction selectivity via a motion-opponent process, and 2) computation of the

direction of a moving pattern as the intersection-of-constraints of the pattern components. In this thesis I present evidence that motion opponency cannot act globally across the MT receptive field, implying that the suppressive (motion opponent) signal is combined and rectified before reaching the soma of MT neurons. Similarly, the pattern computation appears to require the co-localization of motion signals. In their experiments, Majaj et al (Soc for Neurosci Abstracts, 1999) collected a pattern-direction tuning curve by varying the direction of two gratings presented 120 degrees apart (a plaid stimulus). "Component" cells responded to each of the components of the plaid, resulting in bimodal direction tuning curves. "Pattern" cells responded only to the intersection of constraints of the two gratings together, resulting in unimodal tuning curves. In pattern cells, placement of the second grating in a non-localized portion of the receptive field produced "component" behavior, suggesting that the pattern computation requires co-localization of the two motion signals.

The requirement of spatial co-localization for both the motion-opponent and pattern computation is striking. However, multiple lines of evidence suggest that this similarity should not be taken as evidence of the same physiological locus. First, pattern selectivity is found in an extremely small subpopulation of V1 neurons, if at all (Blakemore 1990, Movshon et al 1985); the antridromically activated V1 projecting MT neurons identified by Movshon and Newsome (1996) were component, not pattern selective. In contrast, motion opponency is found in strongly directional V1 neurons, consistent with the subpopulation of

cells that projection to MT. Second, the lack of feedforward inhibitory projections from V1 to MT requires a subpopulation of inhibitory interneurons with receptive fields the size of V1 neurons to explain a local motion opponent computation in MT. No such neurons have ever been recorded. However, local combination of excitatory signals for the pattern computation could hypothetically occur via spatially specific projections onto the dendrites of MT neurons. Alternatively, the pattern computation may require that the two excitatory signals are preprocessed by the same V1 neurons even if the computation is not completely carried out there (e.g. local normalization or masking of the two signals in V1).

Despite the physiological heterogeneity of cortex, different cortical areas are remarkably similar anatomically. Thus it seems possible that cortical areas instantiate their computations according to a generic formulation. Based upon their modeling efforts in V1 and MT, Heeger et al (1996) proposed a potential computational framework for cortical processing. In their model, neurons in a given area implement a three-stage computation: 1) linear combination of input signals; 2) divisive normalization by the pooled signal of neighboring neurons; and 3) spike-generation. In their model of V1, direction tuning is conferred through 1) linear combination of appropriately arranged spatiotemporal inputs to produce space-time oriented receptive fields; followed by 2) an untuned divisive signal that results in contrast-invariant tuning. In MT, motion opponency and the pattern computation are instantiated in their model by 1) convergence of V1

excitatory inputs with different direction preferences to confer pattern selectivity and inhibitory input from neurons with the opposite direction preferences; and 2) divisive normalization by an untuned signal to confer contrast-invariant tuning (Simoncelli and Heeger, 1998). The results I present here suggest modifications to the MT model that have implications for this generalized cortical computational scheme. First, motion opponency likely occurs in V1, not MT. However, a linear combination of excitatory inputs is still required for the pattern computation in MT. Second, the saturation of MT responses cannot be explained by an untuned divisive normalization signal, but rather are transformed through a sigmoidal nonlinearity that is better described as self- (or directionally tuned) normalization. The utility of self-normalization is unclear as it results in reducing the directionality of a cell at high contrasts. Regardless, these results suggest that the normalization stage of the general computational framework should be reconsidered.

5.3 Feature representation and computation: past and future

Sensory processing begins by deconstructing the physical world into the most basic components. From this rudimentary representation, the brain "reconstructs" an amazingly sophisticated understanding of the world around us. One can seek to understand the brain at many different levels, ranging from proteins and their DNA sequences; to cognition. Retinal physiologists of the 1960s in-

troduced the ideal of understanding neural processing in terms of the computational operations performed at each stage. For the last four decades, visual neuroscientists interested in description at this level have been refining the retinal models first proposed by those pioneers and applying extensions of their techniques to successive stages of the visual pathway. V1 simple cells were first described in 1962 and their line-weighting functions first mapped in 1978 (see the Introduction for references). In 1980, Gabor models were suggested as the functional description of simple cell spatial profiles and in 1985 these models were extended to describe direction tuning. Complex cells were first described in 1962 and their subunits first mapped in 1978. Functional models of complex cells (e.g. the Energy model) were first proposed for these cells in the mid-1980s. My work (2004) fits into this context by proposing a unified, multi-subunit functional model to describe both simple and complex cells in V1. In MT, pattern selectivity was first described in 1985 and a model explaining this nonlinear computation proposed in 1998. Thus far, no attempts at mapping and testing functional models of individual MT neurons have been made. In summary, we have reasonable functional models of the computation performed in the retina, LGN, V1, and the beginnings of such models in MT (although refinements to these descriptions are certainly required before we achieve the goal of constructing functional models that can quantitatively predict the response to any stimulus).

Few would argue that we have made any substantial progress toward understanding the computations underlying the perception of objects beyond V1. Why has the description of these neurons (e.g. V4, IT) proven to be such a difficult problem? If one knows the input to a cell, arriving at a description of the input-output relationship of that cell should be relatively straightforward. The problem, as I see it, is two-fold. The extreme selectivity of neurons in visual areas V4 and IT suggests that their computations are highly nonlinear and nonlinearities are difficult to describe systematically. In addition, our mapping techniques (e.g. spike-triggered characterizations) are designed to describe neurons in terms of a linear process performed *directly on the stimulus* and an *instantaneous* nonlinearity. As we ascend the processing pathway, one can envision that single-stage linear-nonlinear models become increasingly inadequate. Thus our efforts must be focused on designing nonlinear systems techniques that characterize *multi-stage* computation. Sophisticated nonlinear systems analysis techniques such as spike-triggered covariance, coupled with techniques that allow us to probe the input-output relationship of multi-stage computations, may be the key to understanding the neural representation of our visual world.

Appendix: Physiological methods

We recorded from isolated single units in MT and primary visual cortex (V1) of adult macaque monkeys (*Macaca fascicularis* and *Macaca Nemestrina*).

Animals were premedicated with atropine sulfate (0.05 mg/kg) and diazepam (1.5 mg/kg) 30 minutes before the induction of anesthesia with 10.0 mg/kg ketamine. During surgery, anesthesia was maintained with 3% isoflourane in a O_2/CO_2 (98%-2%) mixture. We placed cannaulae in the saphenous veins of both hindlimbs and a implanted a trachea tube. The animal was then mounted in a stereotaxic apparatus and the gas anesthesia discontinued. Anesthesia was maintained with continuous infusion of 4-16 μg/kg/hr of sufentinil citrate mixed in a lactated ringer's solution and 2.5% dextrose throughout the experiment.

We performed a craniotomy and durotomy over the region of interest and placed an agar filled chamber over the region to protect the cortical surface and stabilize the region. V1 was targeted ~7mm posterior to the lunate and ~12.5 mm lateral from the midline. MT was targeted by a 20 degree from horizontal penetration at the same location. Along this trajectory the electrode passed through visual areas V1, V2, and V3, followed by a 5-15 mm stretch of white matter, and finally MT. Initial confirmation of MT was made through physiological properties including receptive field size and eccentricity, strong directional selectivity, and robust responses to moving dots and gratings. In both

MT and V1, cells had receptive fields that ranged from 1-20 degrees eccentricity.

During experiments, the animal was artificially respired and body temperature was maintained with a heating pad. Vital signs (heart rate, lung pressure, EEG, ECG, body temperature, and end-tidal CO_2) were monitored continuously. The paralytic norcuron was administered intravenously at a dose of 0.15 ug/kg/hr mixed in a lactated ringer's solution and 2.5% dextrose to prevent involuntary slow drifts of the eyes. Total fluid intake was maintained at approximately 4-8 mg/kg/hr. Gas permeable contact lenses were used to protect the corneas throughout the experiment. Animals received daily injections of the antibiotic Bicillin and the anti-inflammatory agent dexamethasone. Experiments lasted 4-5 days. At the end of the experiment, animals were sacrificed with an overdose of Nembutal and perfused with 4% paraformaldehyde. Confirmation of recording sites was made through histological identification of electrolytic lesions. All experiments were performed in compliance with the National Institutes of Heath *Guide for the Care and Use of Laboratory Animals* and within the guidelines of the New York University Animal Welfare Committee.

We adjusted the focus of the display through supplementary glass lenses chosen initially to bring retinal capillaries in focus with an othalmoscope. Lens strength was confirmed by maximizing the spatial resolution of neuronal responses.

Single unit activity was recorded using platinum-tungsten microelectrodes (Thomas Recordings, Giessen, Germany). Signals were amplified, band-pass filtered, and fed into a time-amplitude window discriminator. Spike arrival times and stimulus synchronization pulses were stored with a resolution of 0.25 msec.

References

Abbott LF, Rolls ET, Tovee MJ. 1996. Representational capacity of face coding in monkeys. *Cereb Cortex* 6: 498-505

Adelson EH, Bergen JR. 1985. Spatiotemporal energy models for the perception of motion. *J Opt Soc Am A* 2: 284-99

Aguera y Arcas B, Fairhall AL. 2003. What causes a neuron to spike? *Neural Comput* 15: 1789-807

Albrecht DG, Geisler WS. 1991. Motion selectivity and the contrast-response function of simple cells in the visual cortex. *Vis Neurosci* 7: 531-46

Anderson JC, Binzegger T, Kahana O, Martin KA, Segev I. 1999. Dendritic asymmetry cannot account for directional responses of neurons in visual cortex. *Nat Neurosci* 2: 820-4

Anderson JS, Carandini M, Ferster D. 2000. Orientation tuning of input conductance, excitation, and inhibition in cat primary visual cortex. *J Neurophysiol* 84: 909-26

Anzai A, Ohzawa I, Freeman RD. 1999. Neural mechanisms for processing binocular information I. Simple cells. *J Neurophysiol* 82: 891-908

Banks MS, Salapatek P. 1981. Infant pattern vision: a new approach based on the contrast sensitivity function. *J Exp Child Psychol* 31: 1-45

Bialek W, Rieke F, de Ruyter van Steveninck RR, Warland D. 1991. Reading a neural code. *Science* 252: 1854-7

Blakemore C. 1990. Matruation of mechanims for efficient spatial vision. In *In: Vision: coding and efficiency*, ed. C Blakemore, pp. 254-66. Cambridge, UK: Cambridge University Press

Boothe RG, Kiorpes L, Williams RA, Teller DY. 1988. Operant measurements of contrast sensitivity in infant macaque monkeys during normal development. *Vision Res* 28: 387-96

Bradley A, Skottun BC, Ohzawa I, Sclar G, Freeman RD. 1987. Visual orientation and spatial frequency discrimination: a comparison of single neurons and behavior. *J Neurophysiol* 57: 755-72

Brenner N, Bialek W, de Ruyter van Steveninck R. 2000. Adaptive rescaling maximizes information transmission. *Neuron* 26: 695-702

Britten KH, Heuer HW. 1999. Spatial summation in the receptive fields of MT neurons. *J Neurosci* 19: 5074-84

Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA. 1996. A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis Neurosci* 13: 87-100

Britten KH, Shadlen MN, Newsome WT, Movshon JA. 1992. The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J Neurosci* 12: 4745-65

Britten KH, Shadlen MN, Newsome WT, Movshon JA. 1993. Responses of neurons in macaque MT to stochastic motion signals. *Vis Neurosci* 10: 1157-69

Buracas GT, Albright TD. 1999. Gauging sensory representations in the brain. *Trends Neurosci* 22: 303-9

Buracas GT, Zador AM, DeWeese MR, Albright TD. 1998. Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron* 20: 959-69

Burgard EC, Hablitz JJ. 1993. Developmental changes in NMDA and non-NMDA receptor-mediated synaptic potentials in rat neocortex. *J Neurophysiol* 69: 230-40

Carandini M, Ferster D. 1997. A tonic hyperpolarization underlying contrast adaptation in cat visual cortex. *Science* 276: 949-52

Carandini M, Heeger DJ. 1994. Summation and division by neurons in primate visual cortex. *Science* 264: 1333-6

Carandini M, Heeger DJ, Movshon JA. 1997. Linearity and normalization in simple cells of the macaque primary visual cortex. *J Neurosci* 17: 8621-44

Carandini M, Heeger DJ, Senn W. 2002. A synaptic explanation of suppression in visual cortex. *J Neurosci* 22: 10053-65

Carlton AG. 1969. On the bias of information estimates. *Psych. Bull.* 71: 108-9.

Chance FS, Abbott LF, Reyes AD. 2002. Gain modulation from background synaptic input. *Neuron* 35: 773-82

Chance FS, Nelson SB, Abbott LF. 1998. Synaptic depression and the temporal response characteristics of V1 cells. *J Neurosci* 18: 4785-99

Chandler JD. 1969. Subroutine STEPIT: find local minima of a smooth function of several parameters. *Behav. Sci.* 14: 81-2

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

Chichilnisky EJ, Baylor DA. 1999. Receptive-field microstructure of blue-yellow ganglion cells in primate retina. *Nat Neurosci* 2: 889-93

Chino YM, Smith EL, 3rd, Hatta S, Cheng H. 1997. Postnatal development of binocular disparity sensitivity in neurons of the primate visual cortex. *J Neurosci* 17: 296-307

Cover TM, Thomas JA. 1991. *Elements of information theory*. New York: Wiley

Daugman JG. 1985. Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. *J Opt Soc Am A* 2: 1160-9

de Ruyter van Steveninck R, Bialek W. 1988a. Real-time performance of a movement-sensitive neuron in the blowfly visual system. *Proc R Soc Lond B* 234: 269-76

de Ruyter van Steveninck RR, Bialek W. 1988b. Real-time performance of a movement-sensitive neuron in the blowfly visual system: coding and information transfer in short spike sequences. *Proc R Soc Lond B* 234: 379-414

de Ruyter van Steveninck RR, Lewen GD, Strong SP, Koberle R, Bialek W. 1997. Reproducibility and variability in neural spike trains. *Science* 275: 1805-8

De Valois RL, Cottaris NP, Mahon LE, Elfar SD, Wilson JA. 2000. Spatial and temporal receptive fields of geniculate and cortical cells and directional selectivity. *Vision Res* 40: 3685-702

Dean AF, Hess RF, Tolhurst DJ. 1981. Divisive inhibition involved in direction selectivity. *J Physiol* 308: 304-5

Dean AF, Tolhurst DJ. 1983. On the distinctness of simple and complex cells in the visual cortex of the cat. *J Physiol* 344: 305-25

DeAngelis GC, Ohzawa I, Freeman RD. 1993. Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. I. General characteristics and postnatal development. *J Neurophysiol* 69: 1091-117

Emerson RC, Bergen JR, Adelson EH. 1992. Directionally selective complex cells and the computation of motion energy in cat visual cortex. *Vision Res* 32: 203-18

Emerson RC, Citron MC, Vaughn WJ, Klein SA. 1987. Nonlinear directionally selective subunits in complex cells of cat striate cortex. *J Neurophysiol* 58: 33-65

Enroth-Cugell C, Robson JG. 1966. The contrast sensitivity of retinal ganglion cells of the cat. *J Physiol* 187: 517-22

Felleman DJ, Kaas JH. 1984. Receptive-field properties of neurons in middle temporal visual area (MT) of owl monkeys. *J Neurophysiol* 52: 488-513

Ferster D, Chung S, Wheat H. 1996. Orientation selectivity of thalamic input to simple cells of cat visual cortex. *Nature* 380: 249-52

Forte J, Peirce JW, Kraft JM, Krauskopf J, Lennie P. 2002. Residual eye-movements in macaque and their effects on visual responses of neurons. *Vis Neurosci* 19: 31-8

Foster KH, Gaska JP, Nagler M, Pollen DA. 1985. Spatial and temporal frequency selectivity of neurones in visual cortical areas V1 and V2 of the macaque monkey. *J Physiol* 365: 331-63

Gabor D. 1946. Theory of communication. *J Inst Electr Eng* 93: 429-57

Gaska JP, Jacobson LD, Chen HW, Pollen DA. 1994. Space-time spectra of complex cell filters in the macaque monkey: a comparison of results obtained with pseudowhite noise and grating stimuli. *Vis Neurosci* 11: 805-21

Gawne TJ, Richmond BJ. 1993. How independent are the messages carried by adjacent inferior temporal cortical neurons? *J Neurosci* 13: 2758-71

Gizzi MS, Katz E, Schumer RA, Movshon JA. 1990. Selectivity for orientation and direction of motion of single neurons in cat striate and extrastriate visual cortex. *J Neurophysiol* 63: 1529-43

Gonzalez F, Perez R, Justo MS, Bermudez MA. 2001. Response latencies to visual stimulation and disparity sensitivity in single cells of the awake Macaca mulatta visual cortex. *Neurosci Lett* 299: 41-4

Heeger DJ. 1992a. Half-squaring in responses of cat striate cells. *Vis Neurosci* 9: 427-43

Heeger DJ. 1992b. Normalization of cell responses in cat striate cortex. *Vis Neurosci* 9: 181-97

Heeger DJ. 1993. Modeling simple-cell direction selectivity with normalized, half-squared, linear operators. *J Neurophysiol* 70: 1885-98

Heeger DJ, Simoncelli EP, Movshon JA. 1996. Computational models of cortical visual processing. *Proc Natl Acad Sci U S A* 93: 623-7

Hertz JA, Kjaer TW, Eskander EN, Richmond BJ. 1995. Meausring natural neural processing with artificial neural networks. *Int. J. Neural Syst. 3 (suppl.)*: 91-103

Hochstein S, Shapley RM. 1976. Linear and nonlinear spatial subunits in Y cat retinal ganglion cells. *J Physiol* 262: 265-84

Hubel DH, Wiesel TN. 1962. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J Physiol* 160: 106-54

Hubel DH, Wiesel TN. 1968. Receptive fields and functional architecture of monkey striate cortex. *J Physiol* 195: 215-43

Jagadeesh B, Wheat HS, Kontsevich LL, Tyler CW, Ferster D. 1997. Direction selectivity of synaptic potentials in simple cells of the cat visual cortex. *J Neurophysiol* 78: 2772-89

Jones JP, Palmer LA. 1987a. An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. *J Neurophysiol* 58: 1233-58

Jones JP, Palmer LA. 1987b. The two-dimensional spatial structure of simple receptive fields in cat striate cortex. *J Neurophysiol* 58: 1187-211

Jones JP, Stepnoski A, Palmer LA. 1987. The two-dimensional spectral structure of simple receptive fields in cat striate cortex. *J Neurophysiol* 58: 1212-32

Kara P, Reinagel P, Reid RC. 2000. Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron* 27: 635-46

Kayser A, Priebe NJ, Miller KD. 2001. Contrast-dependent nonlinearities arise locally in a model of contrast-invariant orientation tuning. *J Neurophysiol* 85: 2130-49

Keat J, Reinagel P, Reid RC, Meister M. 2001. Predicting every spike: a model for the responses of visual neurons. *Neuron* 30: 803-17

Kiorpes L, Movshon JA. 2003. Neural limitations on visual development in primates. In *The visual neurosciences*, ed. LM Chalupa, JS Werner. Cambridge, MA: MIT press

Krukowski AE, Miller KD. 2001. Thalamocortical NMDA conductances and intracortical inhibition can explain cortical temporal tuning. *Nat Neurosci* 4: 424-30

Lau B, Stanley GB, Dan Y. 2002. Computational subunits of visual cortical neurons revealed by artificial neural networks. *Proc Natl Acad Sci U S A* 99: 8974-9

Livingstone MS. 1998. Mechanisms of direction selectivity in macaque V1. *Neuron* 20: 509-26

Livingstone MS, Conway BR. 2003. Substructure of direction-selective receptive fields in macaque V1. *J Neurophysiol* 89: 2743-59

Lund JS, Lund RD, Hendrickson AE, Bunt AH, Fuchs AF. 1976. The origin of efferent pathways from the primary visual cortex, area 17, of the macaque monkey as shown by retrograde transport of horseradish peroxidase. *J Comp Neurol* 164: 287-304

Maex R, Orban GA. 1996. Model circuit of spiking neurons generating directional selectivity in simple cells. *J Neurophysiol* 75: 1515-45

Mainen ZF, Sejnowski TJ. 1995. Reliability of spike timing in neocortical neurons. *Science* 268: 1503-6

Marcelja S. 1980. Mathematical description of the responses of simple cortical cells. *J Opt Soc Am* 70: 1297-300

Marmarelis PZ, Naka K. 1972. White-noise analysis of a neuron chain: an application of the Wiener theory. *Science* 175: 1276-8

Maunsell JH, Gibson JR. 1992. Visual response latencies in striate cortex of the macaque monkey. *J Neurophysiol* 68: 1332-44

Maunsell JH, Van Essen DC. 1983. Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation. *J Neurophysiol* 49: 1127-47

McLean J, Palmer LA. 1989. Contribution of linear spatiotemporal receptive field structure to velocity selectivity of simple cells in area 17 of cat. *Vision Res* 29: 675-9

Miller KD, Troyer TW. 2002. Neural noise can explain expansive, power-law nonlinearities in neural response functions. *J Neurophysiol* 87: 653-9

Morrone MC, Burr DC, Maffei L. 1982. Functional implications of cross-orientation inhibition of cortical visual cells. I. Neurophysiological evidence. *Proc R Soc Lond B Biol Sci* 216: 335-54

Movshon JA, Adelson EH, Gizzi MS, Newsome WT. 1985. The analysis of moving visual patterns. *Pontificiae Academiae Scientarum Scripta Varia* 54: 117-51

Movshon JA, Newsome WT. 1996. Visual response properties of striate cortical neurons projecting to area MT in macaque monkeys. *J Neurosci* 16: 7733-41

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

Movshon JA, Thompson ID, Tolhurst DJ. 1978b. Spatial summation in the receptive fields of simple cells in the cat's striate cortex. *J Physiol* 283: 53-77

Murthy A, Humphrey AL. 1999. Inhibitory contributions to spatiotemporal receptive-field structure and direction selectivity in simple cells of cat area 17. *J Neurophysiol* 81: 1212-24

Murthy A, Humphrey AL, Saul AB, Feidler JC. 1998. Laminar differences in the spatiotemporal structure of simple cell receptive fields in cat area 17. *Vis Neurosci* 15: 239-56

Ohzawa I, Sclar G, Freeman RD. 1982. Contrast gain control in the cat visual cortex. *Nature* 298: 266-8

Paninski L. 2003. Convergence properties of three spike-triggered analysis techniques. *Network* 14: 437-64

Panzeri S, Treves A. 1996. Analytical estimates of limited sampling biases in different information measures. *Network* 7: 87-107

Parker AJ, Newsome WT. 1998. Sense and the single neuron: probing the physiology of perception. *Annu Rev Neurosci* 21: 227-77

Pillow JW, Paninski L, Simoncelli EP. 2004. Maximum likelihood estimation of a stochastic integrate-and-fire neural model. *Advances in neural information processing systems* 16: to appear

Prince SJ, Pointon AD, Cumming BG, Parker AJ. 2000. The precision of single neuron responses in cortical area V1 during stereoscopic depth judgments. *J Neurosci* 20: 3387-400

Prince SJ, Pointon AD, Cumming BG, Parker AJ. 2002. Quantitative analysis of the responses of V1 neurons to horizontal disparity in dynamic random-dot stereograms. *J Neurophysiol* 87: 191-208

Qian N, Andersen RA. 1994. Transparent motion perception as detection of unbalanced motion signals. II. Physiology. *J Neurosci* 14: 7367-80

Qian N, Andersen RA. 1995. V1 responses to transparent and nontransparent motions. *Exp Brain Res* 103: 41-50

Reid RC, Soodak RE, Shapley RM. 1987. Linear mechanisms of directional selectivity in simple cells of cat striate cortex. *Proc Natl Acad Sci U S A* 84: 8740-4

Rieke F, Bodnar DA, Bialek W. 1995. Naturalistic stimuli increase the rate and efficiency of information transmission by primary auditory afferents. *Proc R Soc Lond B Biol Sci* 262: 259-65

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

Rodieck RW, Stone J. 1965a. Analysis of receptive fields of cat retinal ganglion cells. *J Neurophysiol* 28: 833-49

Rodieck RW, Stone J. 1965b. Response of cat retinal ganglion cells to moving visual patterns. *J Neurophysiol* 28: 819-32

Rodman HR, Albright TD. 1987. Coding of visual stimulus velocity in area MT of the macaque. *Vision Res* 27: 2035-48

Rolls ET, Treves A, Tovee MJ. 1997a. The representational capacity of the distributed encoding of information provided by populations of neurons in primate temporal visual cortex. *Exp Brain Res* 114: 149-62

Rolls ET, Treves A, Tovee MJ, Panzeri S. 1997b. Information in the neuronal representation of individual stimuli in the primate temporal visual cortex. *J Comput Neurosci* 4: 309-33

Rust NC, Schultz SR, and Movshon JA (2002) A reciprocal relationship between reliability and responsiveness in developing visual cortical neurons. *J Neurosci* 22:10519:10523

Rust NC, Schwartz O, Movshon JA, Simoncelli EP (2004) Spike-triggered characterization of excitatory and suppressive stimulus dimensions in monkey V1 *Neurocomputing* 58-60: 793-799

Sakai HM. 1992. White-noise analysis in neurophysiology. *Physiol Rev* 72: 491-505

Sanes DH. 1993. The development of synaptic function and integration in the central auditory system. *J Neurosci* 13: 2627-37

Sato H, Katsuyama N, Tamura H, Hata Y, Tsumoto T. 1995. Mechanisms underlying direction selectivity of neurons in the primary visual cortex of the macaque. *J Neurophysiol* 74: 1382-94

Saul AB, Humphrey AL. 1990. Spatial and temporal response properties of lagged and nonlagged cells in cat lateral geniculate nucleus. *J Neurophysiol* 64: 206-24

Saul AB, Humphrey AL. 1992. Temporal-frequency tuning of direction selectivity in cat visual cortex. *Vis Neurosci* 8: 365-72

Schwartz O, Chichilnisky EJ, Simoncelli EP. 2002. Characterizing gain control using spike-triggered covariance. *Advances in neural information processing systems* 14: 269-76

Scobey RP, Gabor AJ. 1989. Orientation discrimination sensitivity of single units in cat primary visual cortex. *Exp Brain Res* 77: 398-406

Shadlen MN, Britten KH, Newsome WT, Movshon JA. 1996. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J Neurosci* 16: 1486-510

Shadlen MN, Newsome WT. 1998. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci* 18: 3870-96

Shannon CE. 1948. The mathematical theory of communication. *Bell Syst Tech J* 27: 379-423

Shapley RM, Victor JD. 1978. The effect of contrast on the transfer properties of cat retinal ganglion cells. *J Physiol* 285: 275-98

Shipp S, Zeki S. 1989. The Organization of Connections between Areas V5 and V1 in Macaque Monkey Visual Cortex. *Eur J Neurosci* 1: 309-32

Sillito AM. 1975. The contribution of inhibitory mechanisms to the receptive field properties of neurones in the striate cortex of the cat. *J Physiol* 250: 305-29

Sillito AM, Kemp JA, Milson JA, Berardi N. 1980. A re-evaluation of the mechanisms underlying simple cell orientation selectivity. *Brain Res* 194: 517-20

Sillito AM, Salt TE, Kemp JA. 1985. Modulatory and inhibitory processes in the visual cortex. *Vision Res* 25: 375-81

Sillito AM, Versiani V. 1977. The contribution of excitatory and inhibitory inputs to the length preference of hypercomplex cells in layers II and III of the cat's striate cortex. *J Physiol* 273: 775-90

Simoncelli EP, Heeger DJ. 1998. A model of neuronal responses in visual area MT. *Vision Res* 38: 743-61

Simoncelli EP, Pillow JW, Paninski L, Schwartz O. 2004. Characterization of neural response with stochastic stimuli. *in The Cognitive Neurosciences, 3rd Edition Ed: M Gazzaniga MIT Press*: to appear

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: 1079-86

Snowden RJ, Treue S, Andersen RA. 1992. The response of neurons in areas V1 and MT of the alert rhesus monkey to moving random dot patterns. *Exp Brain Res* 88: 389-400

Softky WR, Koch C. 1993. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci* 13: 334-50

Spatz WB. 1977. Topographically organized reciprocal connections between areas 17 and MT (visual area of superior temporal sulcus) in the marmoset Callithrix jacchus. *Exp Brain Res* 27: 559-72

Suarez H, Koch C, Douglas R. 1995. Modeling direction selectivity of simple cells in striate visual cortex within the framework of the canonical microcircuit. *J Neurosci* 15: 6700-19

Szulborski RG, Palmer LA. 1990. The two-dimensional spatial structure of nonlinear subunits in the receptive fields of complex cells. *Vision Res* 30: 249-54

Tadmor Y, Tolhurst DJ. 1989. The effect of threshold on the relationship between the receptive-field profile and the spatial-frequency tuning curve in simple cells of the cat's striate cortex. *Vis Neurosci* 3: 445-54

Theunissen F, Roddey JC, Stufflebeam S, Clague H, Miller JP. 1996. Information theoretic analysis of dynamical encoding by four identified primary sensory interneurons in the cricket cercal system. *J Neurophysiol* 75: 1345-64

Tigges J, Tigges M, Anschel S, Cross NA, Letbetter WD, McBride RL. 1981. Areal and laminar distribution of neurons interconnecting the central visual cortical areas 17, 18, 19, and MT in squirrel monkey (Saimiri). *J Comp Neurol* 202: 539-60

Tolhurst DJ. 1989. The amount of information transmitted about contrast by neurones in the cat's visual cortex. *Vis Neurosci* 2: 409-13

Tolhurst DJ, Dean AF. 1991. Evaluation of a linear model of directional selectivity in simple cells of the cat's striate cortex. *Vis Neurosci* 6: 421-8

Tolhurst DJ, Movshon JA, Dean AF. 1983. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res* 23: 775-85

Tolhurst DJ, Movshon JA, Thompson ID. 1981. The dependence of response amplitude and variance of cat visual cortical neurones on stimulus contrast. *Exp Brain Res* 41: 414-9

Touryan J, Lau B, Dan Y. 2002. Isolation of relevant visual features from random stimuli for cortical complex cells. *J Neurosci* 22: 10811-8

Van Essen DC, Maunsell JH, Bixby JL. 1981. The middle temporal visual area in the macaque: myeloarchitecture, connections, functional properties and topographic organization. *J Comp Neurol* 199: 293-326

van Santen JP, Sperling G. 1984. Temporal covariance model of human motion perception. *J Opt Soc Am A* 1: 451-73

van Santen JP, Sperling G. 1985. Elaborated Reichardt detectors. *J Opt Soc Am A* 2: 300-21

Victor JD, Shapley RM. 1979. Receptive field mechanisms of cat X and Y retinal ganglion cells. *J Gen Physiol* 74: 275-98

Vogels R, Spileers W, Orban GA. 1989. The response variability of striate cortical neurons in the behaving monkey. *Exp Brain Res* 77: 432-6

Warland DK, Reinagel P, Meister M. 1997. Decoding visual information from a population of retinal ganglion cells. *J Neurophysiol* 78: 2336-50

Watson AB, Ahumada AJ, Jr. 1985. Model of human visual-motion sensing. *J Opt Soc Am A* 2: 322-41

Werner G, Mountcastle VB. 1965. Neural Activity in Mechanoreceptive Cutaneous Afferents: Stimulus-Response Relations, Weber Functions, and Information Transmission. *J Neurophysiol* 28: 359-97

White EL. 1989. *Cortical circuits: synaptic organization of the cerebral cortex*. Boston: Birkhauser

Wiener N. 1958. *Nonlinear problems in random theory*. Cambridge, MA: MIT Press

Wiesel TN, Hubel DH. 1974. Ordered arrangement of orientation columns in monkeys lacking visual experience. *J Comp Neurol* 158: 307-18

Yabuta NH, Sawatari A, Callaway EM. 2001. Two functional channels from primary visual cortex to dorsal visual cortical areas. *Science* 292: 297-300