Neural Mechanisms of Orientation Selectivity in the Visual Cortex

David Ferster
Department of Neurobiology and Physiology
Northwestern University
Evanston, IL 60208
email: ferster@nwu.edu

Kenneth D. Miller
Departments of Physiology and Otolaryngology,
Keck Center for Integrative Neuroscience,
Sloan Center for Theoretical Biology
University of California, San Francisco 94143-0444
email: ken@phy.ucsf.edu

This is a preprint (final draft) of an article that appeared as
Annual Reviews of Neuroscience, Vol. 23, pp. 441-471 (2000).

ABSTRACT

The origin of orientation selectivity in the responses of simple cells in cat visual cortex serves as a model problem for understanding cortical circuitry and computation. The feedforward model of Hubel and Wiesel posits that this selectivity arises simply from the arrangement of thalamic inputs to a simple cell. Much evidence, including a number of recent intracellular studies, supports a primary role of the thalamic inputs in determining simple cell response properties including orientation tuning. However, this mechanism alone cannot explain the invariance of orientation tuning to changes in stimulus contrast. Simple cells receive push-pull inhibition: ON inhibition in OFF subregions and vice versa. Addition of such inhibition to the feedforward model can account for this contrast invariance, provided the inhibition is sufficiently strong. The predictions of "normalization" and "feedback" models are reviewed and compared to the predictions of this modified feedforward model and to experimental results. The modified feedforward and the feedback models ascribe fundamentally different functions to cortical processing.

INTRODUCTION

No other receptive field property characterizes the neurons of the visual cortex like orientation selectivity. The great majority of neurons in the primary visual cortex of many carnivores and primates are exquisitely sensitive to the orientation of a stimulus. Yet the relay cells of the lateral geniculate nucleus (LGN), which provide the cortex with most of its information about the visual image, respond equally well to a stimulus at any orientation. In at least some species, including cats, this remarkable and quintessentially cortical property emerges fully formed at a single synapse, between thalamic axons and their targets in the cortical layers.

Because it is such a striking phenomenon, because it is relatively easy to measure, and because it is so strongly linked with the function of the visual cortex, orientation selectivity and the mechanisms that give rise to it have been subjected to intense study and debate. Much of the cerebral cortex performs tasks that are dauntingly complex, difficult to characterize, and only just becoming experimentally approachable. Although a complex spatial transformation, extracting the orientation of an image element is still relatively straightforward and tractable. No wonder orientation selectivity has become one of the standard models for how the synaptic circuitry of the cortex performs a complex computation.

The roots of the longstanding controversy over the synaptic mechanisms underlying orientation selectivity lie in the complexity of the cortical circuit. It is easy to say that orientation selectivity in cats emerges at a single synapse between the terminals of geniculate relay cell axons and the cortical cells that they excite. But these same cortical cells receive thousands of synapses altogether, and from many different sources. Determining which of these broad categories of inputs — thalamic excitatory, intracortical excitatory, intracortical inhibitory, or some combination of all three — gives rise to orientation selectivity and how they do so has proven to be a surprisingly difficult task. At issue is not just which of the various pathways contribute, but the entire nature of the cortical computation, whether orientation selectivity arises from a feedforward filtering of the thalamic inputs to the cortex, or from a more dynamic, feedback process that encompasses the entire cortical circuit.

We shall focus our discussion on cat V1, for two reasons. First, the vast majority of cells in layer 4, the cortical layer that receives the dominant LGN input, are orientation selective in cat, though the same is not true in many other species *e.g.* monkeys (Blasdel and Fitzpatrick 1984, Hawken and Parker 1984), ferret (Chapman and Stryker 1993), and tree shrew (Humphrey and Norton 1980). Second, the synaptic physiology underlying orientation selectivity is by far best studied in cats.

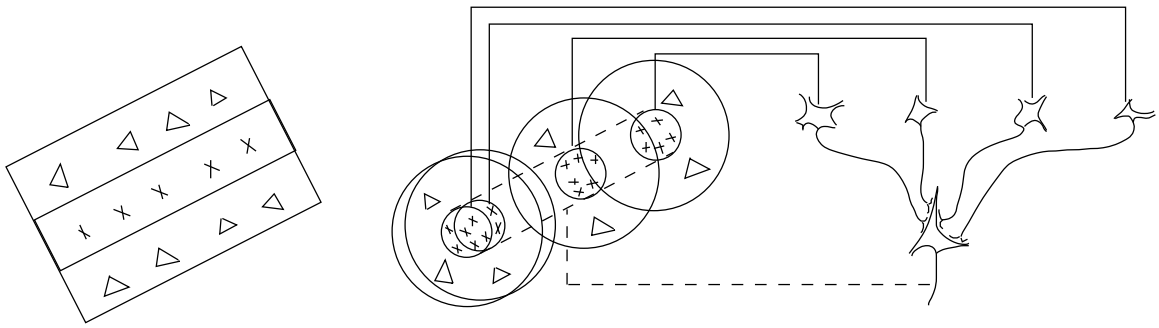

Figure 1:
A. A map of the receptive field of a simple cell in the cat visual cortex. A light flashed in the ON subregion (x) or turned off in an OFF region (triangles) excites the cell, while a light flashed in an OFF region or turned off in the ON region inhibit the cell. Other arrangements of the subregions are possible, such as a central OFF region and flanking ON regions, or one ON and one OFF region. B. Hubel and Weisel's model for how the receptive field of the simple cell can be built from excitatory input from geniculate relay cells. The simple cell (below right) receives input from relay cells (above right) whose receptive field centers are superimposed on the simple cell's central ON region. Not shown are OFF relay cells whose receptive field centers would superimpose on the simple cell's OFF regions.

THE FEEDFORWARD MODEL

When Hubel and Wiesel (1962) first described orientation selectivity in the neurons of the cat visual cortex, they proposed an elegantly direct model that remains at the center of the debate. Their model represents the feedforward model in its simplest form, explaining orientation selectivity solely from the organization of the thalamic input to a simple cell in cortical layer 4.

Simple cells in the cat are defined by the elongated ON and OFF subfields into which their receptive fields can be divided. These subfields are arranged side-by-side, with their long axes parallel to the axis of the preferred orientation of the cell. They are strongly reminiscent, in their width and sensitivity to light, of the ON and OFF centers of the receptive fields of geniculate relay cells. Hubel and Wiesel proposed that they were derived directly from thalamic input. According to their scheme, a cortical ON region arises from the excitatory input from several ON-center relay cells whose receptive field centers lie along the axis of the subfield (Figure 1). Similarly, an OFF region would be derived from the input from several OFF-center neurons.

Orientation selectivity emerges automatically from this simple arrangement. A bar of light at the orientation of an ON subfield that is moved or flashed within the subfield will simultaneously activate all of the presynaptic geniculate ON-center cells. The resulting barrage of synaptic excitation will depolarize the cortical cell and cause it to fire spikes. In

contrast, a bar moved or flashed at right angles to the subfield will only activate a small subset of the underlying geniculate relay cells at one time. The resulting depolarization of the simple cell would be too small to reach threshold, leaving the simple cell inactive. The essence of the feedforward model, then, is a linear summation stage, in which the input from the presynaptic geniculate neurons is summed on the membrane of the simple cell, followed by a non-linear rectification stage, in which the action potential threshold filters out the small synaptic inputs that are evoked by improperly oriented stimuli.

EXPERIMENTAL SUPPORT FOR THE FEEDFORWARD MODEL

Receptive Field Organization and Orientation Tuning Width

If the excitatory input from geniculate relay cells is the dominant input to simple cells and defines their subfields, then many of the response properties of simple cells should resemble those of the relay cells. That both relay cell centers and simple cell subfields come in ON and OFF varieties is one of the resemblances that led Hubel and Wiesel to propose their model, but it is only the most rudimentary resemblance between the two receptive field types. In addition, the widths of simple cell subfields are comparable to those of relay cell receptive field centers at similar visual field eccentricities (Bullier et al. 1982, Mullikin et al. 1984, Reid and Alonso 1995). The resemblance extends to more subtle measures as well, including the dynamics of the responses to flashing bars (Mullikin et al. 1984) and sinusoidally modulated bars (Saul and Humphrey 1992), and to the linearity of spatial summation as measured with sinusoidal gratings (Ferster and Jagadeesh 1991). For each of these measures, simple cell responses fall naturally into the same categories as relay cells, including X and Y, or lagged and non-lagged.

The feedforward model makes an important prediction regarding the relationship between the degree of orientation selectivity and the aspect ratio (length-to-width ratio) of a simple cell's subfields. The longer the subfield in relation to its width, the greater will be the difference in the magnitude of the geniculate excitation evoked by an optimally oriented stimulus and one at right angles. Furthermore, in cells with long narrow subfields, a relatively small shift in stimulus orientation will move a large proportion of the stimulus out of the subfield. Accordingly, the longer a simple cell's subfields are in relation to their width, the more sharply orientation tuned the cell should be. To test these predictions of the feedforward model, (Jones and Palmer 1987b,a) made high resolution maps of the subfields of simple cells, applying their reverse correlation technique to the responses to small dots flashed briefly throughout the receptive field. The average measured subfield aspect ratio was approximately 5, with values as high as 12 in some cells. Similar values were obtained by Gardner et al. (1999) (but see Pei et al. 1994). More critically, Jones and Palmer and Gardner

et al measured orientation tuning in the cells whose receptive fields they had mapped and found a strong relationship between orientation tuning width and receptive field shape: As predicted, the higher the aspect ratio of the subfield, the sharper the orientation tuning. The tuning, however, was sharper than that predicted by the simple feedforward model, a finding that may be related to the effects of the spike threshold (see below).

The feedforward model also predicts that orientation tuning width, when measured with gratings,[1] should decrease with increasing stimulus spatial frequency (see Troyer et al. 1998). Several studies have found this to be the case (Vidyasagar and Sigüenza 1985, Webster and De Valois 1985, Jones et al. 1987, Hammond and Pomfrett 1990).

Synaptic Connections between Geniculate Relay Cells and Cortical Simple Cells

Perhaps the most fundamental prediction of the feedforward model is that simple cells should receive strong excitatory synaptic input from geniculate relay cells. It was found early on in the study of area 17 of the cat that simple cells lie predominately in layers 4 and 6 (Hubel and Wiesel 1962, Kelly and van Essen 1974, Gilbert 1977, Shatz and Stryker 1978, Bullier and Henry 1979), the same layers in which the relay cells terminate (Rosenquist et al. 1975, LeVay and Gilbert 1976). Electrical stimulation of the optic radiations or LGN, combined with intracellular measurements of the latency of the evoked EPSPs confirmed that most simple cells in layer 4 indeed received substantial monosynaptic excitation from relay cells (Bullier and Henry 1979, Ferster and Lindström 1983, Martin and Whitteridge 1984). Cross correlation analysis on simultaneously recorded simple cells and geniculate relay cells also point to the presence of a direct connection between relay cells and simple cells (Tanaka 1983, Reid and Alonso 1995).

The Spatial Organization of Geniculate Input to Simple Cells

One of the most specific predictions of the the feedforward model is that the geniculate input has a definite spatial organization. Is the LGN input segregated into ON and OFF regions that correspond to the simple cell's visually defined ON and OFF subfields, as predicted by Hubel and Wiesel (1962)? An affirmative answer was given by the cross-correlation experiments of Tanaka (1983) and Reid and Alonso (1995). These authors found that functional connections between a simultaneously recorded LGN relay cell and simple cell were present when the center of the relay cell receptive field overlapped a simple cell subfield of similar response polarity. That is, a relay cell and the simple cell were connected if a relay cell's ON center overlapped an ON subfield of the simple cell, or an OFF center overlapped an OFF

[1] Here, a grating is a series of parallel bars drifting across the receptive field. The luminance of the grating varies sinusoidally in the direction perpendicular to the bars, and spatial frequency refers to the number of bars per degree of visual angle, which is inversely proportional to the bar width.

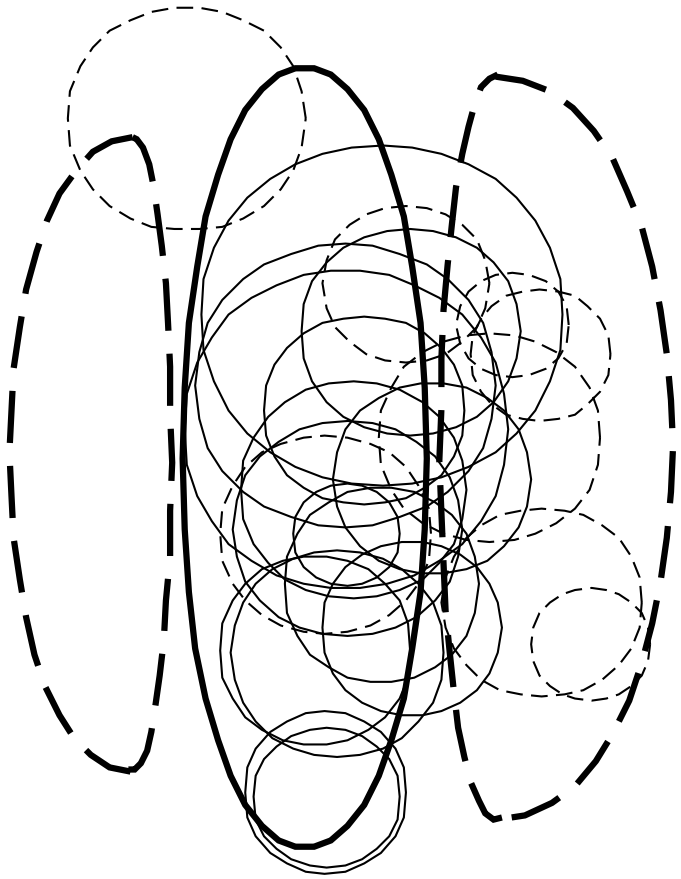

Figure 2:
Summary data from Reid and Alonso (1995). The circles represent receptive fields centers of geniculate relay cells. Each relay cells was recorded simultaneously with a cortical simple cell, and in each case, the two were found to fire in a highly correlated manner that indicated a monosynaptic connection from the relay cell to the simple cell. The receptive field center of each relay cell is plotted on a single idealized simple cell receptive field (thick lines) to indicate its position relative to the receptive field of the simple cell to which it was connected. Solid lines correspond to the strongest of the subfields; dashed lines to the weaker subfields. In almost every case, the receptive field center of the connected relay cell overlapped the subfield of the same polarity, as indicated by the dashed and solid circles.

subfield (Figure 2). Conversely, ON (or OFF) center relay cells rarely connected to a simple cell with an overlapping OFF (or ON) region, and the strength of the synaptic connection was correlated with the degree of overlap of the receptive fields (Reid and Alonso 1995). Thus, the geniculocortical projection to each cat simple cell is wired with extreme precision in just the manner required to support the feedforward model.

Support for a role of the spatial arrangement of LGN inputs in generating orientation selectivity was also provided by Chapman et al. (1991). They found in ferrets that, after

silencing cortical cells with topical application of muscimol, they could record from the LGN relay cell axons terminating within the corresponding region of layer 4. Surprisingly, the receptive fields of LGN axons recorded within a given vertical penetration formed a region in visual space that was elongated parallel to the preferred orientation of cortical cells recorded in the same column prior to muscimol application. Given that only 40% of layer 4 neurons are orientation selective in the ferret (Chapman and Stryker 1993), however, it remains to be determined whether, on a cell-by-cell basis, there is a consistent relationship between the distribution of the receptive fields of the presynaptic afferents and the orientation preference of the cell.

The Relative Strength of Geniculate Input to Simple Cells

In the debate over the relevance of the feedforward model to the origin of orientation selectivity, evaluating the strength of the geniculate excitation to simple cells has become as critical as understanding its spatial organization. Whereas the feedforward model relies on the geniculate projection to provide the predominant excitatory input to simple cells, alternative models, particularly excitatory feedback models of orientation selectivity (see below), assume that the geniculate input is relatively weak and poorly tuned for stimulus orientation compared to inputs from other cortical cells. Anatomical estimates vary widely for the proportion of the total excitatory input in layer 4 contributed by geniculate terminals. Peters and Payne (1993) calculated the proportion to be 5%, judging from the estimates of the density of cells found there, the total number of synapses per cell, the total number of geniculate neurons, and the number of boutons formed by each geniculate arbor. Ahmed et al. (1994) identified geniculate boutons in the electron microscope by their size, after labeling a small sample of them and finding that they were far larger than those originating from other sources. From their counts, they estimated that geniculocortical inputs made up 6% of the population in layer 4. LeVay and Gilbert (1976) and Einstein et al. (1987) counted the proportion of excitatory terminals in layer 4 that were labeled autoradiographically after radioactive tracer injections into the LGN. These counts of directly identified afferent terminals yield the highest numbers of all, between 22 and 26%.

More relevant to the debate than anatomical measures, however, are physiological measures of the relative strength of the geniculate input. Some types of synapses might generate synaptic drive disproportionate to their number given that their neurons of origin might fire at greater rates than other types, that they may simply release more transmitter, have more or larger postsynaptic channels or be located closer to the soma. In particular, thalamocortical synapses are large and specialized, and hence likely to be particularly effective physiologically (Ahmed et al. 1994). Gil et al. (1999) recently found, in studies of thalamocortical slices from the rat somatosensory system, that thalamocortical synapses are about

5 times stronger physiologically than intracortical synapses within layer 4. Similar results were found in studies of cat visual cortical slices in which putative geniculocortical synapses were identified in response to white matter stimulation (Stratford et al. 1996).

The first in vivo physiological estimates of the strength of the geniculate input came from cross correlation studies. Tanaka (1983) and Reid and Alonso (1995) found that a single geniculate afferent could account for between 1% and 20% of the spikes in a simple cell, judging from the proportion of simple cell spikes that were preceded at monosynaptic latencies by a spike in the relay cell. Given that each simple cell likely receives input from multiple relay cells, the total input from all of the presynaptic relay cells could be stronger still, even taking into account the spike correlations among relay cells (Alonso et al. 1996).

Extracellular measurements of spike correlation are suggestive, but could be confounded by nonlinear effects of threshold or by correlated cortical inputs. Ferster et al. (1996) and Chung and Ferster (1998) measured the contribution of the geniculate input to the responses of simple cells intracellularly, recording visually-evoked membrane potential changes in simple cells while inactivating the surrounding cortical neurons. In the first study (Figure 3), the cortex was inactivated by local cooling around the site of the intracellular recording. Control experiments indicated that during cooling, cortical cells ceased to fire action potentials in response to drifting grating or bar stimuli, with the exception of a small fraction of the spikes deep within the cortex in layer 6. During cooling, therefore, the visually evoked activity recorded intracellularly in simple cells was assumed to be of geniculate origin. This remaining activity ranged in amplitude between 5% and 50% of the normal response to the same stimuli. After correcting for the direct effects of cooling on the amplitude of synaptic potentials originating in the geniculate terminals, these authors estimated that the geniculate input was responsible for generating approximately 35% of the visual responses. The remaining 65% of the responses must therefore originate from intracortical sources, which, given the nature of the measurements could include both excitatory and inhibitory inputs.[2]

In a second intracellular study of geniculate input (Chung and Ferster 1998), the cortex was inactivated by electrical stimulation. A single shock to the upper layers of the cortex was found to suppress the response of cortical cells throughout the layers to a briefly flashed, optimally oriented grating. In an important control experiment, the shock also completely suppressed the flash-evoked EPSPs in cortical cells whose input from the LGN was mediated entirely by other cortical cells, a further indication that almost all cortical cells had failed to respond to the visual stimulus. Recordings were then made from simple cells that were known, by their short latency responses to electrical stimulation of the LGN, to receive direct

[2]The measure of visual response used was the amplitude of the temporal modulation of intracellular voltage about its mean in response to drifting sinusoidal gratings of the preferred orientation. Inhibition could contribute to the amplitude of the modulation by increasing the size of its negative-going portion.

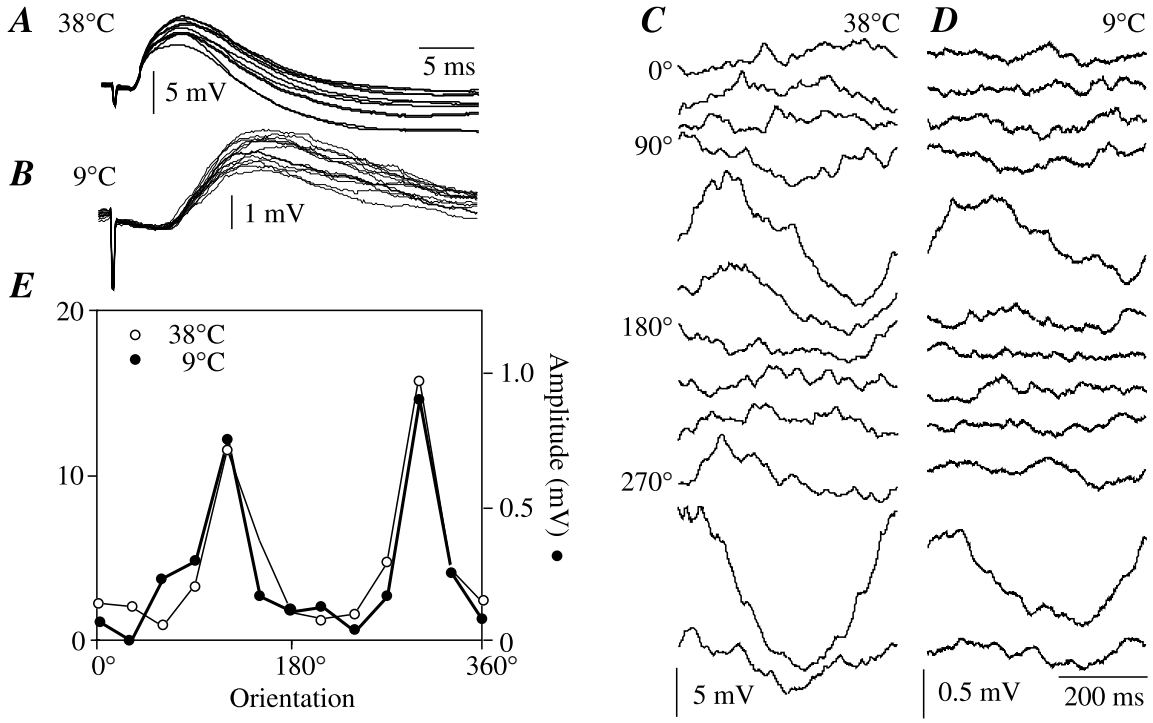

Figure 3:

The effects of cooling on the visually- and electrically-evoked responses of a layer 4 simple cell. A: Response of a simple cell in layer 4 to electrical stimulation of the LGN (1 mA, 200 μs) recorded with the cortex at normal temperature. The short (1.8 ms) and stable latency indicates that this cell received monosynaptic excitation from the LGN. The monosynaptic EPSP was followed by a long-lasting (150 ms) IPSP of disynaptic origin. B: When the cooling plate temperature was lowered to 9^{o}C, the latency of the EPSP increased to 5 ms, its rise time was dramatically slowed, and its amplitude reduced. The IPSP disappeared. C: The cell's response to drifting gratings of 12 different orientations. Each trace is the averaged response to 20 grating cycles. D: Responses to the same visual stimuli as in C, but with the cortex cooled. The responses are similar in shape to those in C, but over 17 times smaller in amplitude (note different vertical scales). E: Orientation tuning curves constructed from the records of C and D. Each point indicates the peak-to-peak amplitude of the first harmonic (2 Hz) component of the corresponding trace. Here again, note the different vertical scales for the two plots.

input from geniculate relay cells . Silencing the cortical circuit by local stimulation reduced the amplitude of the visually evoked EPSPs in these cells to an average of 46% of their normal size. The cooling experiment and the shock-inactivation experiment are therefore in approximate agreement, suggesting that roughly 1/3 to 1/2 of the excitatory input evoked in cortical cells originates in the geniculate, with the remaining inputs arising from nearby cortical cells.

The Orientation Tuning of the Geniculate Input to Simple Cells

Given the measured aspect ratio of the simple cell subfields (Jones and Palmer 1987b,a) and the cross correlation studies suggesting that the subfields arise from geniculate input (Tanaka 1983, Reid and Alonso 1995), it seemed likely that the relay cell input to a simple cell would by itself show significant orientation selectivity. To test this conjecture, in both of the cortical inactivation experiments (Ferster et al. 1996, Chung and Ferster 1998), the orientation selectivity of the membrane potential fluctuations evoked by visual stimulation was measured before and during cortical inactivation. Neither method of cortical inactivation significantly altered the width of orientation tuning of the visually evoked EPSPs. [Ferster et al. (1996) used drifting sinusoidal gratings as the stimulus, and only the tuning of the amplitude of temporal modulations of the voltage was measured; the mean voltage and its tuning were not measured.] Thus, the geniculate input is well tuned for orientation, as is predicted by its spatial organization (Tanaka 1983, Reid and Alonso 1995).

Note that these experiments also indicate that the orientation tuning of the inactivated intracortical input to these cells is very similar to that of the thalamic input. This point lends additional support to the feedforward model. The orientation tuning of the LGN input to a simple cell is not fixed; rather, it depends on the stimuli used. For example, gratings of increasing spatial frequency yield increasingly narrowly tuned LGN input (Troyer et al. 1998). For the LGN input and the cortically induced input to have identical orientation tuning, the cortical input must follow the tuning of the LGN inputs, supporting the primacy of the LGN input in determining orientation tuning.

Spike threshold and orientation tuning

The original feedforward model of Hubel and Wiesel (1962) implicitly consists of two processing stages. The first, the linear summation of input from relay cells whose receptive fields are arranged in rows, has been addressed by many of the experiments discussed so far. The second stage is the nonlinear filtering of the summed inputs by the spike threshold. Threshold is critical to the model in its simplest form: Even when a stimulus is at right angles to the preferred orientation, it will activate a few geniculate neurons since they themselves are

not orientation selective. Hubel and Wiesel therefore invoked threshold to prevent a simple cell from responding to these low amplitude inputs (though other mechanisms, such as inhibition, could also contribute, as discussed below). Until recently, the experiment that could directly test the effects of threshold on orientation selectivity, that is, a comparison of the orientation tuning of the membrane potential with the orientation tuning of spike responses in the same neurons, had not been performed. Doing so requires intracellular recordings that are very stable, and, more importantly, that perturb the relationship between spike frequency and membrane potential as little as possible. The advent of in vivo patch recording (Pei et al. 1991, Jagadeesh et al. 1992) has made such recordings possible, even for the small neurons of layer 4. In a recent experiment, Carandini and Ferster (2000) have shown that the orientation tuning of the spike responses is indeed significantly narrowed relative to the tuning of the synaptic inputs. These authors found that the average half-width at half height of the orientation tuning curve for membrane potential in simple cells was 65% greater than that of the spikes (38^o vs. 23^o), suggesting that threshold plays a significant role in shaping the responses of the simple cells.

FAILURES OF THE FEEDFORWARD MODEL: CONTRAST INVARIANCE OF ORIENTATION TUNING

The above experiments together make a convincing case that the basic organization of the simple cell receptive field — its subfields, its preferred orientation, and some measure of its orientation selectivity — are laid out by the spatial organization of the geniculate input. However, not all of the response properties of simple cells can be successfully predicted from this simple scheme.

Perhaps the most serious inadequacy of the feedforward model is revealed by a striking feature of simple cell responses, the contrast-invariance of orientation tuning of the responses to drifting gratings (Sclar and Freeman 1982, Skottun et al. 1987). As shown in Figure 4A, in real simple cells the width of the orientation tuning curve varies little as the contrast (strength) of the stimulus is varied. Only the height of the tuning curve increases with contrast. This behavior is difficult to explain in Hubel and Wiesel's simple feedforward model, because it is well known that the strengths of the responses of retinal ganglion cells (see Troy and Enroth-Cugell 1993) and geniculate relay cells (Cheng et al. 1995) increase with the contrast of the stimulus. The difficulty is shown in Figure 4D, where the peak excitatory input that would arise from the geniculate input in response to drifting gratings is plotted as a function of orientation for different contrasts. A key feature of the model is that, even at non-preferred orientations, the depolarization at high contrast will be non-zero and will increase with contrast.

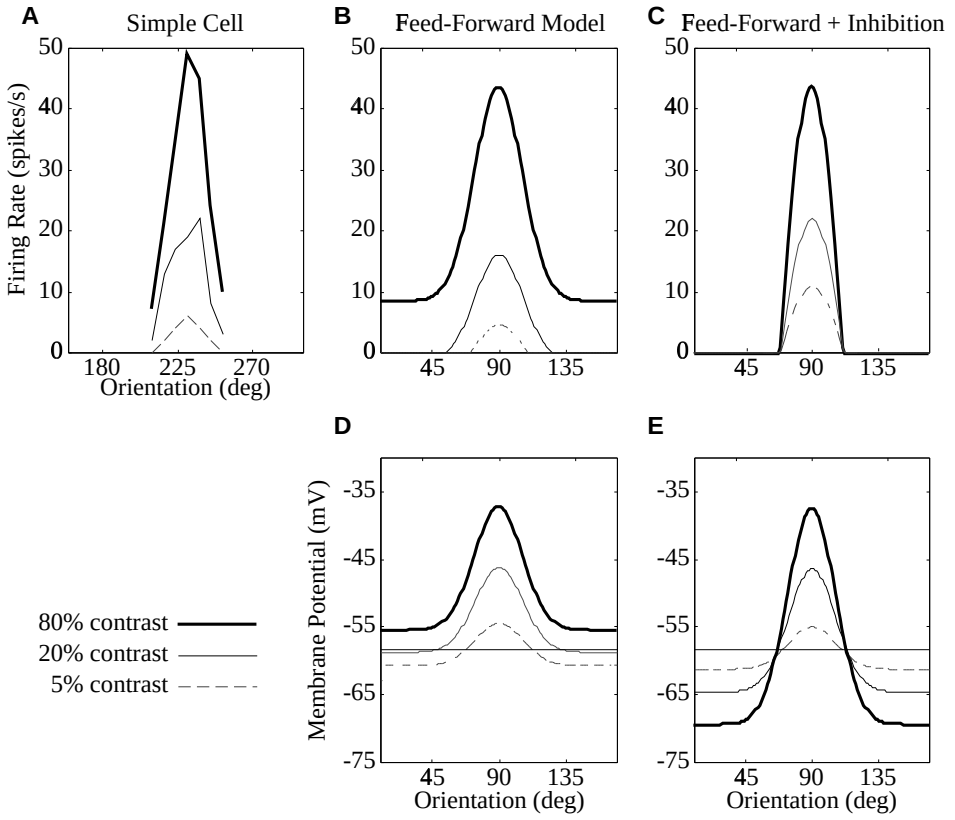

Figure 4:

Orientation tuning curves of a simple cell obtained with drifting simusoidal gratings of 3 different contrasts (adapted from Sclar and Freeman (1982). B and D. Sketch of tuning curves of the synaptic potentials (D) and spikes (B) predicted from the feedforward model of Hubel and Wiesel (1962). C and E. Sketch of tuning curves of the synaptic potentials (E) and spikes (C) predicted by the feedforward model, layered with push-pull inhibition as suggested by Troyer et al. (1998). Horizontal lines represent threshold.

To understand this behavior, consider a grating at right angles to the cell's preferred orientation. Where it crosses an ON-region of the receptive field, the bright portion of the grating will activate the underlying ON-center geniculate relay cells at rates as high as 100 spikes per second. At the same time, the bright portion of the grating will suppress OFF-center inputs where it crosses the OFF region of the receptive field. But the reduction in the responses of the suppressed OFF-center cells saturates at zero spikes per second, and since the activity starts from a relatively low spontaneous rate of 10-15 spikes per second, the reduction cannot balance out the excitation of the ON-center cells. So stimuli of the non-preferred orientation above a certain contrast evoke a net excitation from the relay cells. Furthermore, the amplitude of these non-optimal relay cell responses grows with contrast, since the responses of the relay cells themselves grow with contrast. As a result (Figure 4D), the model predicts that the net excitation evoked by non-optimal orientations at high contrasts can exceed that evoked by the preferred orientation at low contrast.

Because both the baseline level of the predicted membrane potential tuning curve and the size of its peak grow with stimulus contrast, the feedforward model breaks down when an unvarying threshold is applied to derive the spike responses of the cell (Figure 4B). No matter where threshold is placed, the orientation tuning width of the resulting firing rate responses will show considerable broadening with increasing contrast, unlike the responses of real simple cells. This is the "iceberg effect" — beneath the spiking threshold lies more broadly tuned excitatory input, which would be revealed by increasing contrast in the feedforward model. Furthermore, if the threshold is too high, stimuli at low contrasts will evoke no spike responses at all, whereas with low thresholds, stimuli at sufficiently high contrast will evoke responses at all orientations. Yet few simple cells show these behaviors (Figure 4A).

The contrast-invariance of orientation tuning has not been studied for stimuli other than drifting gratings, but similar problems exist for any stimulus. For example, consider a drifting bar. Even a low-contrast bar evokes spikes at the preferred orientation, while even a high-contrast bar does not evoke spikes at the orthogonal orientation. Yet the high-contrast bar will evoke net LGN excitation, because it will raise the firing rates of the LGN inputs it excites much more than it can lower the firing rates of the LGN inputs it suppresses. Since this net excitation increases with contrast, a sufficiently low-contrast bar of the preferred-orientation and a high-contrast orthogonally-oriented bar will yield the same peak level of LGN input, a level that must be suprathreshold given the spiking responses evoked by the preferred-orientation stimulus. Thus, the feedforward model would predict spiking responses at non-preferred orientations to high-contrast drifting bars.

In summary, if the signal underlying simple cell orientation tuning originates in feedforward LGN input, then to explain contrast-invariant tuning, some mechanism must raise the amount of LGN input required to yield a spiking response at higher stimulus contrasts.

Possible mechanisms for a contrast-dependent effective threshold

Perhaps the simplest mechanism for raising the amount of geniculate input required to evoke spikes would be to raise the spike threshold with increasing stimulus contrast. Based on the biophysical properties of cortical cells (see, for example, McCormick et al. 1985), it seems unlikely that the actual voltage threshold for spikes would change with contrast. Indeed, intracellular recordings from simple cells have shown directly that threshold voltage is invariant with contrast (Carandini and Ferster 2000).

A second potential source for a contrast-dependent change in the effectiveness of the geniculate input is the frequency-dependent depression of LGN synaptic efficacy (Markram and Tsodyks 1996, Stratford et al. 1996, Abbott et al. 1997, Gil et al. 1997). This depression increases with input firing rate, and thus would increase with contrast. While synaptic depression might partially alleviate the problem of orientation tuning widening with contrast, it is unlikely to solve the problem, for several reasons. First, the primary effect of depression is to lessen the difference between the tuning curves in Figure 4D. Depression cannot be strong enough to eliminate these differences, however, since in real simple cells spiking responses increase with increasing stimulus contrast (Figure 4A). Hence, while the problem may be alleviated, it will not be eliminated. Second, synaptic depression builds over a number of presynaptic spikes, so it would be unlikely to affect the response to transient stimuli such as a flashing bar.

A third way in which contrast could change the effectiveness of the geniculate input is for it to evoke a contrast-dependent hyperpolarization of the resting potential, thus increasing the size of the visually evoked depolarization needed to reach threshold. Contrast adaptation evokes just such a hyperpolarization (Carandini and Ferster 1997), which may be due in part to a long-lasting potassium conductance (Sanchez-Vives et al. 1997). Such adaptation is orientation-tuned, however: it is not induced by stimuli with orientation orthogonal to the preferred (Allison and Martin 1997). It also requires several seconds to develop fully (Albrecht et al. 1984, Ohzawa et al. 1985, McLean and Palmer 1996, Carandini and Ferster 1997). Thus, adaptation-induced hyperpolarization is unlikely to provide the required suppression of geniculate input.

One of the few remaining possibilities, and the one that we favor, is that the contrast-related modulation of the efficacy of the thalamic input necessary to explain contrast-invariant orientation tuning is supplied by stimulus-induced synaptic inhibition. When the feedforward model is extended to include inhibition, not only does it explain how inhibition could provide contrast invariance, but it also accounts for some of the known receptive field properties of inhibition in simple cells.

Strong push-pull inhibition can solve the problem of contrast-invariant tuning

Consider layering an inhibitory input on top of the input from relay cells, an input with a receptive field identical to that of the excitatory input, but with opposite response polarity (ON instead of OFF, and OFF instead of ON). Each ON region would then receive ON excitation and OFF inhibition, and each OFF region would receive OFF excitation and ON inhibition. With these inhibitory inputs in place, the peak response to a bar or grating of the preferred orientation would be unaffected because the excitation and inhibition would occur out of phase with one another: Whenever the excitation was peaking, the inhibition would be at its minimum, and vice versa. The response to orthogonally oriented bars, on the other hand, would be strongly affected. If the inhibition had the same contrast sensitivity and strength as the excitation, and if ON and OFF regions had equal strength, an orthogonally oriented bar would activate ON inhibition from the OFF region that exactly cancelled the ON excitation from the ON region, no matter what the contrast of the stimulus.

With excitation and inhibition of exactly equal strength, the baseline in the tuning curve of the synaptic input to a simple cell would be invariant with contrast: At non-preferred orientations, the excitation and inhibition would exactly cancel one another at all contrasts. But even if ON and OFF regions had equal strengths, this antiphase inhibition would not entirely solve the problem of contrast-invariant tuning (Troyer et al. 1998). The height of the curves would still grow in amplitude with contrast so the portion of the peak that reaches above threshold would still widen with contrast. The orientation tuning of the spike responses would then still not be completely contrast invariant.

A full solution can be found by assuming that the antiphase inhibition is sufficiently stronger than the relay-cell excitation (Troyer et al. 1998). This has the effect of actually pulling the baseline in Figure 4D down as contrast is increased and the peak grows, so that the width of the peak at the point that it crosses threshold remains constant (Figure 4E). In this way, the feedforward model with added inhibition can achieve contrast invariant orientation tuning. Furthermore, this mechanism is surprisingly insensitive to the strength of the inhibition: once the inhibition is sufficiently strong, increasing it further simply sharpens the orientation tuning, while maintaining the contrast invariance of the tuning (Troyer et al. 1998).

Simple cells receive strong push-pull inhibition

Antiphase or push-pull inhibition was first proposed by Hubel and Wiesel (1962), though they did not discuss its importance for contrast invariance. There is now ample evidence both from extracellular (Heggelund 1981, Palmer and Davis 1981) and intracellular (Ferster 1988, Borg-Graham et al. 1998, Hirsch et al. 1998, Anderson et al. 1999) recordings that simple cells

do, in fact, receive strong OFF inhibition in their ON subfields, and strong ON inhibition in their OFF subfields. This push-pull inhibition is stronger than relay-cell excitation, as required by Troyer et al. (1998)'s model. This was demonstrated in intracellular recordings by Hirsch et al. (1998), who showed that push-pull inhibition can completely suppress the response of simple cells to excitatory inputs. When a flashed spot of lightthat evoked a strong excitation from one subregion was moved slightly to encroach on a neighboring subregion of opposite polarity, the resulting push-pull inhibition overwhelmed the excitatory response, instead yielding hyperpolarization.

The overall dominance of cortical inhibition over excitation is also suggested by other experiments. Electrical stimulation of the LGN (Ferster 1986, Douglas and Martin 1991) or of the cortex (Hirsch and Gilbert 1991, Chung and Ferster 1998) evokes a brief excitation followed by a long lasting and often stronger inhibition. The same is often true of briefly flashed visual stimuli (Hirsch et al. 1998). Measurements of excitatory and inhibitory conductances evoked by drifting gratings show that the latter can be 2-5 times as large as the former (Anderson et al. 1999).

Push-pull inhibition is by far the dominant if not the sole form of inhibition received by simple cells. Any inhibition that is not in a push-pull arrangement, e.g. ON inhibition in an ON subfield, must be far weaker. To see this, note that a light spot flashed in an ON subregion evokes a strong depolarizing response. Therefore any inhibition evoked by the spot must be much weaker than the evoked excitation. Another important feature of the inhibition received by simple cells is its orientation tuning. Intracellular recordings show that this inhibition is maximal at the preferred orientation and falls off strongly away from the preferred. The orientation tuning width of the inhibition received by a cell appears to be nearly identical to that of the cell's excitatory inputs (Ferster 1986, Douglas et al. 1991, Nelson et al. 1994, Anderson et al. 1999).

What is the source of the push-pull inhibition?

There are a number of possible sources for push-pull inhibition with the properties predicted by Troyer et al. (1998)'s model. One possible source is direct input from inhibitory geniculate relay cells (Carandini and Heeger 1994). The receptive fields of these relay cells would have the same arrangement into elongated rows as that proposed by Hubel and Wiesel for the excitatory input. Were they to exist, they would fit exactly the criteria of a contrast response function similar to that of the relay cells, and complementary spatial organization to the excitatory input. Unfortunately, no physiological evidence for direct inhibition from relay cells has been found (Ferster and Lindström 1983, Martin and Whitteridge 1984), though some anatomical evidence has been reported (Einstein et al. 1987).

Thus, it is more realistic to fashion push-pull inhibition out of input from cortical in-

hibitory interneurons. An obvious problem, however, is that most cortical cells are strongly orientation selective. The inhibitory simple cells, if they were like the cortical cells most often recorded, would not normally fire in response to an orthogonally oriented stimulus, and so they would fail to counteract the non-zero baseline in the orientation tuning curves of Figure 4D.

One fix would be to arrange for each simple cell to receive inhibition from a large pool of other simple cells of all preferred orientations. At the preferred orientation of the postsynaptic cell, the subfields of the inhibitory interneurons would be aligned with those of the postsynaptic cell but have opposite response polarity, and so generate the observed push-pull arrangement. These interneurons would provide strong inhibition. At non-preferred orientations, many different interneurons with many different receptive field locations would counterbalance the relay cell input evoked by a bar of any orientation and location. [Note that this pool of inhibitory inputs would need to provide a signal that is subtracted from the feedforward geniculate input, rather than dividing it as in "normalization" models (Carandini and Heeger 1994) discussed below.] Compared to the push-pull inhibition at the preferred orientation, which is known to be quite strong, these interneurons would provide only weak inhibition at the orthogonal orientation, in accordance with the inhibition observed in simple cells. The inhibition would only need to be strong enough to overcome the relatively small relay-cell excitation predicted by the feedforward model to come from the relay cells in response to orthogonal stimuli.

A second possibility has been proposed by (Troyer et al. 1998). In their model, the interneurons providing push-pull inhibition to simple cells all have preferred orientations similar to that of the postsynaptic cell ($\pm 30^o$), but respond much like the uncompensated feedforward input itself. These interneurons, unlike the classical simple cell, would respond much like the one depicted in Figure 4B and D. For stimuli of the preferred orientation, they would act like normal simple cells, responding in the normal, position sensitive way only when bright stimuli were located in an ON region and/or dark stimuli were in an OFF region. For stimuli of the non-preferred orientation, these cells would show weak responses that grow with contrast, much like the LGN input in Figure 4B and D. The orientation tuning of these inhibitory simple cells, then, would not be invariant with contrast. They would retain the properties of the pure feedforward input.

While these inhibitory interneurons are themselves not contrast invariant in orientation, they would be in a position to generate contrast invariance in their neighbors. Consider a group of simple cells with a given preferred orientation. A high-contrast stimulus orthogonal to this preferred orientation will provide LGN input strong enough to evoke a response in simple cells unless the simple cells are sufficiently inhibited. In the model, some of the simple cells in the group — the postulated interneurons — do respond, and provide the inhibition

that prevents the remainder of the cells from firing.

A critical prediction of this model is the existence of inhibitory interneurons in layer 4 with simple receptive fields that respond to stimuli of the non-preferred orientation in a manner that increases with increasing stimulus contrast. These interneurons would respond more or less as predicted by the feed-forward model in its simplest form (Hubel and Wiesel 1962). It is difficult, however, to assess from current data whether or not these neurons exist. A large number of simple cells have been described in extracellular studies of the cortex. About 10-15% of visually responsive neurons in cat V1 are poorly tuned or non-selective for orientation (earlier work reviewed in Fregnac and Imbert (1984); Maldonado et al. (1997)). This percentage could easily include the postulated layer 4 interneurons, which are probably underrepresented in the sample. Interneurons form only 15-20% of cortical cells and, being small, are likely to have been recorded less frequently than excitatory cells. Nor is it often possible to identify interneurons from extracellular recordings. A few interneurons have been identified intracellularly, and are reported to have simple receptive fields with orientation tuning, though details of tuning were not reported (Gilbert and Wiesel 1979). A recent such study of 8 identified interneurons found, suggestively, that all respond to orientations orthogonal to their preferred (Azouz et al. 1997). Unfortunately, no interneurons have yet been studied thoroughly enough to determine whether they possess all the properties predicted by Troyer et al. (1998). A search for the proposed interneurons remains a key experiment with which to test the model.

Development of the contrast-invariant circuit

One of the questions most often asked about the feedforward model of orientation selectivity is how the underlying circuitry develops. The question becomes particularly puzzling in light of the observation that orientation selectivity develops in the absence of patterned visual input, before a kitten's eyes open (Hubel and Wiesel 1970, Crair et al. 1998, Fregnac and Imbert 1984, Movshon and Van Sluyters 1981). Miller (1994) showed that the arrangement of ON and OFF inputs to a simple cell predicted by the feedforward model can develop from activity-dependent rules of synaptic modification, provided that spontaneous LGN activity patterns in the developing animal have certain simple structures. The development of intracortical connections can be explained by these same rules (Kayser and Miller 1998, Miller et al. 1999). Excitatory cells would develop connections to correlated cells, whereas inhibitory cells would develop connections with anticorrelated cells. These rules yields the required push-pull inhibition. In addition, they yield excitation among cells that have roughly superimposed ON regions and superimposed OFF regions (same-phase excitation). The inhibition yields contrast-invariant tuning as we have seen, while the excitation amplifies the tuned responses (Ferster et al. 1996, Chung and Ferster 1998, Troyer et al. 1998).

Experimental support for the feedforward model with push-pull inhibition and same-phase excitation

The feedforward model with push-pull inhibition and same-phase excitation is consistent with a number of observations obtained from intracellular recordings of simple cells and described above. First there is the motivating observation that excitation and inhibition are spatially opponent (push-pull inhibition). Second, the orientation tuning of inhibition to simple cells in the model is identical to that of excitation: Both peak at the preferred orientation and fall off to small values at the orthogonal orientation, as observed in vivo. Third, the model is consistent with the cortical inactivation experiments (Ferster et al. 1996, Chung and Ferster 1998), which show that the orientation tuning of visually-evoked responses does not change when cortical cells are silenced, and show additionally that the cortex amplifies responses to the LGN inputs 2- to 3-fold. Fourth, the model explains why push-pull inhibition must be so strong relative to feedforward excitation. The behavior of Troyer et al. (1998)'s model conforms to that of simple cells in two other important ways. Iontophoresis of $GABA_A$ antagonists into the visual cortex have long been known to degrade or even abolish orientation selectivity (Sillito 1975, Tsumoto et al. 1979), whereas intracellular blockade of $GABA_A$ inhibition in a single cell has little effect on the cell's orientation selectivity (Nelson et al. 1994). The model closely mimics these behaviors. Because it is a variant of the feedforward model and because it relies on the spatial organization of the geniculate input for establishing orientation selectivity, Troyer et al's model also exhibits a decrease in orientation tuning width with increasing stimulus spatial frequency (Vidyasagar and Sigüenza 1985, Webster and De Valois 1985, Jones et al. 1987, Hammond and Pomfrett 1990). Again, the existence of interneurons with the proposed response properties is a key test of the model.

NORMALIZATION AND THE FEEDFORWARD MODEL

In addition to the contrast invariance of orientation selectivity, there are several other properties of simple cells that must be addressed by any model of cortical function. First, cortical responses saturate as the contrast of a stimulus increases. The responses do not, however, saturate at a fixed firing rate determined by the electrical properties of the cell, but at a rate that changes with the stimulus (Maffei et al. 1973, Dean 1981, Albrecht and Hamilton 1982). Responses to stimuli of non-optimal orientation or spatial frequency saturate at lower firing rates than do the responses to optimal stimuli. Second, as the contrast of a stimulus increases, the time course of the responses of simple cells changes. Specifically, the temporal phase (or latency) of the response to a sinusoidal grating advances in time (Dean and Tolhurst 1986, Albrecht 1995, Carandini et al. 1997). Third, responses to gratings of high temporal frequency increase more with increasing contrast than do responses to grat-

ings of low temporal frequency. As a consequence, the temporal frequency tuning of simple cells changes with contrast (Albrecht 1995). Finally, the response to a superposition of two stimuli is often less than the sum of the responses to each stimulus alone, even when one of the component stimuli evokes no response at all. A prominent example of this effect is cross-orientation inhibition: Two gratings, one of the preferred-orientation and one of the orthogonal orientation, evoke a smaller response than does the preferred-orientation grating alone (Bishop et al. 1973, Ferster 1981, Hammond and MacKay 1981, Morrone et al. 1982, Li and Creutzfeldt 1984, De Valois et al. 1985, Gulyas et al. 1987, Bonds 1989, Bauman and Bonds 1991, Nelson 1991, DeAngelis et al. 1992, Geisler and Albrecht 1992)

The normalization models were proposed (Albrecht and Geisler 1991, Heeger 1992) to explain some of these effects. In these models, the feedforward geniculate input is assumed to grow linearly with stimulus contrast, but is divided or normalized just prior to threshold by an inhibitory input whose strength also increases with contrast. The combination of the inhibition and excitation yields a sigmoidal, saturating function of contrast. Carandini and Heeger (1994) proposed that the normalization signal would take the form of a shunting inhibition driven by the pooled responses of surrounding neurons of many different preferred orientations and spatial frequencies. This shunting inhibition, which would thus be orientation independent but increase with stimulus contrast, would increase the conductance of a cell. This shunt would lower the cell's membrane time constant, thereby lowering its integration time, with the resulting effect of advancing the phase of responses to sinusoidal gratings and enhancing responses to higher temporal frequencies (Carandini and Heeger 1994, Carandini et al. 1997). Finally, a shunting inhibition at all orientations could explain cross-orientation inhibition.

While the normalization models can be made to fit the spike responses of simple cells, two key predictions have not been borne out in intracellular experiments. First, the models require a large stimulus-evoked shunting conductance that depends only on stimulus contrast, and hence is independent of stimulus orientation. It now seems clear, despite initial indications to the contrary (Douglas et al. 1988, Ferster and Jagadeesh 1992), that visual stimuli do indeed evoke large increases in the input conductance of simple cells (Borg-Graham et al. 1998, Hirsch et al. 1998, Anderson et al. 1999). The amplitudes of stimulus-induced conductance changes are strongly dependent on orientation, however, with a preferred orientation and tuning width similar to that of the membrane potential responses (Anderson et al. 1999). Second, the membrane time constant does not change sufficiently to explain contrast-dependent changes in temporal properties. Normal resting time constants are in the range of 20 ms (Hirsch et al. 1998, Anderson et al. 1999). The normalization model, however, requires decreases in time constant of 60 ms or more, which is clearly impossible since time constants cannot go below zero.

How, then, can we explain the response properties that are described by the normalization framework? It turns out these properties can be explained in the context of the feedforward model with push-pull inhibition by a variety of existing nonlinearities in the electrical properties of cells and the temporal properties of synaptic inputs (Kayser et al. 1999). These include saturation and contrast dependent phase advance in the responses of the LGN inputs themselves, frequency-dependent synaptic depression, spike-rate adaptation, non-zero threshold, and small stimulus-induced changes in membrane time constant. Since LGN inputs saturate with contrast (Sclar 1987, Cheng et al. 1995), we need only explain the difference between LGN and cortical saturating contrasts. This can be accounted for by synaptic depression (Stratford et al. 1996, Markram and Tsodyks 1996), which causes the geniculate synaptic input to saturate before presynaptic LGN firing rates saturate (Abbott et al. 1997, Tsodyks and Markram 1997). Synaptic depression, together with the other nonlinearities listed above, can also explain the difference between LGN and cortical phase advances (Chance et al. 1998) and changes in temporal frequency tuning with contrast. Synaptic depression and spike-rate adaptation act as contrast-dependent high-pass filters, allowing high-temporal-frequency responses to grow with contrast at a faster rate than low-temporal-frequency responses. Contrast-induced decrease in the membrane time constant also causes a relative enhancement of high-temporal-frequency responses with increasing contrast. Finally, cross-orientation inhibition and other two-stimulus suppression effects can be accounted for, in part, by the push-pull inhibition in the feedforward model: the non-preferred grating evokes strong push-pull inhibition, reducing responses to the preferred orientation (Krukowski et al. 1998).

It should be noted that even if the feedforward model with push-pull inhibition and same-phase-excitation were to prove largely correct, it is still at best incomplete. A number of issues need to be addressed. The model does not currently deal with direction selectivity, for example, though it is likely that the addition of lagged and non-lagged input using the scheme proposed by Saul and Humphrey (1992) would fit well with the current model. The model also does not deal with the diversity of cortical inhibitory interneurons. Studies in rat somatomotor thalamocortical slices show two classes of inhibitory neurons in layer 4 (Gibson et al. 1999): a feedforward class receiving strong thalamic input, like the inhibitory neurons of the model, and a feedback class receiving little or no thalamic input, which is not incorporated in the model. Furthermore, cells of each class show strong gap junction coupling among themselves. Intracellular studies in cat V1 layer 4 have found some inhibitory neurons that respond primarily to stimulus contrast: they are complex cells (not selective for stimulus polarity) and unselective or only weakly selective for stimulus orientation (J. Hirsch, private communication). Their role is unknown. While the simplicity of the model gives order to a wide variety of findings, the complexity of the cortical circuit should not be

underestimated.

FEEDBACK MODELS OF CORTICAL FUNCTION

So far we have focused on feedforward models in which orientation selectivity is generated by the spatial organization of the receptive fields of presynaptic geniculate relay cells. In these models, intracortical connections serve to sharpen and render contrast-invariant the orientation selectivity specified by the feedforward connections, and to amplify responses. In an alternative series of models, the cortical circuitry plays a much more central role in establishing orientation selectivity. These are the feedback models (Ben-Yishai et al. 1995, Somers et al. 1995, Hansel and Sompolinsky 1996, Ben-Yishai et al. 1997, Adorján et al. 1999), the first set of models developed to address the problem of contrast-invariance of orientation selectivity (Figure 4). (See also related models of other visual cortical phenomena: Douglas and Martin 1991, Suarez et al. 1995, Douglas et al. 1995).

In the feedback models, the geniculate input to simple cells is assumed to be relatively weak, compared to what is assumed by the feedforward model, and more importantly compared to the input from other cortical cells. The geniculate input is typically also assumed to be very poorly tuned for orientation, based on those experiments that find very small aspect ratios for the subfields of the simple cells (Pei et al. 1994). Sharp orientation tuning arises instead from excitatory interconnections among cortical cells with similar orientation preference and inhibitory interconnections among cells with more wide ranging orientation preferences. In this scheme, a plot of the strength and sign of connections between neurons against the difference in their preferred orientations forms a Mexican hat function (Somers et al. 1995, Sompolinsky and Shapley 1997): Cells with nearby preferred orientations have net excitatory connections, while cells with more disparate preferred orientations have net inhibitory connections.

A key feature of the feedback models is that the mutually excitatory, or feedback connections among cells with similar preferred orientations dramatically amplify any suprathreshold input that they receive from the LGN. A suprathreshold geniculate input triggers a small amount of activity within the feedback loop, which is enhanced by reverberations within the loop. Key to the model is that with sufficiently strong feedback, the cortex acquires its own intrinsic spatial pattern of response that isindependent of the input . The only stable response is to have a "bubble" of cortical cells exciting one another while inhibiting those cortical cells immediately outside the bubble. The geometry of the lateral excitatory and inhibitory connections controls the size of the "bubble" or region that is activated by a geniculate input. The larger the spread of excitatory connections and the stronger they are, the larger the lateral spread of activity; conversely, the stronger the lateral inhibitory connections, the smaller the lateral spread of activity. The pattern of activity that develops

in the cortex is therefore a uniquely cortical property, depending little on the pattern (orientation tuning width) of the initial, triggering input. A stimulus of higher strength (contrast) will evoke stronger activity without changing the shape of the activity bubble.

This last point means that any suprathreshold stimulus will evoke the appearance of a stereotypical pattern of activity in the cortex. The only aspects of the pattern that can change are its height (the strength of the activity), and its location on the surface of the cortex, which, by virtue of the columnar organization of the cortex, corresponds to its characteristic orientation. Both of these parameters are selected by the stimulus. The cortical circuit, in formal terms, forms a multistable attractor, in that there are many equally-favored possible states of activity, in this case all of which are identical in shape. It is this property — and this is the point of the model — that leads to the contrast invariance of the tuning width of individual cells. Tuning width is merely a function of the width of the activity pattern in the orientation domain: The wider the pattern, the farther away a stimulus can be from the preferred orientation of a cell, and still cause some activation in that cell. Since the width of the pattern of activity is a function of cortical connectivity alone and is therefore independent of the contrast (or any other attribute) of the stimulus, the orientation tuning width of the individual cells is independent of contrast and of other stimulus attributes.

It now becomes clear why the feedback models can tolerate very broad orientation tuning of the geniculate input and still maintain very sharp tuning in the cortical cells. Imagine that each simple cell receives geniculate input arranged in subfields, as originally suggested by Hubel and Wiesel (1962), but that these subfields have very small aspect ratios. As a result, the orientation tuning of the thalamic input to each simple cell will be extremely broad. A large number of simple cells will receive significant excitation when an oriented stimulus is presented. But one set of cells, whose preferred orientations are the same as the orientation of the stimulus, will receive slightly more excitation than all the rest. They will generate the strongest mutual excitation and the strongest inhibition of their neighbors. And when the cortical circuitry takes over, the stereotyped cortical activity pattern will form with fixed width, centered on the orientation of those cells that received the strongest excitation. The winner takes all. In this way, the final pattern of cortical activity can be much sharper in orientation than the geniculate input that triggered it.

Experimental support for the feedback models

The feedback models can account for many experimental observations. 1) The most prominent observation is the contrast invariance of orientation selectivity (Sclar and Freeman 1982, Skottun et al. 1987). 2) Feedback models explain an experiment in which it was found that the orientation tuning of synaptic potentials in some cortical cells sharpened over time (Pei et al. 1994), presumably as the cortical circuitry took over and sharpened the input from

the LGN. Note, however, that the feedforward model with push-pull inhibition also predicts that LGN synaptic potentials are sharpened by inhibition at high contrasts. Other time-dependent changes in orientation preferences have been seen extracellularly in monkey visual cortex, though these effects may in part involve input from beyond the classical receptive field (Ringach et al. 1997). 3) The feedback models are roughly consistent with the observed orientation tuning of excitatory and inhibitory inputs to simple cells (Ferster 1986, Douglas et al. 1991, Nelson et al. 1994, Anderson et al. 1999). Strong inhibition at the preferred orientation is required to keep the excitatory feedback under control, while only weak inhibition at the orthogonal orientation is required to counteract weak excitation. However, the feedback models predict that the inhibition received by a cell should have broader orientation tuning than the excitation it receives. Experimentally, excitation and inhibition appear to have very similar orientation tuning, although small differences cannot be ruled out. 4) Local inactivation of the cortical circuit with injections of GABA can disrupt orientation selectivity of cells hundreds of microns away (Eysel et al. 1990, Crook et al. 1995). This behavior is easily understood when it is the local cortical circuitry that determines orientation selectivity, but is less easily understood in the context of the feedforward models, where the thalamic input is dominant. 5) Because intracortical inhibitory connections are critical to specifying orientation selectivity in the feedback models, experiments in which GABA-antagonists disrupt orientation selectivity are easily understood (Sillito 1975, Tsumoto et al. 1979, Sillito et al. 1980, Bonds 1989, Pfluger and Bonds 1995). 6) The natural tendency of the cortex to amplify small inputs and converge to one of the stable attractors (the stereotyped patterns of activity) brings to mind experiments of Arieli et al. (1995), who found that even in the absence of a stimulus, waves of activity appear and propagate across the cortical surface.

Contradictions of the feedback models

The experiments most difficult to reconcile with the feedback models are the cortical inactivation experiments of Ferster et al. (1996) and of Chung and Ferster (1998). In each of these experiments, cortical activity was severely disrupted throughout the layers with one of two different methods. The amplitude of the spiking responses to visual and electrical stimuli were reduced by 90% or more (Figure 4), yet the width of orientation tuning of the EPSPs remaining in simple cells, assessed using two different types of visual stimuli, was unaffected. These results strongly suggest that the cortical circuit does not actually sharpen orientation tuning beyond what is provided by the thalamic input, other than through application of the spike threshold within individual simple cells (Carandini and Ferster 2000). Secondly, the experiments suggest that the geniculate input to simple cells constitutes a large fraction of the total input, 35-50%, so that the cortical inputs amplify the thalamic input only by a small factor of 2-3. And they do so in a manner that is independent of orientation.

These results do not rule out a strong role for feedback elsewhere in the visual cortex. The inactivation experiments were limited to simple cells in the cat visual cortex. Some of the experimental evidence for feedback, however, comes from other cell types and other species. Changes of orientation tuning over time, for example, are best documented outside of layer 4 in the monkey visual cortex (Ringach et al. 1997). These layers in the cat visual cortex have more prominent lateral connections than layer 4 and so might be in a better position to participate in feedback circuitry. Feedback may contribute to responses in many ways without completely determining the structure of cortical activity patterns as in the feedback models of orientation tuning.

These feedback models in their simplest form produce orientation tuning that is independent of all other stimulus attributes. Yet orientation tuning width narrows with increasing spatial frequency (Vidyasagar and Sigüenza 1985, Webster and De Valois 1985, Jones et al. 1987, Hammond and Pomfrett 1990). To correct the models, it is possible to separate the cortical circuit into subsets, each with its own preferred spatial frequency and its own pattern of excitatory and inhibitory connections that determine its orientation tuning width. The question arises, however, as to how many different stimulus parameters, and therefore how many different circuits, must be built into the cortex to account completely for the variability in cortical responses. A consequence of this limitation is explored in detail by Carandini and Ringach (1997), who show that the feedback models are unable to distinguish overlapping gratings of different orientations. Despite the presence of two orientations in the stimulus, the models often converge on their standard pattern of activity within a single group of active cells with a single range of orientation preferences, again defined by the geometry of the excitatory and inhibitory connections. Furthermore, adding noise to a stimulus of a single orientation results in spurious responses in cells tuned to the orthogonal orientation. Such behaviors have not been reported in cortical cells. Rigorous tests for such behavior have yet to be carried out, however, and would form an extremely strong test of the feedback models.

ORIENTATION SELECTIVITY AND CORTICAL COMPUTATION

At the outset of this review, we suggested that orientation selectivity serves as a model system for understanding cortical computation. What conclusions can we draw from our view of the function of the striate cortex? Our survey suggests a set of provocative, if frankly speculative, ideas.

The three salient features of the model of cat layer 4 for which there is strong experimental evidence are orientation-specific feedforward excitation, strong push-pull inhibition, and a weaker recurrent excitation to amplify responses. The intracortical circuitry can be

summarized as "correlation-based": excitatory cells connect to cells that are well correlated in activity; inhibitory cells connect to cells that are anti-correlated, or minimally coactive. Furthermore, a subset of the inhibitory cells must be more directly responsive to the inputs, and thus have broader tuning, than the excitatory cells. This suggests several candidate principles for the layer 4 cortical circuit. First, the entire circuit, including both feedforward and intracortical connections, can develop based on activity-dependent rules guided simply by the activity patterns of the feedforward inputs. Second, the circuit is very local: Cells need not integrate information even across an entire hypercolumn, but may restrict interactions to only a very local region, about 1/3 of a hypercolumn, representing $\pm 30^o$. Third, the pattern of activity is input-driven: e.g., inputs that stimulate cells with a broader or narrower range of preferred orientations elicit correspondingly broader or narrower activity patterns in the cortex. Fourth, the feedforward inhibition, which is directly driven by the input, is stronger than feedforward excitation and responds more like this input than do other cell. Studies in the rat whisker barrel system also indicate that the layer 4 computation is local, input-driven, and dependent on inhibitory responses that more directly reflect the thalamic input than do excitatory responses (Simons and Carvell 1989, Brumberg et al. 1996, Pinto et al. 1996, Goldreich et al. 1999). It will be of great interest to determine if an analogue of push-pull inhibition can be found in layer 4 of this and other systems.

What computation might this circuit perform? It can allow layer 4 cells in visual cortex to recognize a given orientation independent of the stimulus contrast (Troyer et al. 1998). But this specific task can be abstracted to encompass more general rules of feature extraction, in particular the task of recognizing stimulus form independent of stimulu magnitude. Call the input set driving a given cell $\mathcal{A}$, which for a simple cell is the activity generated in the relay cells by optimally oriented light and dark bars in the ON and OFF subfields. Push-pull inhibition generalizes to inhibitory input from the pattern $\overline{\mathcal{A}}$, the set of inputs most anticorrelated, or least coactive, with $\mathcal{A}$. For a simple cell, $\mathcal{A}$ is the same as $\overline{\mathcal{A}}$ except generated by stimuli with light and dark bars reversed. Finally, there is a large set of patterns, $\mathcal{B}$, that share some inputs with both $\mathcal{A}$ and $\overline{\mathcal{A}}$, but are uncorrelated — only randomly coactive — with each. In simple cells, these patterns correspond to input activity generated by stimuli of the orthogonal orientation. In simple cells receiving the input $\mathcal{A}$ alone (Figure 4D), we have seen that orientation selectivity becomes contrast dependent, because input pattern $\mathcal{B}$ of sufficiently large amplitude (an orthogonal stimulus of high contrast) can activate the cell. Adding strong push-pull inhibition translates into making the cell selective for the pattern $\left(\mathcal{A} \text{ AND NOT } \overline{\mathcal{A}}\right)$. As a result, $\mathcal{B}$ of any strength, since it activates both $\mathcal{A}$ and $\overline{\mathcal{A}}$ to some degree, can no longer activate the cell when push-pull inhibition is present. The cell becomes selective for pattern $\mathcal{A}$, independent of stimulus magnitude.

Thus, we postulate that layer 4 locally divides its inputs into opposing pairs of correlated

input structures such that a cell responds only when one is present without the other. Layer 4, in turn, projects to layers 2/3 where in cat V1 we find complex cells that respond to a given stimulus orientation independent of its polarity. That is, while layer 4 cells seem to respond to $\left(\mathcal{A} \text{ AND NOT } \overline{\mathcal{A}}\right)$ layer 2/3 cells respond to something more like $\left(\mathcal{A} \text{ OR } \overline{\mathcal{A}}\right)$, extracting the element that the two opposites have in common (orientation), while discarding the elements that distinguish them (polarity). These ideas of opposition followed by synthesis as possible roots of mental processing are reminiscent both of many eastern philosophies and of "dialectical" western philosophies (*e.g.*, Merleau-Ponty 1962).

The feedback models, in contrast, incorporate a completely different philosophy of cortical processing. In these models, the cortex converges on stereotypical patterns of activity in response to a variety of stimuli. The architecture of the cortical circuit determines in advance how many different modes of response, and therefore how many different stimuli, can be encoded by the cortex. Those stimuli that do not conform to the predefined patterns will be represented as the nearest such pattern, and two or more patterns cannot easily be simultaneously represented. In the feed-forward model with strong push-pull inhibition, the cortex is more flexible in its response to the visual image. Different stimuli that evoke sufficiently different patterns of thalamic input will almost invariably evoke different patterns of cortical activity. Given the fundamental differences between the two models, determining which mode of operation the cortex uses (if indeed it uses either one) becomes all the more interesting.

References

L. F. Abbott, J. A. Varela, K. Sen, and S. B. Nelson. Synaptic depression and cortical gain control. *Science*, 275:220–224, 1997.

P. Adorján, J. B. Levitt, J. S. Lund, and K. Obermayer. A model for the intracortical origin of orientation preference and tuning in macaque striate cortex. *Vis. Neurosci.*, 16:303–318, 1999.

B. Ahmed, J. C. Anderson, R. J. Douglas, K. A. Martin, and J. C. Nelson. Polyneuronal innervation of spiny stellate neurons in cat visual cortex. *J. Comp. Neurol.*, 341:39–49, 1994.

D. G. Albrecht. Visual cortex neurons in monkey and cat: Effect of contrast on the spatial and temporal phase transfer functions. *Vis. Neurosci.*, 12:1191–1210, 1995.

D. G. Albrecht, S. B. Farrar, and D. B. Hamilton. Spatial contrast adaptation characteristics of neurones recorded in the cat's visual cortex. *J. Physiol.*, 347:713–739, 1984.

D. G. Albrecht and W. S. Geisler. Motion selectivity and the contrast-response function of simple cells in the visual cortex. *Vis. Neurosci.*, 7:531–546, 1991.

D. G. Albrecht and D. B. Hamilton. Striate cortex of monkey and cat: contrast response function. *J. Neurophysiol.*, 48:217–237, 1982.

J. D. Allison and K. A. C. Martin. Contrast adaptation produced by null direction and cross orientation stimulation of neurons in cat visual cortex. *Soc. Neurosci. Abstr.*, 23: 454, 1997.

J. M. Alonso, W. M. Usrey, and R. C. Reid. Precisely correlated firing in cells of the lateral geniculate nucleus. *Nature*, 383:815–819, 1996.

J. Anderson, M. Carandini, and D. Ferster. Orientation tuning of input conductance in cat primary visual cortex. *J. Neurosci.*, (submitted), 1999.

A. Arieli, D. Shoham, R. Hildesheim, and A. Grinvald. Coherent spatiotemporal patterns of ongoing activity revealed by real-time optical imaging coupled with single-unit recording in the cat visual cortex. *J. Neurophysiol.*, 73:2072–2093, 1995.

R. Azouz, C. M. Gray, L. G. Nowak, and D. A. McCormick. Physiological properties of inhibitory interneurons in cat striate cortex. *Cerebral Cortex*, 7:534–545, 1997.

L. A. Bauman and A. B. Bonds. Inhibitory refinement of spatial frequency selectivity in single cells of the cat striate cortex. *Vis. Res.*, 31:933–944, 1991.

R. Ben-Yishai, R. L. Bar-Or, and H. Sompolinsky. Theory of orientation tuning in visual cortex. *Proc. Natl. Acad. Sci. USA*, 92:3844–3848, 1995.

R. Ben-Yishai, D. Hansel, and H. Sompolinsky. Traveling waves and the processing of weakly tuned inputs in a cortical network module. *J. Comput. Neurosci.*, 4:57–77, 1997.

P. O. Bishop, J. S. Coombs, Henry, and G. H. Receptive fields of simple cells in the cat striate cortex. *J. Physiol.*, 231:31–60, 1973.

G. G. Blasdel and D. Fitzpatrick. Physiological organization of layer 4 in macaque striate cortex. *J. Neurosci.*, 4:880–895, 1984.

A. B. Bonds. Role of inhibition in the specification of orientation selectivity of cells in the cat striate cortex. *Vis. Neurosci.*, 2:41–55, 1989.

L. J. Borg-Graham, C. Monier, and Y. Frégnac. Visual input evokes transient and strong shunting inhibition in visual cortical neurons. *Nature*, 393:369–373, 1998.

J. C. Brumberg, D. J. Pinto, and D. J. Simons. Spatial gradients and inhibitory summation in the rat whisker barrel system. *J. Neurophysiol.*, 76:130–140, 1996.

J. Bullier and G. H. Henry. Laminar distribution of first-order neurons and afferent terminals in cat striate cortex. *J. Neurophysiol.*, 42:1271–1281, 1979.

J. Bullier, M. J. Mustari, and G. H. Henry. Receptive-field transformations between LGN neurons and S-cells of cat striate cortex. *J. Neurophysiol.*, 47:417–438, 1982.

M. Carandini and D. Ferster. A tonic hyperpolarization underlying contrast adaptation in cat visual cortex. *Science*, 276:949–952, 1997.

M. Carandini and D. Ferster. The contribution of firing threshold to orientation tuning in cat primary visual cortex. *J. Neurosci.*, In Press, 2000.

M. Carandini and D. J. Heeger. Summation and division by neurons in visual cortex. *Science*, 264:1333–1336, 1994.

M. Carandini, D. J. Heeger, and J. A. Movshon. Linearity and normalization in simple cells of the macaque primary visual cortex. *J. Neurosci.*, 17:8621–8644, 1997.

M. Carandini and D. Ringach. Predictions of a recurrent model of orientation selectivity. *Vision Res.*, 37:3061–3071, 1997.

F. S. Chance, S. B. Nelson, and L. F. Abbott. Synaptic depression and the temporal response characteristics of V1 cells. *J. Neurosci.*, 18:4785–4799, 1998.

B. Chapman and M. P. Stryker. Development of orientation selectivity in ferret visual cortex and effects of deprivation. *J. Neurosci.*, 13:5251–5262, 1993.

B. Chapman, K. R. Zahs, and M. P. Stryker. Relation of cortical cell orientation selectivity to alignment of receptive fields of the geniculocortical afferents that arborize within a single orientation column in ferret visual cortex. *J. Neurosci.*, 11:1347–1358, 1991.

H. Cheng, Y. M. Chino, E. L. Smith, J. Hamamoto, et al. Transfer characteristics of lateral geniculate nucleus X neurons in the cat: Effects of spatial frequency and contrast. *J. Neurophysiol.*, 74:2548–2557, 1995.

S. Chung and D. Ferster. Strength and orientation tuning of the thalamic input to simple cells revealed by electrically evoked cortical suppression. *Neuron*, 20:1177–89, 1998.

M. C. Crair, D. C. Gillespie, and M. P. Stryker. The role of visual experience in the development of columns in cat visual cortex. *Science*, 279:566–570, 1998.

J. M. Crook, Z. F. Kisvardy, and U. T. Eysel. GABA induced inactivation of functionally characterized sites in cat striate cortex: effects on orientation and direction selectivity. *Vis. Neurosci.*, 14:141–158, 1995.

R. L. De Valois, L. G. Thorell, and D. G. Albrecht. Periodicity of striate-cortex-cell receptive fields. *J. Opt. Soc. Am. A*, 2:1115–1123, 1985.

A. F. Dean. The relationship between response amplitude and contrast for cat striate cortical neurons. *J. Physiol.*, 318:413–427, 1981.

A. F. Dean and D. J. Tolhurst. Factors influencing the temporal phase of response to bar and grating stimuli for simple cells in the cat straite cortex. *Exp. Brain Res.*, 62:143–151, 1986.

G. C. DeAngelis, J. G. Robson, I. Ohzawa, and R. D. Freeman. Organization of suppression in receptive fields of neurons in cat visual cortex. *J. Neurophysiol.*, 68:144–163, 1992.

R. J. Douglas, C. Koch, M. Mahowald, K. A. Martin, and H. H. Suarez. Recurrent excitation in neocortical circuits. *Science*, 269:981–985, 1995.

R. J. Douglas and K. A. C. Martin. A functional microcircuit for cat visual cortex. *J. Physiol.*, 440:735–769, 1991.

R. J. Douglas, K. A. C. Martin, and D. Whitteridge. Selective responses of visual cortical cells do not depend on shunting inhibition. *Nature*, 332:642–644, 1988.

R. J. Douglas, K. A. C. Martin, and D. Whitteridge. An intracellular analysis of the visual responses of neurones in cat visual cortex. *J. Physiol.*, 440:659–696, 1991.

G. Einstein, T. L. Davis, and P. Sterling. Ultrastructure of synapses from the a-laminae of the lateral geniculate nucleus in layer iv of the cat striate cortex. *J. Comp. Neurol.*, 260: 63–75, 1987.

U. T. Eysel, J. M. Crook, and H. F. Machemer. GABA-induced remote inactivation reveals cross-orientation inhibition in the cat striate cortex. *Exp. Brain Res.*, 80:626–630, 1990.

D. Ferster. A comparison of binocular depth mechanisms in areas 17 and 18 of the cat visual cortex. *J. Physiol.*, 311:623–655, 1981.

D. Ferster. Orientation selectivity of synaptic potentials in neurons of cat primary visual cortex. *J. Neurosci.*, 6:1284–1301, 1986.

D. Ferster. Spatially opponent excitation and inhibition in simple cells of the cat visual cortex. *J. Neurosci.*, 8:1172–1180, 1988.

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

D. Ferster and B. Jagadeesh. Nonlinearity of spatial summation in simple cells of areas 17 and 18 of cat visual cortex. *J. Neurophysiol.*, 66:1667–1679, 1991.

D. Ferster and B. Jagadeesh. EPSP-IPSP interactions in cat visual cortex studied with in vivo whole-cell patch recording. *J. Neurosci.*, 12(4):1262–1274, 1992.

D. Ferster and S. Lindström. An intracellular analysis of geniculo-cortical connectivity in area 17 of the cat. *J. Physiol.*, 342:181–215, 1983.

Y. Fregnac and M. Imbert. Development of neuronal selectivity in the primary visual cortex of the cat. *Physiol. Rev.*, 64:325–434, 1984.

J. L. Gardner, A. Anzai, I. Ohzawa, and R. D. Freeman. Linear and nonlinear contributions to orientation tuning of simple cells in the cat's striate cortex. *Vis. Neurosci.*, In press, 1999.

W. S. Geisler and D. G. Albrecht. Cortical neurons: isolation of contrast gain control. *Vis. Res.*, 32:1409–1410, 1992.

J. R. Gibson, M. Beierlein, and B. W. Connors. Two networks of electrically coupled inhibitory neurons in neocortex. *Nature*, 462:75–79, 1999.

Z. Gil, B. Connors, and Y. Amitai. Differential regulation of neocortical synapses by neuromodulators and activity. *Neuron*, 19:679–686, 1997.

Z. Gil, B. Connors, and Y. Amitai. Efficacy of thalamocortical and intracortical synaptic connections: Quanta, innervation, and reliability. *Neuron*, 23:385–397, 1999.

C. D. Gilbert. Laminar differences in receptive field properties of cells in cat primary visual cortex. *J. Physiol.*, 268:391–421, 1977.

C. D. Gilbert and T. N. Wiesel. Morphology and intracortical projections of functionally characterised neurones in the cat visual cortex. *Nature*, 280:120–125, 1979.

D. Goldreich, H. T. Kyriazi, and D. J. Simons. Functional independence of layer IV barrels in rodent somatosensory cortex. *J. Neurophysiol.*, 82:1311–1316, 1999.

B. Gulyas, G. A. Orban, J. Duysens, and H. Maes. The suppressive influence of moving textured backgrounds on responses of cat striate neurons to moving bars. *J. Neurophysiol.*, 57:1767–1791, 1987.

P. Hammond and D. M. MacKay. Modulatory influences of moving textured backgrounds on responsiveness of simple cells in feline striate cortex. *J. Physiol.*, 319:431–442, 1981.

P. Hammond and C. J. Pomfrett. Influence of spatial frequency on tuning and bias for orientation and direction in the cat's striate cortex. *Vision Res.*, 30:359–369, 1990.

D. Hansel and H. Sompolinsky. Chaos and synchrony in a model of a hypercolumn in visual cortex. *J. Comput. Neurosci.*, 3:7–34, 1996.

M. J. Hawken and A. J. Parker. Contrast sensitivity and orientation selectivity in lamina IV of the striate cortex of old world monkeys. *Exp. Brain Res.*, 54:367–372, 1984.

D. J. Heeger. Normalization of cell responses in cat striate cortex. *Vis. Neurosci.*, 9:181–198, 1992.

P. Heggelund. Receptive field organization of simple cells in cat striate cortex. *Exp. Brain Res.*, 42:89–98, 1981.

J. A. Hirsch, J.-M. Alonso, R. C. Reid, and L. Martinez. Synaptic integration in striate cortical simple cells. *J. Neurosci.*, 18:9517–9528, 1998.

J. A. Hirsch and C. D. Gilbert. Synaptic physiology of horizontal connections in the cat's visual cortex. *J. Neurosci.*, 11:1800–1809, 1991.

D. H. Hubel and T. N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.*, 160:106–154, 1962.

D. H. Hubel and T. N. Wiesel. The period of susceptibility to the physiological effects of unilateral eye closure in kittens. *J. Physiol.*, 206:419–436, 1970.

A. L. Humphrey and T. T. Norton. Topographic organization of the orientation column system in the striate cortex of the tree shrew (tupaia glis). I. Microelectrode recording. *J. Comp. Neurol.*, 192:531–547, 1980.

B. Jagadeesh, C. M. Gray, and D. Ferster. Visually evoked oscillations of membrane potential in cells of cat visual cortex. *Science*, 257:552–554, 1992.

J. P. Jones and L. A. Palmer. An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. *J. Neurophysiol.*, 58:1233–1258, 1987a.

J. P. Jones and L. A. Palmer. The two-dimensional spatial structure of simple receptive fields in cat striate cortex. *J. Neurophysiol.*, 58:1187–1211, 1987b.

J. P. Jones, A. Stepnoski, and L. A. Palmer. The two-dimensional spectral structure of simple receptive fields in cat striate cortex. *J. Neurophysiol.*, 59:1212–1232, 1987.

A. S. Kayser and K. D. Miller. The development of the cat layer 4 orientation circuit: Origins of columnar organization and contrast invariance. *Soc. Neurosci. Abstr.*, 24:261, 1998.

A. S. Kayser, N. J. Priebe, and K. D. Miller. Contrast-dependent nonlinearities arise locally in a model of contrast-invariant orientation tuning. *J. Neurosci.*, (submitted), 1999.

J. P. Kelly and D. C. van Essen. Cell structure and function in the visual cortex of the cat. *J. Physiol.*, 238:515–547, 1974.

A. E. Krukowski, A. Hoffman, and K. D. Miller. Correlation-based intracortical connectivity in striate cortex can account for temporal frequency tuning and 'cross-orientation' inhibition. *Soc. Neurosci. Abstr.*, 24:261, 1998.

S. LeVay and C. D. Gilbert. Laminar patterns of geniculocortical projection in the cat. *Brain Res.*, 113:1–19, 1976.

C. Y. Li and O. Creutzfeldt. The representation of contrast and other stimulus parameters by single neurons in area 17 of the cat. *Pflugers. Arch.*, 401:304–314, 1984.

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

P. E. Maldonado, I. Gödecke, C. M. Gray, and T. Bonhoeffer. Orientation selectivity in pinwheel centers in cat striate cortex. *Science*, 276:1551–1555, 1997.

H. Markram and M. Tsodyks. Redistribution of synaptic efficacy between neocortical pyramidal neurons. *Nature*, 382:807–810, 1996.

K. A. Martin and D. Whitteridge. Form, function and intracortical projections of spiny neurones in the striate visual cortex of the cat. *J. Physiol.*, 353:463–504, 1984.

D. A. McCormick, B. W. Connors, J. W. Lighthall, and D. A. Prince. Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *J. Neurophysiol.*, 54:782–805, 1985.

J. McLean and L. A. Palmer. Contrast adaptation and excitatory amino acid receptors in cat striate cortex. *Vis. Neurosci.*, 13:1069–1087, 1996.

M. Merleau-Ponty. *Phenomenology of perception*. Humanities Press, New York, 1962.

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

K. D. Miller, E. Erwin, and A. Kayser. Is the development of orientation selectivity instructed by activity? *J. Neurobiol.*, 41:44–57, 1999.

M. C. Morrone, D. C. Burr, and L. Maffei. Functional implications of cross-orientation inhibition of cortical visual cells. I. Neurophysiological evidence. *Proc. R. Soc. Lond. B*, 216:335–354, 1982.

J. A. Movshon and R. C. Van Sluyters. Visual neural development. *Ann. Rev. Psychol.*, 32: 477–522, 1981.

W. H. Mullikin, J. P. Jones, and L. A. Palmer. Receptive-field properties and laminar distribution of X-like and Y-like simple cells in cat area 17. *J. Neurophysiol.*, 52:351–371, 1984.

S. Nelson, L. Toth, B. Sheth, and M. Sur. Orientation selectivity of cortical neurons during intracellular blockade of inhibition. *Science*, 265:774–777, 1994.

S. B. Nelson. Temporal interactions in the cat visual system. I. Orientation-selective suppression in the visual cortex. *J. Neurosci.*, 11:344–356, 1991.

I. Ohzawa, G. Sclar, and R. D. Freeman. Contrast gain control in the cat's visual system. *J. Neurophysiol.*, 54:651–667, 1985.

L. A. Palmer and T. L. Davis. Receptive-field structure in cat striate cortex. *J. Neurophysiol.*, 46:260–276, 1981.

X. Pei, T. R. Vidyasagar, M. Volgushev, and O. D. Creutzfeldt. Receptive field analysis and orientation selectivity of postsynaptic potentials of simple cells in cat visual cortex. *J. Neurosci.*, 14:7130–7140, 1994.

X. Pei, M. Volgushev, T. R. Vidyasagar, and O. D. Creutzfeldt. Whole cell recording and conductance measurements in cat visual cortex in-vivo. *NeuroReport*, 2:485–488, 1991.

A. Peters and B. R. Payne. Numerical relationships between geniculocortical afferents and pyramidal cell modules in cat primary visual cortex. *Cerebral Cortex*, 3:69–78, 1993.

B. Pfluger and A. B. Bonds. Dynamic diferentiation of GABAA-sensitive influences on orientation selectivity of complex cells in the cat visual cortex. *Exp. Brain Res.*, 104: 81–88, 1995.

D. J. Pinto, J. C. Brumberg, D. J. Simons, and G. B. Ermentrout. A quantitative population model of whisker barrels: Re-examining the Wilson-Cowan equations. *J. Comput. Neurosci.*, 3:247–264, 1996.

R. C. Reid and J. M. Alonso. Specificity of monosynaptic connections from thalamus to visual cortex. *Nature*, 378:281–284, 1995.

D. L. Ringach, M. J. Hawken, and R. Shapley. Dynamics of orientation tuning in macaque primary visual cortex. *Nature*, 387:281–284, 1997.

A. C. Rosenquist, S. B. Edwards, and L. A. Palmer. An autoradiographic study of the projections of the dorsal lateral geniculate nucleus and the posterior nucleus in the cat. *Brain Res.*, 80:71–93, 1975.

M. V. Sanchez-Vives, L. G. Nowak, and D. A. McCormick. Cellular and network mechanisms generating adaptation to contrast in the visual cortex: An in vivo and in vitro study. *Soc. Neurosci. Abstr.*, 23:1944, 1997.

A. B. Saul and A. L. Humphrey. Evidence of input from lagged cells in the lateral geniculate nucleus to simple cells in cortical area 17 of the cat. *J. Neurophysiol.*, 68:1190–1208, 1992.

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

G. Sclar and R. D. Freeman. Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. *Exp. Brain Res.*, 46:457–461, 1982.

C. J. Shatz and M. P. Stryker. Ocular dominance in layer IV of the cat's visual cortex and the effects of monocular deprivation. *J. Physiol.*, 281:267–283, 1978.

A. M. Sillito, J. A. Kemp, J. A. Milson, and N. Berardi. A re-evaluation of the mechanisms underlying simple cell orientation selectivity. *Brain Res.*, 194:517–520, 1980.

M. Sillito, A. The contribution of inhibitory mechanisms to the receptive field properties of neurones in the striate cortex of the cat. *J. Physiol.*, 250:305–329, 1975.

D. J. Simons and G. E. Carvell. Thalamocortical response transformation in the rat vibrissa/barrel system. *J. Neurophysiol.*, 61:311–330, 1989.

B. C. Skottun, A. Bradley, G. Sclar, I. Ohzawa, and R. D. Freeman. The effects of contrast on visual orientation and spatial frequency discrimination: A comparison of single cells and behavior. *J. Neurophysiol.*, 57:773–786, 1987.

D. Somers, S. B. Nelson, and M. Sur. An emergent model of orientation selectivity in cat visual cortical simple cells. *J. Neurosci.*, 15:5448–5465, 1995.

H. Sompolinsky and R. Shapley. New perspectives on the mechanisms for orientation selectivity. *Curr. Opin. Neurobiol.*, 7:514–522, 1997.

K. J. Stratford, K. Tarczy-Hornoch, K. A. Martin, N. J. Bannister, and J. J. Jack. Excitatory synaptic inputs to spiny stellate cells in cat visual cortex. *Nature*, 382:258–261, 1996.

H. Suarez, C. Koch, and R. Douglas. Modeling direction selectivity of simple cells in striate visual cortex within the framework of the canonical microcircuit. *J. Neurosci.*, 15:6700–6719, 1995.

K. Tanaka. Cross-correlation analysis of geniculostriate neuronal relationships in cats. *J. Neurophysiol.*, 49:1303–1318, 1983.

J. B. Troy and C. Enroth-Cugell. X and Y ganglion cells inform the cat's brain about contrast in the retinal image. *Exp. Brain Res.*, 93:383–390, 1993.

T. W. Troyer, A. Krukowski, N. J. Priebe, and K. D. Miller. Contrast-invariant orientation tuning in cat visual cortex: Feedforward tuning and correlation-based intracortical connectivity. *J. Neurosci.*, 18:5908–5927, 1998.

M. V. Tsodyks and H. Markram. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proc. Natl. Acad. Sci. USA*, 94, 1997.

T. Tsumoto, W. Eckart, and O. D. Creutzfeld. Modification of orientation sensitivity of cat visual cortex neurons by removal of GABA-mediated inhibition. *Exp. Brain Res.*, 34: 351–363, 1979.

T. R. Vidyasagar and J. A. Sigüenza. Relationship between orientation tuning and spatial frequency in neurones of cat area 17. *Exp. Brain Res.*, 57:628–631, 1985.

M. A. Webster and R. L. De Valois. Relationship between spatial-frequency and orientation tuning of striate-cortex cells. *J. Opt. Soc. Am. A.*, 2:1124–1132, 1985.