

Understanding Layer 4 of the Cortical Circuit: A Model Based on Cat V1

Kenneth D. Miller

Departments of Physiology and Otolaryngology, W.M. Keck Center for Integrative Neuroscience, Sloan-Swartz Center for Theoretical Neurobiology at UCSF, University of California, San Francisco, CA 94143-0444, USA

This paper reviews theoretical and experimental results on the processing of layer 4, the input-recipient layer, of cat primary visual cortex (V1). A wide range of experimental data can be understood from a model in which response tuning of layer 4 cells is largely determined by a local interplay of feedforward excitation (from thalamus) and feedforward inhibition (from layer 4 inhibitory interneurons driven by thalamus). Feedforward inhibition dominates excitation, inherits its tuning from the thalamic input and sharpens the tuning of excitatory cells. At least a strong component of the feedforward inhibition received by a cell is spatially opponent to the excitation it receives, meaning that inhibition is driven by dark in regions of the visual field in which excitation is driven by light, and vice versa. The idea of opponent inhibition can be generalized to mean inhibition driven by input patterns that are strongly anticorrelated with the patterns that excite a cell. This paper argues that dominant feedforward opponent inhibition may be a general principle of cortical layer 4. This leads to the suggestion that the properties that show columnar organization – invariance across the vertical depth of cortex – may be properties that are shared by ‘opposite’ (anticorrelated) stimulus pairs. This contrasts with the more common idea that a column represents a set of cells that all share similar stimulus preferences.

Introduction

The cerebral cortex has a stereotyped six-layer structure [reviewed by Callaway (Callaway, 1998)]. ‘Feedforward’ inputs, which for primary sensory cortex come from thalamus and represent the sensory periphery, primarily innervate layer 4. Layer 4 cells project strongly to layers 2/3, which in turn provide feedforward input to layer 4 of the next higher cortical area as well as projecting to the deep layers. The deep layers in turn provide feedback to layers 2–4 and thalamus, and provide output to non-thalamic subcortical structures.

To understand the computations being performed by cortex, we need to understand the nature of the processing undertaken by each layer. The natural starting place in thinking about sensory processing is layer 4, the primary layer in which sensory input first arrives. Here I outline a picture of the processing taking place in cortical layer 4 in cat primary visual cortex (V1) that has been emerging from both experimental and theoretical work in recent years. This picture is intriguingly similar to that emerging from studies of layer 4 of rodent primary somatosensory cortex (S1), as reviewed elsewhere (Miller *et al.*, 2001; Pinto *et al.*, 2002; Swadlow, 2002). As befits the position of layer 4 as the recipient of feedforward input, this picture suggests that the response tuning of layer 4 cells is largely determined by feedforward input, including feedforward inhibition (inhibition from interneurons driven by the thalamus) as well as feedforward excitation (from the thalamus). The inhibition dominates, so that a cell can only be excited by stimuli that cause the effects of feedforward excitation and inhibition to be separated in time; concurrent engagement of the two yields a net

inhibition. Neurons providing feedforward inhibition follow the tuning of the thalamic inputs, thereby sculpting the responses of excitatory cells to have tighter tuning than the thalamic inputs. Both the feedforward excitation and inhibition that a cell receives are evoked locally, from cells preferring nearby orientations. While the feedforward input establishes initial response tuning, local recurrent excitation and neuronal non-linearities (e.g. spike threshold) enhance responses evoked by preferred versus non-preferred stimuli.

In this article I review the evidence leading to this picture, along with countervailing evidence that renders it still controversial.

The Problem Posed by the Thalamic Input

Cells in layer 4 of cat V1 are predominantly *simple cells*: cells with receptive fields (RFs) consisting of aligned, alternating ON (light-preferring) and OFF (dark-preferring) subregions (Hubel and Wiesel, 1962; Gilbert, 1977; Bullier and Henry, 1979). The subregions share a common axis of elongation, which defines the cell’s preferred orientation – the orientation of a light/dark edge that best drives the cell. A simple cell’s thalamic afferents, which come from the lateral geniculate nucleus (LGN), form a pattern matching the cell’s subregion structure: ON-center afferents have RF centers overlying ON subregions, and OFF-center afferents similarly overlie OFF subregions (Tanaka, 1983; Reid and Alonso, 1995) (Fig. 1, left).

The degree to which the feedforward excitation from thalamus determines a simple cell’s response properties has been a subject of great controversy (Sompolinsky and Shapley, 1997; Ferster and Miller, 2000). That this thalamic excitation is not sufficient to specify a simple cell’s responses can be seen from the invariance of a cell’s orientation tuning under changes in stimulus contrast (contrast is the light/dark difference in the stimulus relative to the mean luminance). A dim light bar of the cell’s preferred orientation, flashed over an ON-subregion, will weakly activate the corresponding ON-center LGN afferents (Fig. 1, middle). A bright light bar flashed orthogonal to the preferred orientation will strongly activate a subset of ON-center afferents, while suppressing the spontaneous activity of a subset of OFF-center afferents (Fig. 1, right). LGN afferents have spontaneous rates of 10–15 Hz, and can fire at rates of 100 Hz or more when stimulated. Thus, even if the number of suppressed OFF-center inputs were equal to the number of excited ON-center inputs, the bright orthogonal bar would yield a net positive LGN input, because stimulated ON-center cells raise their firing rates much more than suppressed OFF-center cells can lower their firing rates. With proper choices of brightness, one can in principle arrange that the dim preferred-orientation bar and the bright orthogonal bar elicit the same temporal pulse of LGN input. Yet a typical simple cell will respond to even a dim preferred-orientation bar, and will not respond to even a bright




Figure 1. LGN inputs to simple cells and the problem of the contrast invariance of orientation tuning. (Left) Thick lines show an idealized simple cell receptive field: solid oval in center represents an ON-subregion, dashed subregions to either side represent OFF-subregions. Thin circles represent the receptive field centers of the LGN cells found to connect to the simple cell: solid circles, centers of ON-center inputs; dashed circles, centers of OFF-center inputs. Simple cell receives input from ON-center LGN cells with centers overlying its ON-subregions, and from OFF-center LGN cells with centers overlying its OFF-subregions. Composite data from Reid and Alonso (Reid and Alonso, 1995). (Center and right) The problem of contrast-invariant orientation tuning. (Center) A dim, vertically oriented flash covering the ON-subregion (vertical rectangle) will weakly excite the ON-center inputs with underlying centers. (Right) A bright, horizontally oriented flash (horizontal rectangle) will strongly excite the subset of ON-center inputs with underlying centers, and will suppress the OFF-center inputs with underlying centers. Excited cells can have their firing rates increased greatly, while suppressed cells can only have their firing rates decreased to zero. Accordingly, the two stimuli can be arranged to give the same total pulse of LGN input, but a typical simple cell will respond to the dim vertical flash and not to the bright horizontal flash. See text for further discussion.




Figure 2. An example of contrast-invariant orientation tuning. Orientation tuning curves of a simple cell obtained with drifting sinusoidal gratings of three different contrasts (5% contrast, dashed line; 20% contrast, thin solid line; 80% contrast, thick solid line). [Adapted from Sclar and Freeman (Sclar and Freeman, 1982).]

orthogonally oriented bar. This is an example of the contrast invariance of orientation tuning (Fig. 2): the tuning curves of response versus orientation have a similar shape at low contrast (dim bars) as at high contrast (bright bars). Such contrast-invariance has been quantitatively demonstrated for steady-state responses to drifting sinusoidal luminance gratings (Sclar and Freeman, 1982; Skottun *et al.*, 1987; Anderson *et al.*, 2000b). This property demonstrates that the arrangement of LGN inputs alone is not sufficient to explain simple-cell orientation tuning.

This problem is a quite general one. Thalamic afferents that are strongly excited can have their firing rates greatly increased, while those that are suppressed can only have their firing rates reduced to zero. As a result, a non-preferred stimulus of high magnitude (high contrast) that strongly excites a subset of a cell's thalamic inputs while suppressing other inputs can yield a net positive thalamic input. This input in response to a non-preferred stimulus can be as strong as, or stronger than, the input in response to a preferred stimulus of low magnitude. Yet




Figure 3. Diagram of the model circuit for simple cells in V1 layer 4 proposed by Troyer *et al.* (Troyer *et al.*, 1998). Gray circles represent receptive fields of two excitatory neurons (top) and two inhibitory neurons; white represents ON-subregions, black represents OFF-subregions. The four receptive fields are meant to be centered at the same retinotopic point. Excitatory cells receive both feedforward LGN excitation corresponding to their receptive fields (ON-center LGN inputs with centers overlying ON-subregions, OFF-center inputs with centers overlying OFF-subregions) and antiphase feedforward inhibition. Excitatory cells also provide recurrent excitation to other cells of similar preferred orientation and phase. In the full model circuit, cells connect probabilistically, so that the illustrated connections are the most probable but other less specific connections are also made, with connections being made between neurons differing in preferred orientation by up to $\sim 30^\circ$. In addition, excitatory-to-inhibitory and inhibitory-to-inhibitory connections may exist, following the scheme that excitatory cells tend to connect to cells with similar preferred orientation and similar phase while inhibitory cells tend to connect to cells of similar preferred orientation but roughly opposite phase.

the cell should respond to its preferred stimulus even at low magnitude, and should not respond to a non-preferred stimulus even at high magnitude. This is the general form of the problem posed by the thalamic afferents: how can the cortex distinguish

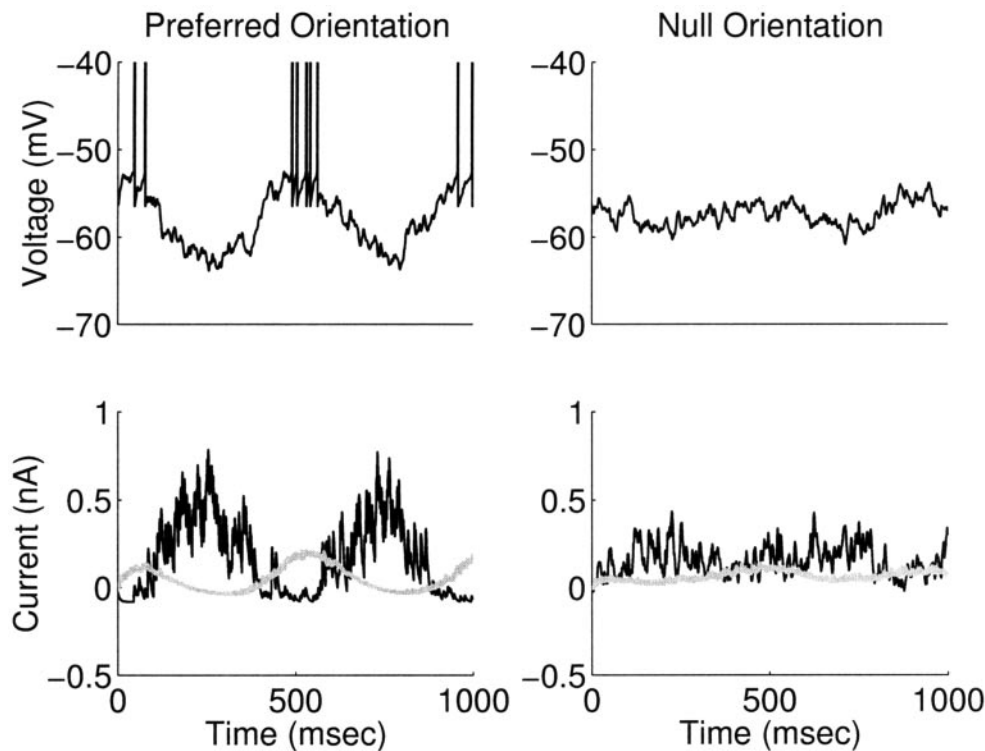


Figure 4. Illustration of responses in the model circuit to 2 Hz drifting sinusoidal gratings of the preferred orientation (left) or the orientation orthogonal to the preferred ('null' orientation, right). Top traces show membrane voltage, bottom traces show excitatory conductances (gray) and inhibitory conductances (black). Conductances are expressed as the current that would flow if the cell were voltage clamped at spike threshold voltage, and are measured relative to background [methods and parameters as previously (Krukowski and Miller, 2001), with LGN synapses having 90% NMDA onto excitatory cells and no NMDA onto inhibitory cells, no synaptic depression, and no feedback excitation]. A preferred orientation stimulus causes strong temporal modulation of the excitatory input and of the inhibitory input. Because the inhibitory input is antiphase to the excitatory input, these modulations are out of phase with one another, allowing excitation to periodically dominate and drive response. In response to the null orientation, there is little temporal modulation of the input. Because the mean inhibition is stronger than the mean excitation, the cell does not respond.

these two stimuli that yield similar total strengths of thalamic input?

Opponent Inhibition Provides a Solution to the Problem Posed by the Thalamic Input

We have recently proposed (Troyer *et al.*, 1998) that consideration of feedforward inhibition along with feedforward LGN excitation suffices to explain the contrast-invariance of orientation tuning and a variety of other response properties. To understand this proposal, we need to first define some terms. We shall refer to two simple cells of the same preferred orientation as having the same phase if, in the region in which their receptive fields overlap, their ON-subregions overlap and their OFF-subregions overlap. We shall refer to two such neurons as having opposite phase or being antiphase to one another if, in the region of overlap, ON-subregions of one overlap OFF-subregions of the other. (To be precise: we are using 'phase' to refer to absolute phase in the visual world, rather than phase relative to the receptive field center.)

The model of Troyer *et al.* (Troyer *et al.*, 1998) was inspired by intracellular recordings demonstrating (i) that the inhibition and the excitation that a simple cell receives have similar orientation tuning, with both being maximal at the preferred orientation (Ferster, 1986; Martinez *et al.*, 1998; Anderson *et al.*, 2000a), but (ii) that inhibition and excitation are evoked by stimuli of opposite light/dark polarity at any given point in the receptive field (Ferster, 1988; Hirsch *et al.*, 1998) [but see (Borg-Graham *et al.*, 1998)], i.e. in an ON-subregion, light evokes excitation and

dark evokes inhibition, while in an OFF-subregion dark evokes excitation and light evokes inhibition. This can be summarized by saying that the receptive field of the inhibition received by a simple cell is antiphase to the receptive field of the excitation the cell receives. This is also described by saying that the inhibition a cell receives is spatially *opponent* to the excitation it receives. These findings inspired a circuit model (Troyer *et al.*, 1998) in which inhibitory cells tend to project to cells of similar preferred orientation but roughly opposite phase, while excitatory cells tend to project to cells of similar preferred orientation and phase (Fig. 3). A key feature is that the feedforward (LGN-driven) antiphase inhibition is stronger than the feedforward LGN excitation; this is consistent with the experimental fact that an electrical shock to LGN, which indiscriminately activates both feedforward excitation and feedforward inhibition, yields strong inhibition in cortex (Ferster and Jagadeesh, 1992).

The strong antiphase inhibition solves the problem posed above. A bright bar orthogonal to the preferred orientation will roughly equally excite both the excitatory LGN input to a cell and the input to the cell's antiphase feedforward inhibitory inputs. Because the inhibition dominates, the cell will not fire (Fig. 4). More generally, the antiphase inhibition achieves contrast-invariant orientation tuning (Troyer *et al.*, 1998). For a stimulus to excite a cell, it must excite the cell's inputs much more strongly than it excites the cell's antiphase inhibition. This can only be achieved by a narrow range of stimulus orientations around the preferred orientation, and this range stays roughly invariant under changes of stimulus contrast. Note that this picture requires that the antiphase inhibition be evoked even by

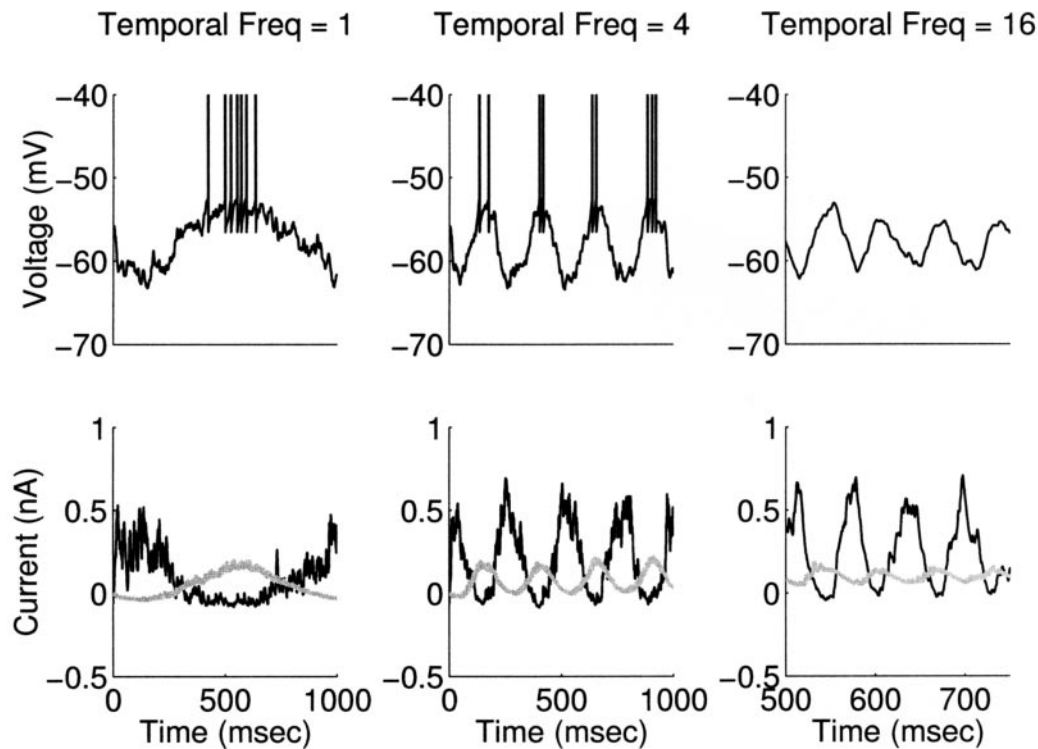


Figure 5. Illustration of responses in the model circuit to drifting sinusoidal gratings of the preferred orientation at varying temporal frequencies (1, 4, 16 Hz from left to right). Conventions and parameters as in Figure 4. The circuit includes NMDA conductances in the thalamocortical synapses onto excitatory cells, with the proportion of NMDA to AMPA conductances matched to that observed at the oldest ages studied in thalamocortical synapses in slices (Crair and Malenka, 1995). The slow NMDA conductances low-pass filter the excitatory input to the cell, so that the excitation fails to show much temporal modulation at higher frequencies. This along with the overall dominance of the inhibition causes the cell to fail to respond to higher temporal frequencies, even though these frequencies strongly drive the cell's LGN inputs. Comparison to Figure 4 should make clear that the same basic mechanism accounts for both orientation tuning and temporal frequency tuning.

stimuli orthogonal to the preferred. That is, the feedforward inhibition has orientation tuning like that of the thalamic inputs: it responds to *all* orientations, although it is driven best by the preferred orientation. The feedforward inhibition has tuning that mirrors the thalamic tuning, allowing it to sharpen the tuning of the excitatory cells.

We can generalize this idea. Suppose that a cell's preferred stimulus excites some set of inputs \mathcal{A} , and that it suppresses another set of inputs $\bar{\mathcal{A}}$. For a simple cell, the excited inputs \mathcal{A} are ON-center cells overlying an ON subregion and OFF-center cells overlying an OFF subregion, while the suppressed inputs $\bar{\mathcal{A}}$ are OFF-center cells overlying an ON subregion and ON-center inputs overlying an OFF subregion. Suppose the cell receives strong feedforward inhibition driven by the inputs $\bar{\mathcal{A}}$ that are suppressed by a preferred stimulus. In response to a preferred stimulus, this causes disinhibition that adds to the response. But a non-preferred stimulus that excites some elements of \mathcal{A} will also excite some elements of $\bar{\mathcal{A}}$ – e.g. the bright bar orthogonal to the preferred orientation excites some ON-center cells overlying ON subregions (\mathcal{A}) and also excites some ON-center cells overlying OFF subregions ($\bar{\mathcal{A}}$). The dominant inhibition driven by $\bar{\mathcal{A}}$ will prevent the cell from responding to the non-preferred stimulus, at any magnitude. Thus, we generalize the idea of opponent inhibition to denote *inhibition driven by the stimulus pattern most anticorrelated with the pattern that excites the cell* – i.e. inhibition driven by the inputs that are suppressed by the preferred stimulus. Opponent inhibition solves the problem posed by the thalamic inputs – it filters out responses to non-preferred stimuli at any magnitude, while allowing responses to preferred stimuli even at low magnitude.

The model circuit also includes recurrent excitation among cells of similar orientation and phase preference, i.e. among cells with similar preferred stimuli. This serves to amplify responses to effective stimuli without altering tuning.

Opponent Inhibition Can Explain a Range of Other Response Properties

The same model circuit can also account for the temporal frequency tuning of cortical cells (Kukowski and Miller, 2001), which cuts off at much lower frequencies than LGN tuning. An essential idea of the model circuit is that feedforward inhibition dominates over LGN excitation, so that any stimulus that causes the two to arrive together will fail to elicit a response. The reason why a cell will not respond to a non-preferred orientation is that it evokes feedforward inhibition at the same time as it evokes feedforward excitation. It turns out the very same principle can explain cortical temporal tuning (Fig. 5).

To understand this, consider what happens as we start increasing the temporal frequency of a preferred-orientation drifting grating stimulus. If the time over which excitation or inhibition persists becomes comparable to the period of the grating, then excitation and inhibition will come to overlap in time. In this case, since inhibition dominates, the cell will fail to respond. (Another way to say the same thing is that if the excitation or inhibition are low-pass filtered so that they lose their antiphase temporal modulations and become reduced to their means, the mean inhibition will dominate the mean excitation.) What sets the time over which excitation and inhibition persist? One factor is the membrane time constant of the cell, the time over which synaptic currents are integrated into membrane

voltage. It turns out this is too short to explain the temporal frequencies at which cortical response cuts off. Another factor is the time course of the synaptic conductances. Both *N*-methyl-D-aspartate (NMDA)-receptor-mediated excitatory conductances and γ -aminobutyric acid (GABA_B)-receptor-mediated inhibitory conductances have long time courses, persisting for times of the order of 100 ms. These have the right time course to explain cortical temporal frequency tuning. Indeed we found (Krukowski and Miller, 2001) that if NMDA is present in thalamocortical synapses onto excitatory cells in the proportions observed in thalamocortical slices (Crair and Malenka, 1995), this, along with our inhibition-dominated model circuit, suffices to explain the temporal frequency cutoffs of cortical cells. The fact that the same principle explains two such different response properties – orientation tuning and temporal tuning (cf. Figs 4 and 5) – increases confidence that the model captures something correct about the biology.

The model circuit can also explain a number of contrast-dependent non-linearities (Kayser *et al.*, 2001; Lauritzen *et al.*, 2001) that had previously been proposed to require ‘normalizing’ inhibition derived equally from cells of all stimulus preferences [the ‘normalization’ model (Carandini *et al.*, 1999)]. The normalization model begins with the idea that the input to simple cells derives from a linear filtering of the stimulus. This accords with the many response properties of simple cells that appear linear (up to rectification). For example, a linear model predicts that orientation tuning curves simply scale with contrast, i.e. that orientation tuning is contrast-invariant. However, some properties of simple cell responses are non-linear, and the normalization model posits that an additional cortical circuit – a normalizing circuit – is needed to ‘correct’ the linear input and explain these response properties. These non-linear response properties include: an advance with increasing contrast in the phase of response to sinusoidal gratings (a linear model would show the same phase of response at all contrasts); an emergence with increasing contrast of responses to higher temporal frequencies that evoke little or no response at low contrast (in a linear model, temporal frequency curves would simply scale with contrast); saturation of cortical responses at contrasts lower than those at which the LGN inputs saturate; and cross-orientation inhibition, the reduction of response to a preferred-orientation stimulus by simultaneous presentation of an orthogonal stimulus which by itself evokes no response (in a linear model, responses to the two stimuli would add).

We propose a different viewpoint from that of the normalization model. It is not the case that a simple cell receives linear input that must be corrected to account for non-linearities. Rather, a simple cell receives non-linear input and processes it through non-linear machinery, and what is needed is an explanation of how the responses of the simple cell nonetheless come to appear linear. The most obvious non-linearity in the input to a simple cell is caused by the rectification of LGN responses – the fact that LGN responses cannot be decreased below zero. We saw above that this rectification can cause a stimulus orthogonal to a cell’s preferred orientation to evoke a strong LGN input to the cell. There are a multitude of other non-linearities in the circuit, including frequency-dependent synaptic depression in LGN and cortical synapses, spike-rate adaptation currents in cortical cells, stimulus-induced conductance changes in cortical cells, and the cortical spike threshold. We argue that the approximately linear response of simple cells is achieved, in spite of these non-linearities, by the dominant opponent inhibition, which filters out the input caused by LGN rectification while leaving the linear component of LGN input [a similar explanation of

simple cell linearity, but using phase-non-specific feedforward inhibition rather than antiphase feedforward inhibition, is found in the model of Wielaard *et al.* (Wielaard *et al.*, 2001)]. The remaining non-linearities in the input and the circuit can explain the non-linear aspects of simple cell response – no separate normalizing circuit is needed, instead the non-linearities are present from the outset. We showed that our model circuit can indeed explain all of the non-linear response properties described above (Kayser *et al.*, 2001; Lauritzen *et al.*, 2001).

In sum, this simple model circuit promises to provide a unified account of classical receptive field properties of simple cells, although many response properties such as direction selectivity, end stopping and beyond-the-classical-receptive-field effects remain to be addressed.

Experimental Results that Functionally Characterize Inhibitory Neurons

Our model predicts that inhibitory interneurons in layer 4 should provide strong feedforward inhibition that has orientation tuning like that of a simple cell’s thalamic inputs. In particular, such cells should respond in a contrast-dependent manner to all orientations.

In a recent study using intracellular recording *in vivo*, roughly ten inhibitory neurons were recorded in layer 4 of cat V1, and these were found to be of two types: simple cells showing good orientation tuning (studied with moving bars at one contrast), and complex cells – cells responding either to light or dark throughout their receptive field – showing roughly equal responses to all orientations (Hirsch *et al.*, 2000). This raises the possibility that the tuning attributed in the antiphase model to simple inhibitory cells – response to all orientations, though tuned for the preferred – might actually be achieved by the combination of two inhibitory populations. The simple cells would provide opponent inhibition, but would not respond to orientations far from the preferred. The complex cells would provide the broadly tuned inhibition that prevents simple cells, both excitatory and inhibitory, from responding to orientations far from the preferred.

Numerous studies of ‘suspected inhibitory neurons’ (SINs) in rodent and rabbit cortex also suggest that layer 4 neurons receive strong and broadly tuned feedforward inhibition (Miller *et al.*, 2001; Pinto *et al.*, 2002; Swadlow, 2002).

In slice recordings from layer 4 of rodent somatosensory cortex, two biophysical types of interneurons were found: fast-spiking (FS) neurons receive strong feedforward input from thalamus, while low-threshold-spiking (LTS) neurons receive no feedforward input and so provide only feedback inhibition (Gibson *et al.*, 1999) [however, Porter *et al.* found that both types of interneurons can provide feedforward inhibition (Porter *et al.*, 2001)]. Furthermore there is extensive gap-junction coupling within each type, but not between the two types. It is tempting to guess that these two biophysical types correspond to the two functional types, simple and complex, found in layer 4 of V1, but this appears not to be the case (J.A. Hirsch, private communication.) Our model interneurons had parameters corresponding to FS neurons, and lacked gap-junction coupling. The roles of LTS interneurons, of purely feedback inhibition and of gap-junction coupling in layer 4 functional responses remain to be explored. The high rate of gap-junction coupling suggests that these cells should have rather non-specific functional responses, consistent with the complex inhibitory cells seen in layer 4 (though not the simple inhibitory cells) and consistent with properties reported for SINs.

Experimental Results that Suggest Feedforward Processing in Layer 4

The picture we have presented suggests that the response properties of layer 4 simple cells should be dominantly determined by feedforward processing – i.e. by the combination of the LGN inputs and LGN-driven inhibition. A series of intra-cellular recording experiments from David Ferster's laboratory in recent years have provided compelling evidence that the processing underlying simple cell orientation selectivity is indeed dominantly feedforward.

Ferster *et al.* attempted to directly compare the orientation tuning of the thalamic input to that induced by the full cortical circuit (Ferster *et al.*, 1996). To achieve this, they compared the tuning of the voltage responses of simple cells in two conditions: the normal condition, with the full cortical circuit intact; and after cortical cooling, which blocked cortical spiking, leaving transmission along and vesicle release from thalamic axons intact (though slowed and diminished). By eliminating cortical spiking, the cooling should allow isolation of the voltage responses induced by the thalamic input alone. The temporal modulations of voltage in simple cells induced by high-contrast drifting sinusoidal gratings, though smaller in the cooled condition, showed identical orientation tuning in the control and cooled conditions, suggesting that the tuning of the full cortical circuit followed that of the thalamic inputs. This result is accounted for by the model of Troyer *et al.* (Troyer *et al.*, 1998): the voltage modulations follow the LGN inputs, while the inhibition and threshold sharpen spiking tuning relative to voltage tuning [a sharpening observed experimentally (Carandini and Ferster, 2000; Volgushev *et al.*, 2000)]. Note that the tuning of the voltage modulations induced by the thalamic inputs should depend on the stimulus; sinusoidal gratings of higher spatial frequencies should evoke narrower thalamic orientation tuning than gratings of lower spatial frequencies (Troyer *et al.*, 1998) [and indeed, the voltage modulations induced by the full circuit show narrower orientation tuning with increasing grating spatial frequency (Lampl *et al.*, 2001)]. Thus, a match of thalamic and full-circuit tuning for a particular choice of stimulus suggests that the full-circuit tuning follows the thalamic more generally.

The cooling did not entirely eradicate cortical spiking; cells in layer 6, farthest from the cooling plate, showed perhaps 15% of their normal spiking responses. Ferster's group therefore assayed the same question by an independent technique, using a shock to the cortex to induce hyperpolarization and thus suppress cortical spiking for a period of >100 ms, and examining the tuning of voltage responses to flashed gratings during the period of suppression (Chung and Ferster, 1998). Again, voltage responses showed the same orientation tuning in control and suppressed conditions. This experiment showed that transient responses, like the steady-state responses observed in the cooling experiment, appear to be largely determined by feedforward processing.

An argument against a feedforward computation of orientation tuning has been that orientation tuning width is narrower than would be expected from a semi-linear prediction based on the arrangement of the cell's ON and OFF subregions (Gardner *et al.*, 1999). (We use 'semi-linear' to refer to a prediction that may take into account rectification of neuronal responses.) However, the antiphase model predicts that inhibition and threshold sharpen spiking tuning relative to voltage tuning; it is voltage tuning that would be expected to follow a semi-linear prediction. Ferster's group tested this by mapping the cell's receptive field intracellularly with flashed spots, and found that the

orientation tuning of the voltage response to a drifting sinusoidal luminance grating could be well predicted from the receptive field map (Lampl *et al.*, 2001). For two cases in which two spatial frequencies were tested on the same cell, both the broader voltage tuning for the lower-frequency grating and the narrower voltage tuning for the higher-frequency grating were correctly predicted. However, the semi-linear prediction tended to predict a greater response orthogonal to the preferred orientation than is actually observed, in agreement with earlier results (Volgushev *et al.*, 1996).

Finally, Ferster's group examined the intracellular basis of contrast-invariant orientation tuning (Anderson *et al.*, 2000b). They examined two aspects of the voltage response to drifting sinusoidal gratings of various orientations and contrasts: the amplitude of the temporal modulation of voltage induced by the grating ('voltage modulation'); and the mean depolarization induced by the stimulus ('voltage mean'). They found that the voltage modulation and the voltage mean each showed similar orientation tuning that simply scaled with changes in stimulus contrast. In combination with their previous finding that the orientation tuning of the voltage modulation at high contrast followed the tuning of the thalamic inputs (Ferster *et al.*, 1996), this suggests that the voltage orientation tuning across contrasts follows the tuning of the thalamic inputs.

These results, while not a necessary consequence of the antiphase inhibition model, are consistent with it. The model predicts that the voltage modulation will have orientation tuning that scales with contrast, as observed, but is more agnostic about the tuning of the voltage mean. The model predicts that the mean LGN input to a simple cell should be untuned for orientation, because a grating stimulus raises LGN firing rates by an amount that depends on contrast but is independent of orientation. If not opposed by inhibition, this would lead to a mean voltage response that is depolarizing at all orientations. However, the dominant feedforward inhibition in the model adds to the direct LGN input to produce a total mean feedforward input that is inhibitory, meaning that it has a subthreshold reversal potential. In response to a null stimulus (a stimulus oriented orthogonal to the preferred), one should see only this mean feedforward input (because cortical cells are not driven to spike, so there is no local feedback input, only feedforward input). The voltage response induced by this input depends on the location of its reversal potential relative to rest. Empirically, little voltage change was observed in response to a null stimulus, suggesting that the mean feedforward input has a reversal potential near rest. Since rest is near the inhibitory reversal potential, this is consistent with this mean input being inhibition dominated. In sum, the lack of a voltage response to a null-oriented stimulus, despite the increase in LGN firing rates evoked by that stimulus, suggests the presence of dominant feedforward inhibition, as we have posited. However, a further complication is that short-term synaptic depression of thalamocortical synapses can eliminate a significant fraction of the feedforward mean input at the temporal frequencies studied, but at higher temporal frequencies (e.g. 8 Hz) the inhibitory mean should be strongly present, and so should be visible as a conductance change in response to a null stimulus even if no voltage change is apparent (Krukowski, 2000); this remains to be tested. Although the mean feedforward input is predicted to be untuned for orientation, at least two effects could lead the mean voltage response to have orientation tuning like that of the voltage modulation. First, voltage can modulate up much further than it can modulate down, because the excitatory reversal potential is much further from rest than is the inhibitory reversal potential; as a result,

voltage modulation will induce a mean depolarization with an orientation tuning identical to that of the modulation. Second, spiking tuning follows (but is narrower than) the tuning of the voltage modulations, and recurrent excitatory connections will contribute a mean depolarization whenever spiking occurs.

Anderson *et al.* also found that voltage noise – the trial-by-trial fluctuations about the average stimulus-induced voltage response for a given stimulus – was critical to turning the contrast-invariant voltage tuning that they observed into contrast-invariant spiking tuning (Anderson *et al.*, 2000b). A simple picture of this effect (Hansel and van Vreeswijk, 2002; Miller and Troyer, 2002) is given by assuming that the average spiking rate R is some instantaneous function $R(V)$ of the average voltage V . In the absence of noise, this function would be linear above some threshold voltage and zero below the threshold. Such a linear-threshold function would convert the contrast-invariant voltage tuning into spiking tuning that broadens with contrast, because at higher contrasts more orientations would produce suprathreshold voltages. Noise smooths this linear threshold function, because a subthreshold average voltage will sometimes fluctuate above threshold. In particular, noise converts the linear threshold function into a power law $R \propto V^n$ over some range of voltages, and a power law converts contrast-invariant voltage tuning into contrast-invariant spiking tuning, with tuning sharpened by a factor of \sqrt{n} . (Contrast-invariant voltage tuning means that the voltage response factors into a function of orientation θ times a function of contrast C , $V = f_1(\theta)f_2(C)$. Raising this to a power n preserves the factoring, $R = V^n = [f_1(\theta)]^n [f_2(C)]^n$, and thus preserves contrast invariance. Orientation tuning curves are reasonably described by Gaussians, and raising a Gaussian to a power n reduces the standard deviation of the Gaussian by a factor \sqrt{n}). Our model of contrast-invariant orientation tuning did not rely on such noise smoothing except at very low contrasts (Troyer *et al.*, 2002), and so needs some revision in light of Anderson *et al.*'s finding that the full range of contrasts lies in the noise-smoothed regime. However, the noise smoothing achieves contrast-invariant tuning only if the voltage shows contrast-invariant tuning, and as discussed above this in turn requires dominant feedforward inhibition to suppress mean voltage responses to non-preferred orientations. Thus we expect the basic ideas of our model, including the role of dominant feedforward inhibition, to remain intact (also supported by preliminary results: S.E. Palmer and K.D. Miller, unpublished).

The results presented thus far have focused on orientation tuning. Another property of simple cells in layer 4 is direction selectivity: preference for stimulus movement in one of the two opposite directions orthogonal to the preferred orientation. Physiological evidence is suggestive that this property might also be understood in layer 4 from the structure of the feedforward input received by a cell along with the effects of the spike threshold nonlinearity. Voltage responses to moving stimuli can be predicted as a simple linear sum of inputs; stimuli moving in the two directions can be decomposed into a sum of stationary stimuli, and the voltage responses to the moving stimuli can correspondingly be predicted from a sum of the voltage responses to stationary stimuli (Jagadeesh *et al.*, 1997). Furthermore, the voltage responses could be understood as arising from sums of only two input components, with properties that closely resemble those of two temporal types of LGN inputs: non-lagged cells and lagged cells (Jagadeesh *et al.*, 1997). Just as adjacent rows of ON- and OFF-center inputs can explain a simple cell's spatial response profile, an appropriate spatial mix of lagged and non-lagged input can produce cells whose space-time receptive fields show preference for one direction. Studies of temporal

response profiles of simple cell receptive fields found timing corresponding to lagged-type input only in cells of layer 4B (Saul and Humphrey, 1992), and correspondingly cells in layer 4B show the strongest direction preference in their linear space-time receptive fields (Murthy *et al.*, 1998). Strobe-rearing greatly reduces direction selectivity in cat V1 cells (Humphrey and Saul, 1998), and correspondingly eliminates the convergence of non-lagged-like and lagged-like temporal responses in individual simple cells (Humphrey *et al.*, 1998). Studies of adaptation suggest that direction-selective simple cells receive inhibition from other simple cells preferring the same direction but with different space-time phases (Saul, 1999), which suggests a generalization to space-time receptive fields of the spatial antiphase inhibition posited thus far.

Experimental Results that Argue for Other Contributions to Orientation Tuning

A number of observations are suggestive of a role of recurrent connections, cross-orientation inhibition and/or phase-nonspecific inhibition in generating orientation selectivity [reviewed by Sompolinsky and Shapley (Sompolinsky and Shapley, 1997) and by Ferster and Miller (Ferster and Miller, 2000)]. When recording from a site preferring one orientation, GABA-induced inactivation of a site preferring the orthogonal orientation 350–700 μm away leads to a broadening of orientation tuning at the recorded site, and this was true in particular at many recording sites in layer 4 (Crook *et al.*, 1996, 1997). Furthermore, anatomical studies confirm the existence of inhibitory neurons in the vicinity of inactivation sites that project to the vicinity of the corresponding recording site (Crook *et al.*, 2000). Anatomical labeling combined with optical imaging shows that sites in layer 4 in cat area 18 receive connections from proximal sites (roughly, within 500 μm) that are strongly biased towards similar orientation preferences, as expected from the antiphase model, but long-range connections over distances up to 2–3 mm are fairly uniformly distributed across orientations (Yousef *et al.*, 1999). Adaptation to an orientation to one side of the preferred orientation can induce a shift in orientation tuning toward, and an increase in response to, orientations to the opposite side of the preferred orientation, and this effect shows little dependence on cortical depth and hence appears likely to hold in layer 4 (Dragoi *et al.*, 2000). Intracellular studies of transient responses to a flashed bar of the preferred orientation show an initial conductance increase with sub- or peri-threshold reversal potential, before the response becomes either excitatory or inhibitory (depending on whether the bar was flashed over an appropriate or inappropriate subregion) (Borg-Graham *et al.*, 1998); however, cells were not identified by layer, so the applicability to layer 4 is uncertain. Finally, as already mentioned, a linear model of voltage responses based on responses to flashed spots predicts larger voltage responses to the null orientation than are actually observed (Volgushev *et al.*, 1996; Lampl *et al.*, 2001).

Studies of the dynamics of orientation tuning in response to flashed stimuli have also been argued to support a role for feedback, but at least some of these results may instead be compatible with the results of feedforward inhibition. A recent intracellular study divided the orientation tuning curve of voltage responses into a tuned component and an untuned component, where the latter is a constant voltage response across orientations. The study found no statistically significant changes with time after stimulus onset in the width of the tuned component, but in many cells the untuned component grew more negative over time (Gillespie *et al.*, 2001). This increasing

negativity of the untuned component is expected if feedforward inhibition follows feedforward excitation. The overall voltage tuning curve – tuned plus untuned component – would narrow with time, as reported for some cells in another study (Volgushev *et al.*, 1995). An extracellular study in monkey reported that perhaps half of cells studied showed changes in the tuned response component with post-stimulus time, but these effects were not seen in thalamic-recipient portions of layer 4 (Ringach *et al.*, 1997). This study used stimuli several times larger than the classical receptive field, so surround suppression effects may have played a role.

Alternative Models

A major alternative model of V1 circuitry posits that strong, localized feedback excitation and more widespread feedback inhibition create orientation tuning that is an intrinsic property of cortex, independent of the tuning of the thalamic input (Ben-Yishai *et al.*, 1995; Somers *et al.*, 1995). This yields contrast-invariant orientation tuning – the width of orientation tuning is a cortical property, independent of any stimulus property, including stimulus contrast. In this model, factors that change the tuning of a cell's thalamic input are predicted to have no effect on its orientation tuning. This is contradicted by Ferster's findings that a cell's voltage orientation tuning follows the voltage tuning of its thalamic inputs and that it has the tuning predicted from its spatial receptive field, including narrower tuning for higher spatial frequency gratings. It is also contradicted by findings that spiking orientation tuning narrows with increasing spatial frequency of a grating stimulus [reviewed by Troyer *et al.* (Troyer *et al.*, 1998)] and with increasing length of a bar stimulus (Orban, 1991), in both cases as predicted if orientation tuning follows the tuning of the thalamic inputs.

McLaughlin *et al.* (McLaughlin *et al.*, 2000) and Wiesel *et al.* (Wiesel *et al.*, 2001) have proposed a model of responses in layer 4C α of monkey V1. This model also relies on strong feedforward inhibition to cancel the non-linear component of the LGN input, but it assumes the inhibition has no phase specificity, coming equally from cells of all preferred phases. This is motivated in part by experiments reporting transient phase-non-specific inhibitory responses to flashed stimuli (Borg-Graham *et al.*, 1998), in contrast to the phase-specific opponent arrangement seen by others (Ferster, 1988; Hirsch *et al.*, 1998). Phase-non-specific feedforward inhibition can also solve the problem posed by the thalamic inputs, by setting a contrast-dependent threshold for response – a high-contrast stimulus orthogonal to the preferred would evoke stronger inhibition than a low-contrast preferred stimulus, allowing the cell to respond to the latter and not the former. However, the model of McLaughlin *et al.* and Wiesel *et al.* actually operates in a parameter regime in which the inhibition is not strong enough to fully cancel the non-linear component of the LGN input, so that many cells respond to stimuli of all orientations. In this model, cells are assumed to receive input from all other cells within a given distance, and as a result a cell's orientation tuning depends on its location in the orientation map. Cells located in 'linear' regions of the map, where nearby cells all have similar preferred orientations, receive inhibition only from cells of similar preferred orientation, and these cells respond to all orientations although showing a tuning peak at the preferred. Cells located near orientation 'pinwheels', points where cells of all preferred orientations converge, receive inhibition from cells of all preferred orientations and hence show sharp orientation tuning. The prediction that cells in linear regions show broader orientation tuning than cells in pinwheels seems not to be

correct in cats (Ruthazer *et al.*, 1996; Maldonado *et al.*, 1997), but the case in monkeys is not known.

A Developmental Model of Cortical Layer 4 and Columnar Invariance

We have recently shown (Kayser and Miller, 2002) that the model functional circuit of Figure 3, including both the pattern of LGN inputs to simple cells and the intracortical connectivity between the excitatory and inhibitory simple cells, will all develop under simple Hebb-like rules of activity-instructed synaptic modification. The only requirement is that LGN input activities during development should show a simple statistical structure that is likely to arise in spontaneous activity driven by quantal events in photoreceptors (Mastrorade, 1989). In addition, one must assume that inhibition is stronger than excitation in order for the resulting circuit to show the functional response properties of simple cells.

This suggests the more general hypothesis that layer 4 of any piece of cortex may develop through simple Hebb-like rules, guided simply by the statistical structure of its inputs' activities. This and the dominance of inhibition leads naturally to opponent inhibition, in the generalized sense in which we defined it above: a cell becomes selective for a preferred pattern of inputs, and also becomes strongly inhibited by the input pattern that is most anticorrelated with the preferred pattern, which we can call the 'opposite' pattern. As we argued above, this solves the problem posed by the rectification of the thalamic input and endows layer 4 with magnitude-invariant form recognition: it enables a cell to respond to its preferred stimulus even at low magnitude, and not to respond to a non-preferred stimulus even at high magnitude, even though the latter stimulus may provide a cell with as much thalamic input as the former stimulus. To this basic idea must be added a role for non-specific, broadly tuned inhibitory cells, such as the complex cells reported by Hirsch *et al.* (2000). Whether opponent inhibition is indeed an idea that generalizes across cortical areas, and how the opponent inhibition and the non-specific inhibition relate to one another, remain to be worked out.

The hypothesis that layer 4 develops through Hebb-like rules and develops opponent inhibition leads to a more general hypothesis about cortical columnar organization (Kayser and Miller, 2002). Which cortical properties should show columnar invariance, i.e. an invariance across the cortical layers at a given tangential position? Such properties should in particular be locally invariant in layer 4. Given Hebbian development resulting in opponent inhibition, it turns out that if a given stimulus pattern is represented in a local region of layer 4, the opposite pattern will also be represented in the same local region. That is, a local region will include cells that represent stimulus pairs that are as dissimilar as possible, where 'dissimilarity' is measured by anticorrelation of the input patterns evoked by the stimuli. This contrasts with the more common idea that cells in a column all represent a similar set of response properties. As a result, the only properties that can be locally invariant in layer 4, and hence that are candidates for being invariant across a column, are properties that are *shared* by a stimulus and its opposite. Properties that differ between a stimulus and its opposite cannot show columnar invariance by this reasoning.

For simple cells in cat V1 layer 4, a preferred stimulus is a pattern of light and dark bars matching the cell's subregions, while its opposite is a pattern of the same orientation but with opposite phase – light in place of dark and vice versa. Thus, for V1, the prediction is that orientation, which is shared by an input pattern and its opposite, should show local invariance in

layer 4, while phase, which differs between an input pattern and its opposite, should not. It is well known that preferred orientation shows columnar invariance in cat V1, and it appears that preferred phase does not (DeAngelis *et al.*, 1999). It remains to be seen whether this hypothesis can account for the properties that show columnar invariance in other cortical areas.

Conclusion: Understanding Layer 4

We have described a simple model of layer 4 of cat V1, based on the ideas of dominant feedforward inhibition and opponent inhibition. We have described a number of experimental results that fit nicely within this framework, and others that do not. The complexity of the biological circuit remains greater than any single simple model can fully capture. But, overall, the picture of strong feedforward antiphase inhibition supplementing the tuning of the thalamic inputs can explain a large body of diverse data in cat V1 layer 4. We suggest that dominant feedforward inhibition and opponent inhibition may be general features of the circuitry of layer 4 of cerebral cortex (Miller *et al.*, 2001).

Notes

This article was adapted and expanded from K.D. Miller, D.J. Simons and D.J. Pinto (2001) Processing in layer 4 of the neocortical circuit: new insights from visual and somatosensory cortex. *Curr. Opin. Neurobiol.* 11:488–497, copyright © 2001 by and with permission from Elsevier Science, and with permission from Dan Simons and David Pinto. I thank Anton Krukowski and Todd Troyer for helpful comments on the manuscript and Anton Krukowski for making several of the figures. I was supported by R01-EY11001 from the NEI.

Address correspondence to: Kenneth Miller, Department of Physiology, University of California at San Francisco, 513 Parnassus, San Francisco, CA 94143–0444, USA. Email: ken@phy.ucsf.edu.

References

- Anderson JS, Carandini M, Ferster D (2000a) Orientation tuning of input conductance, excitation, and inhibition in cat primary visual cortex. *J Neurophysiol* 84:909–926.
- Anderson JS, Lampl I, Gillespie D, Ferster D (2000b) The contribution of noise to contrast invariance of orientation tuning in cat visual cortex. *Science* 290:1968–1972.
- Ben-Yishai R, Bar-Or RL, Sompolinsky H (1995) Theory of orientation tuning in visual cortex. *Proc Natl Acad Sci USA* 92:3844–3848.
- Borg-Graham LJ, Monier C, Fregnac Y (1998) Visual input evokes transient and strong shunting inhibition in visual cortical neurons. *Nature* 393:369–373.
- Bullier J, Henry GH (1979) Laminar distribution of first-order neurons and afferent terminals in cat striate cortex. *J Neurophysiol* 42:1271–1281.
- Callaway E (1998) Local circuits in primary visual cortex of the macaque monkey. *Annu Rev Neurosci* 21:47–74.
- Carandini M, Ferster D (2000) Membrane potential and firing rate in cat primary visual cortex. *J Neurosci* 20:470–484.
- Carandini M, Heeger DJ, Movshon JA (1999) Linearity and gain control in V1 simple cells. In: *Cerebral cortex*, Vol. 13: Models of cortical circuits (Ulinski P, Jones EG, Peters A, eds), pp. 401–443. New York: Kluwer Academic/Plenum Press.
- Chung S, Ferster D (1998) Strength and orientation tuning of the thalamic input to simple cells revealed by electrically evoked cortical suppression. *Neuron* 20:1177–1189.
- Crair MC, Malenka RC (1995) A critical period for long-term potentiation at thalamocortical synapses. *Nature* 375:325–328.
- Crook JM, Kisvárdy ZF, Eysel UT (1996) GABA-induced inactivation of functionally characterized sites in cat visual cortex (area 18): effects on direction selectivity. *J Neurophysiol* 75:2071–2088.
- Crook JM, Kisvárdy ZF, Eysel UT (1997) GABA-induced inactivation of functionally characterized sites in cat striate cortex: effects on orientation and direction selectivity. *Vis Neurosci* 14:141–158.
- Crook JM, Kisvárdy ZF, Eysel UT (2000) Evidence for a contribution of lateral inhibition to orientation tuning and direction selectivity in cat visual cortex: reversible inactivation of functionally characterized sites combined with neuroanatomical tracing techniques. *Eur J Neurosci* 10:2056–2075.
- DeAngelis GC, Ghose GM, Ohzawa I, Freeman RD (1999) Functional micro-organization of primary visual cortex: receptive field analysis of nearby neurons. *J Neurosci* 19:4046–4064.
- Dragoi V, Sharma J, Sur M (2000) Adaptation-induced plasticity of orientation tuning in adult visual cortex. *Neuron* 28:287–298.
- Ferster D (1986) Orientation selectivity of synaptic potentials in neurons of cat primary visual cortex. *J Neurosci* 6:1284–1301.
- Ferster D (1988) Spatially opponent excitation and inhibition in simple cells of the cat visual cortex. *J Neurosci* 8:1172–1180.
- Ferster D, Jagadeesh B (1992) EPSP–IPSP interactions in cat visual cortex studied with *in vivo* whole-cell patch recording. *J Neurosci* 12:1262–1274.
- Ferster D, Miller KD (2000) Neural mechanisms of orientation selectivity in the visual cortex. *Annu Rev Neurosci* 23:441–471.
- Ferster D, Chung S, Wheat H (1996) Orientation selectivity of thalamic input to simple cells of cat visual cortex. *Nature* 380:249–252.
- Gardner JL, Anzai A, Ohzawa I, Freeman RD (1999) Linear and nonlinear contributions to orientation tuning of simple cells in the cat's striate cortex. *Vis Neurosci* 16:1115–1121.
- Gibson JR, Beierlein M, Connors BW (1999) Two networks of electrically coupled inhibitory neurons in neocortex. *Nature* 462:75–79.
- Gilbert CD (1977) Laminar differences in receptive field properties of cells in cat primary visual cortex. *J Physiol* 268:391–421.
- Gillespie DC, Lampl I, Anderson JS, Ferster D (2001) Dynamics of the orientation-tuned membrane potential response in cat primary visual cortex. *Nat. Neurosci.* 4:1014–1019.
- Hansel D, van Vreeswijk C (2002) How noise contributes to contrast invariance of orientation tuning in cat visual cortex. *J Neurosci* 22:5118–5128.
- Hirsch JA, Alonso J-M., Reid RC, Martinez L (1998) Synaptic integration in striate cortical simple cells. *J Neurosci* 18:9517–9528.
- Hirsch JA, Alonso J-M., Pillai C, Pierre C (2000) Simple and complex inhibitory cells in layer 4 of cat visual cortex. *Soc Neurosci Abstr* 26:1083.
- Hubel DH and Wiesel TN (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J Physiol* 160:106–154.
- Humphrey AL, Saul AB (1998) Strobe rearing reduces direction selectivity in area 17 by altering spatiotemporal receptive-field structure. *J Neurophysiol* 80:2991–3004.
- Humphrey AL, Saul AB, Feidler JC (1998) Strobe rearing prevents the convergence of inputs with different response timings onto area 17 simple cells. *J Neurophysiol* 80:3005–3020.
- Jagadeesh B, Wheat HS, Kontsevich LL, Tyler CW, Ferster D (1997) Direction selectivity of synaptic potentials in simple cells of the cat visual cortex. *J Neurophysiol* 78:2772–2789.
- Kayser AS, Miller KD (2002) Opponent inhibition: a developmental model of layer 4 of the neocortical circuit. *Neuron* 33:131–142.
- Kayser AS, Priebe NJ, Miller KD (2001) Contrast-dependent nonlinearities arise locally in a model of contrast-invariant orientation tuning. *J Neurophysiol* 85:2130–2149.
- Krukowski AE (2000) A model of cat primary visual cortex and its thalamic input. PhD thesis, University of California, San Francisco.
- Krukowski AE, Miller KD (2001) Thalamocortical NMDA conductances and intracortical inhibition can explain cortical temporal tuning. *Nat. Neurosci.* 4:424–430.
- Lampl I, Anderson J, Gillespie D, Ferster D (2001) Prediction of orientation selectivity from receptive field architecture in simple cells of cat visual cortex. *Neuron* 30:263–274.
- Lauritzen TZ, Krukowski AE, Miller KD (2001) Local correlation-based circuitry can account for responses to multi-grating stimuli in a model of cat V1. *J Neurophysiol* 86:1803–1815.
- Maldonado PE, Godecke I, Gray CM, Bonhoeffer T (1997) Orientation selectivity in pinwheel centers in cat striate cortex. *Science* 276:1551–1555.
- Martinez LM, Alonso J-M., Reid RC, Hirsch JA (1998) The role of excitation and inhibition in the orientation tuning of simple and complex cells in cat striate cortex. *Soc Neurosci Abstr* 24:766.
- Mastrorade DN (1989) Correlated firing of retinal ganglion cells. *Trends Neurosci* 12:75–80.
- McLaughlin D, Shapley R, Shelley M, Wielaard DJ (2000) A neuronal network model of macaque primary visual cortex (V1): orientation

- selectivity and dynamics in the input layer 4Cα. *Proc Natl Acad Sci USA* 97:8087–8092.
- Miller KD, Troyer TW (2002) Neural noise can explain expansive, power-law nonlinearities in neural response functions. *J Neurophysiol* 87:653–659.
- Miller KD, Simons DJ, Pinto DJ (2001) Processing in layer 4 of the neocortical circuit: new insights from visual and somatosensory cortex. *Curr Opin Neurobiol* 11:488–497.
- Murthy A, Humphrey AL, Saul AB, Feidler JC (1998) Laminar differences in the spatiotemporal structure of simple cell receptive fields in cat area 17. *Vis Neurosci* 15:239–256.
- Orban GA (1991) Quantitative electrophysiology of visual cortical neurones. In: *Vision and visual dysfunction. The neural basis of visual function* (Cronly-Dillon J, Leventhal AG, eds), pp. 173–222. London: Macmillan.
- Porter JT, Johnson CK, Agmon A (2001) Diverse types of interneurons generate thalamus-evoked feedforward inhibition in the mouse barrel cortex. *J Neurosci* 21:2699–2710.
- Reid RC, Alonso JM (1995) Specificity of monosynaptic connections from thalamus to visual cortex. *Nature* 378:281–284.
- Ringach DL, Hawken MJ, Shapley R (1997) Dynamics of orientation tuning in macaque primary visual cortex. *Nature* 387:281–284.
- Ruthazer ES, Crair MC, Gillespie D, Stryker MP (1996) Orientation tuning of single cells at 'pinwheel' orientation singularities in kitten visual cortex. *Soc Neurosci Abstr* 22:1996.
- Saul AB (1999) Visual cortical simple cells: who inhibits whom. *Vis Neurosci* 16:667–673.
- Saul AB, Humphrey AL (1992) 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.
- Sclar G, Freeman RD (1982) Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. *Exp Brain Res* 46:457–461.
- Skottun BC, Bradley A, Sclar G, Ohzawa I, Freeman RD (1987) The effects of contrast on visual orientation and spatial frequency discrimination: a comparison of single cells and behavior. *J Neurophysiol* 57:773–786.
- Somers D, Nelson SB, Sur M (1995) An emergent model of orientation selectivity in cat visual cortical simple cells. *J Neurosci* 15:5448–5465.
- Sompolinsky H, Shapley R (1997) New perspectives on the mechanisms for orientation selectivity. *Curr Opin Neurobiol* 7:514–522.
- Tanaka K (1983) Cross-correlation analysis of geniculostriate neuronal relationships in cats. *J Neurophysiol* 49:1303–1318.
- Troyer TW, Krukowski AE, Priebe NJ, Miller KD (1998) Contrast-invariant orientation tuning in cat visual cortex: feedforward tuning and correlation-based intracortical connectivity. *J Neurosci* 18:5908–5927.
- Troyer TW, Krukowski AE, and Miller KD (2002) LGN input to simple cells and contrast-invariant orientation tuning: an analysis. *J Neurophysiol* 87: (in press).
- Volgushev M, Vidyasagar TR, Pei X (1995) Dynamics of the orientation tuning of postsynaptic potentials in the cat visual cortex. *Vis Neurosci* 12:621–628.
- Volgushev M, Vidyasagar TR, Pei X (1996) A linear model fails to predict orientation selectivity of cells in the cat visual cortex. *J Physiol* 496:597–606.
- Volgushev M, Pernberg J, Eysel U (2000) Comparison of the selectivity of postsynaptic potentials and spike responses in cat visual cortex. *Eur J Neurosci* 12:257–263.
- Wieland DJ, Shelley M, McLaughlin D, Shapley R (2001) How simple cells are made in a nonlinear network model of the visual cortex. *J Neurosci* 21:5203–5211.
- Yousef T, Bonhoeffer T, Kim DS, Eysel UT, Toth E, Kisvarday ZF (1999) Orientation topography of layer 4 lateral networks revealed by optical imaging in cat visual cortex (area 18). *Eur J Neurosci* 11:4291–4308.