

Mechanisms of Neuronal Computation in Mammalian Visual Cortex

Nicholas J. Priebe¹ and David Ferster^{2,*}

¹Section of Neurobiology, Center for Perceptual Systems, University of Texas at Austin, 2401 Speedway, Austin, TX 78705, USA

²Department of Neurobiology, Northwestern University, 2205 Tech Drive, Evanston, IL 60208, USA

*Correspondence: ferster@northwestern.edu

<http://dx.doi.org/10.1016/j.neuron.2012.06.011>

Orientation selectivity in the primary visual cortex (V1) is a receptive field property that is at once simple enough to make it amenable to experimental and theoretical approaches and yet complex enough to represent a significant transformation in the representation of the visual image. As a result, V1 has become an area of choice for studying cortical computation and its underlying mechanisms. Here we consider the receptive field properties of the simple cells in cat V1—the cells that receive direct input from thalamic relay cells—and explore how these properties, many of which are highly nonlinear, arise. We have found that many receptive field properties of V1 simple cells fall directly out of Hubel and Wiesel's feedforward model when the model incorporates realistic neuronal and synaptic mechanisms, including threshold, synaptic depression, response variability, and the membrane time constant.

Introduction

In most sensory areas of the brain, the local circuit transforms its input to generate a novel representation of the external world. The sensory receptive fields that are produced represent the visible result of a neuronal computation. Sensory transformations can be subtle, as in the case of the lateral geniculate nucleus (LGN), in which the center-surround structure of the input from retinal ganglion cells is largely preserved in the output from the geniculate relay cells (Hubel and Wiesel, 1962). Or transformations can be dramatic, as in the case of the retina, in which the pixel-like representation of the visual image by retinal photoreceptors is transformed into the center-surround receptive fields of retinal ganglion cells (Kuffler, 1953).

The quintessential example of a complex sensory computation is the one performed by the primary visual cortex (V1). There, selectivity for a range of image properties emerges from relatively unselective inputs. Simple cells in layer 4 of V1, unlike their LGN inputs, are sensitive to contour length, direction of motion, size, depth, and most famously, orientation (Hubel and Wiesel, 1962). As striking as the cortical transformation is, the resulting changes in the visual representation can be measured experimentally in quantitative detail and described with mathematical precision. Few areas outside the visual cortex have been described so comprehensively and on so many levels, from basic neuronal response properties, to anatomical connectivity, to functional architecture. Since the cerebral cortex is thought to be the primary locus of high-level processes such as perception, cognition, language, and decision making, it is no wonder that the visual cortex has become the most widely studied proxy for computation in the cerebral cortex. Not only does it lend itself to questions of how its sensory transformation contributes to visual perception (Gilbert and Li, 2012), but the emergence of orientation selectivity is the model system for studying how cortical circuitry performs a neuronal computation.

Models of Orientation Selectivity

Few computational models have the elegance, simplicity, and longevity of Hubel and Wiesel's proposal for how the cortical circuit generates orientation selectivity. In their 1962 paper, they proposed that a simple cell becomes orientation selective by virtue of the excitation it receives from LGN relay cells whose receptive fields are aligned parallel to the simple cell's preferred orientation (Figure 1A). The total excitatory input integrated over an oriented stimulus that moves across the receptive field will be nearly identical at all orientations, because the geniculate inputs respond identically at each stimulus orientation. What varies instead is their relative timing, which will be nearly simultaneous for the preferred orientation but spread out in time for the non-preferred orientations (Figure 1B). Even for nonpreferred stimuli, however, the total excitatory input is nonzero. A threshold is therefore required to render the spike output of the cell perfectly orientation selective, with no response at the orthogonal orientation (Figure 1B, bottom).

One feature of simple cells that surely prompted Hubel and Wiesel to propose the feedforward model is the similarity between the ON and OFF subfields of simple cells and the ON and OFF centers and surrounds of geniculate relay cells. That ON subfields of simple cells are in fact driven from input from ON-center LGN relay cells (and OFF from OFF) was demonstrated convincingly by spike-triggered averaging of the spike responses of a simple cell from a simultaneously recorded LGN cell (Tanaka, 1983). If an excitatory connection is detected, the receptive field center of the presynaptic LGN cell almost invariably overlaps a subfield in the simple cell of the same polarity (Figure 1C), and the stronger the connection, the more closely aligned the receptive fields (Reid and Alonso, 1995). Further confirmation of the feedforward model comes from experiments showing that the LGN relay cell axons that project into a cortical orientation column—recorded while the cortical neurons are silenced pharmacologically—have their receptive fields aligned parallel to the preferred orientation of nearby cells

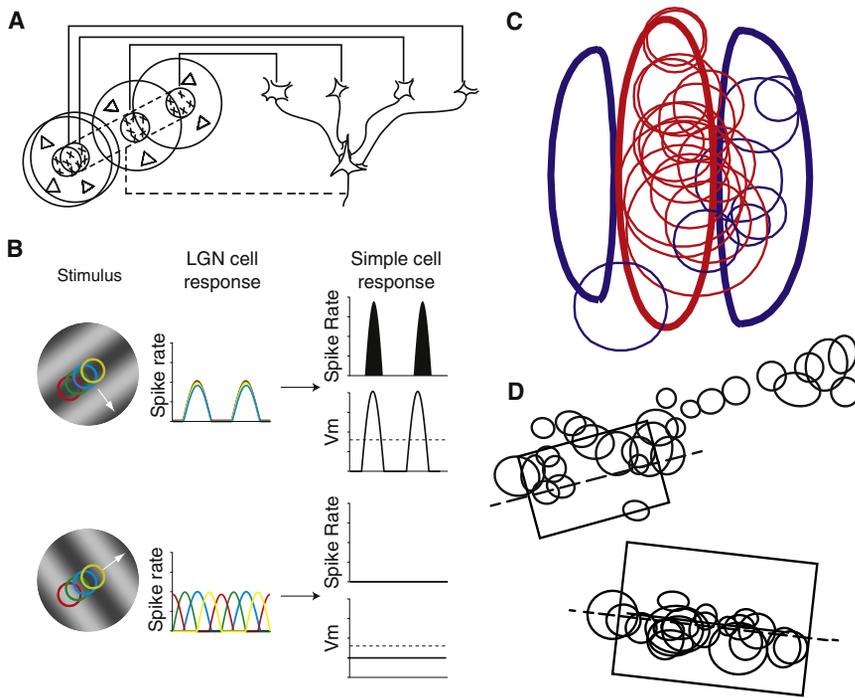


Figure 1. The Feedforward Model of Orientation Selectivity in Primary Visual Cortex

(A) The feedforward model as originally proposed by Hubel and Wiesel (1962). Four relay cells from the LGN (top right), whose receptive fields are shown to the left, synapse onto a V1 simple cell (bottom right). The simple cell derives its preferred orientation from the axis of alignment of these relay cell receptive fields and others like them that are not shown.

(B) The response of the feedforward model to drifting gratings in the preferred orientation (top) and the orthogonal orientation (bottom). LGN neurons with spatially offset receptive fields respond synchronously for the preferred orientation and asynchronously for the orthogonal orientation (middle panels). The average feedforward input increases in response to both stimuli, but only the preferred orientation response is sufficient to cross threshold (dotted lines) and evoke action potentials (right panels).

(C) The spatial relationship between the receptive fields of 23 recorded LGN relay cells (circles) and the receptive field of their postsynaptic simple cells (ovals). Each simple cell receptive field, along with its presynaptic LGN cell receptive fields, have been scaled and shifted to superimpose on an idealized receptive field. The image is adapted from Reid and Alonso (1995). Not shown is a tendency for LGN cells overlapping the center of a simple cell subregion to make stronger connections than those overlapping the periphery of the subregions.

(D) The receptive fields of two different sets of LGN relay cells that terminate in one column of V1 (circles), superimposed on the receptive field of a V1 simple cell recorded in layer 4 of the same column (square). The image is adapted from Chapman et al. (1991).

recorded prior to silencing (Figure 1D) (Chapman et al., 1991). Third, the summed receptive fields of a group of LGN cells projecting to a single orientation column—identified by spike-triggered averaging of cortical field potentials—form a simple-like receptive field aligned with the column's preferred orientation (Jin et al., 2011).

While there is little disagreement that a simple cell's preferred orientation is laid out by its geniculate input, less certain is whether feedforward input is sufficient to explain all of a simple cell's behavior, or whether additional circuit elements and mechanisms are required. Hints supporting the latter interpretation started to emerge soon after the 1962 paper. Hubel and Wiesel had made their observations delivering visual stimuli by hand and judging neuronal responses by ear. The subsequent introduction of methods for precise stimulus delivery and response measurement made possible a more quantitative description of simple cell response properties. A number of these properties appeared to be inconsistent with a purely feedforward model, at least a model in which all the elements were linear, as was commonly assumed. These properties, which we will address below, include (1) cross-orientation suppression, (2) contrast invariance of orientation tuning width, (3) contrast-dependent changes in response timing and in temporal frequency preference, and (4) the mismatch between measured orientation tuning and the tuning predicted by a simple cell's receptive field organization. Uncovering the origin of these properties has proven to be one of the keys to understanding the nature of the cortical computation.

One comprehensive solution to the origin of simple cell nonlinearities was suggested by psychophysics: in the tilt aftereffect illusion, the perceived orientation of a vertical stimulus is shifted away from vertical after prolonged viewing of a slightly oblique stimulus. This result was interpreted to mean that intracortical inhibition, specifically inhibition between cortical neurons of different preferred orientations, sharpened orientation tuning or even created it de novo (Blakemore and Tobin, 1972). This proposal was strengthened by pharmacological experiments: cortical application of GABA_A antagonists cause a broadening of orientation tuning (Sillito, 1975). Cross-orientation inhibition, a form of lateral inhibition (Hartline, 1949), but in the orientation domain rather than the spatial domain, is considered a natural extension of similar mechanisms either observed or proposed to operate throughout the brain. Because of the columnar organization of orientation preference in the cortex, the orientation domain translates into the spatial domain on the cortical surface. Cross-orientation inhibition can then emerge from simple, spatially defined rules of cortical connectivity.

Cross-orientation inhibition has been proposed to operate in several distinct modes, depending on the orientation dependence and amplitude of inhibitory interconnections. In attractor models, feedback inhibition forms a set of multistable attractors (Ben-Yishai et al., 1995; Somers et al., 1995), in which the width of orientation tuning of cortical cells is determined by the lateral extent of cortico-cortical connections. In recurrent models, recurrent excitatory connections amplify feedforward inputs in a way that is sculpted by lateral inhibitory connections (Douglas et al., 1995). Here again, the width of tuning and other aspects of

cortical responses are set by intracortical rather than thalamo-cortical interconnections. In balanced models, strong recurrent excitation and inhibition are thought to balance one another tightly (van Vreeswijk and Sompolinsky, 1998). In addition to explaining many aspects of simple cell behavior, this balance can explain the large variability of cortical spiking responses (Shadlen and Newsome, 1998). In push-pull models, cross-orientation inhibition arises from feedforward inhibition from simple cell-like inhibitory interneurons (Troyer et al., 1998, 2002), which themselves receive no inhibition and so fire at the null orientation, helping to establish contrast-invariant orientation tuning. In normalization models, a large pool of cortical interneurons of all different preferred orientations generates shunting inhibition proportional in strength to stimulus contrast at all orientations (Carandini et al., 1997; DeAngelis et al., 1992; Heeger, 1992). The excitatory thalamic inputs are therefore normalized (divided) by a signal proportional to contrast. Normalization models have been highly successful in explaining many of the contrast-dependent, nonlinear properties of simple cells and will be considered below in more detail.

One central driving force for inhibition-based models of cortical computation has been how well they can account for all of the simple cell's response nonlinearities (Carandini and Heeger, 2012). Aside from pharmacological experiments showing a degradation of orientation selectivity under GABA_A blockade, however, direct experimental evidence for strong cross-orientation inhibition in cat V1 is equivocal. Intracellular recording of membrane potential (V_m) in simple cells shows little hyperpolarization in response to nonpreferred stimuli (Ferster, 1986). Measurements of V_m alone, however, cannot rule out the presence of shunting inhibition; an increase in membrane conductance with a reversal potential at rest would generate no hyperpolarization yet would reduce the effectiveness of excitatory current in depolarizing the membrane. Detecting the presence of shunting inhibition requires injecting current into a cell while presenting visual stimuli to move the membrane potential away from the reversal potential of inhibitory synapses. Such experiments suggest that inhibition in simple cells has the same preferred orientation and tuning width as excitation (Anderson et al., 2000; Douglas et al., 1995; Ferster, 1986; Martinez et al., 2002; though see Monier et al., 2003). Overall, it appears that whatever shunting inhibition is present at the non-preferred orientation is too small to support the inhibitory models of orientation tuning.

Additional evidence that visually selective synaptic inhibition does not contribute directly to shaping orientation selectivity comes from experiments in which visually evoked action potentials in cortical cells are suppressed. During inactivation, either by cooling (Ferster et al., 1996) or by electrical stimulation (Chung and Ferster, 1998), orientation selectivity of the remaining excitatory input, the majority of which probably arises from the LGN, changes little. That is, the LGN inputs alone generate membrane potential responses that are as well tuned for orientation as the inputs from the fully functioning cortical circuit.

These results give rise to an apparent contradiction. The feedforward input to simple cells is probably organized very much as Hubel and Wiesel proposed but apparently fails to account for many properties of simple cells. Inhibition-based models can

account for these properties but lack definitive experimental support.

The resolution to this contradiction lies in the specific implementation of the feedforward model. The most common implementations tend to be highly simplified: LGN responses are assumed to be linearly related to stimulus strength; LGN cells excite simple cells in proportion to their spike rates. In contrast, real neurons are filled with nonlinear processes: spike threshold, synaptic depression, trial-to-trial response variability, driving force nonlinearities on synaptic currents, response saturation, and more. These nonlinear processes, it turns out, are critical in generating simple cell behavior: when we incorporate them into the feedforward model, almost all of the nonlinear properties of simple cells emerge in quantitative detail. Indeed, these properties are nearly unavoidable when the model is based on realistic synaptic and cellular mechanisms. Unlike many neuronal models, the resulting feedforward model is heavily constrained by experimental data. There are few intrinsic assumptions, few parameters, and all but two parameters are experimentally constrained; and the two unconstrained parameters can vary over a wide range without affecting the model's fit to the data. We will consider each of the nonlinear response properties of simple cells in turn and discuss how an amended feedforward model accounts for them.

Cross-orientation Suppression

Cortical spiking responses to a preferred ("test") grating (Figure 2A) are profoundly attenuated, or even completely extinguished, by simultaneous presentation of an orthogonal ("mask") grating (Figure 2B). This cross-orientation suppression has long been considered functional evidence for inhibition between neurons of distinct orientation preferences (cross-orientation inhibition). In support of this interpretation, antagonists of GABA_A-mediated inhibition reduce cross-orientation suppression in visually evoked potentials (Morrone et al., 1982, 1987). Cross-orientation suppression is also sensitive to the mask orientation, suggesting that neurons selective for orientation, such as those found in cortex, are key circuit elements underlying suppression.

Nevertheless, aspects of cross-orientation suppression appear to be at odds with a cortical mechanism. First, cross-orientation suppression is largely monocular (Ferster, 1981; Walker et al., 1998); a null-oriented mask stimulus presented to one eye has little effect on a preferred-oriented test stimulus presented to the other eye, whereas the majority of cortical neurons—presumably including inhibitory interneurons—are binocular. Second, strong suppression can be evoked by mask stimuli of high temporal frequency, beyond the frequencies to which most cortical neurons can respond (Freeman et al., 2002). Third, unlike most cortical cells, suppression is relatively unaffected by contrast adaptation (Freeman et al., 2002). Fourth, the onset of suppression is coincident with the onset of excitatory neuronal responses, earlier than the onset of inhibition (Smith et al., 2006). And finally, in intracellular recordings, inhibition appears to decrease, rather than increase, when a mask stimulus is superimposed on a test stimulus (Priebe and Ferster, 2006).

All of these features of cross-orientation suppression are more reminiscent of LGN relay cells than they are of V1 cells; relay cells

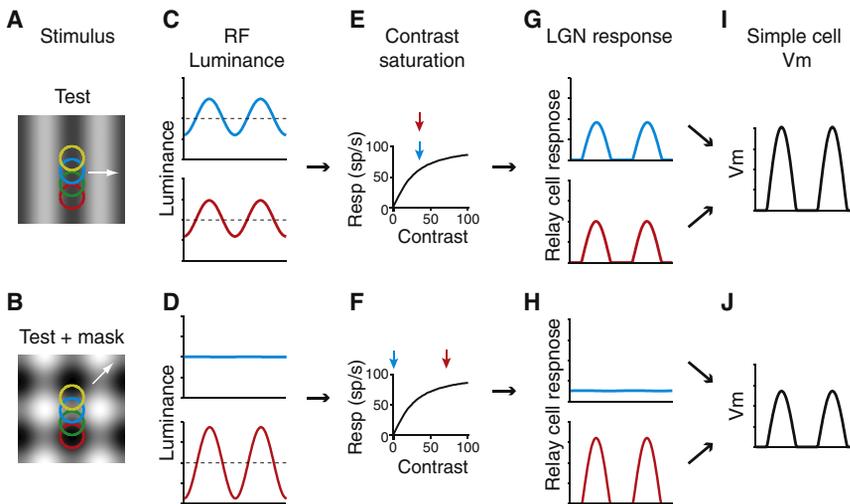


Figure 2. Cross-orientation Suppression in a Feedforward Model of Visual Cortex

(A and B) The spatial receptive fields of LGN relay cells (colored circles) are superimposed on top of a 32% contrast vertical grating (A) or a plaid composed of 32% horizontal and vertical gratings (B). (C and D) Stimulus luminance is plotted as a function of time for two LGN relay cells, indicated by color (C, grating; D, plaid). (E and F) The contrast response curve of LGN relay cells. The arrows indicate the contrast passing over each relay cell's receptive field (E, grating; F, plaid). (G and H) The modeled responses of the relay cells based on the contrast passing over their receptive fields include both saturation and rectification (G, grating; H, plaid). (I and J) The average input to a target V1 simple cell. The average relay cell input is about 10% less for the plaid stimulus (I) than for the grating stimulus (J).

are monocular, respond at high temporal frequency, adapt little to contrast, and, by definition, provide the excitatory input to the cortex. It has been proposed, therefore, that cross-orientation suppression arises from nonlinear interactions within the thalamocortical projection itself, rather than from within the cortex (Carandini et al., 2002; Ferster, 1986). One nonlinearity is synaptic depression: by increasing the overall level of activity in LGN cells, the mask stimulus could increase the overall level of depression at the thalamocortical synapses, thereby reducing the excitatory drive evoked by the test stimulus. Thalamocortical depression, however, may not be strong enough to account fully for cross-orientation suppression (Boudreau and Ferster, 2005; Li et al., 2006; Reig et al., 2006).

Alternatively, cross-orientation suppression may arise from two simple and well-described response nonlinearities of LGN relay cells: contrast saturation and firing-rate rectification (Ferster, 1986; Li et al., 2006; Priebe and Ferster, 2006). In response to drifting gratings, LGN relay cells modulate their firing rates in synchrony with the grating cycles, but because LGN relay cells have low spontaneous firing rates, high-contrast stimuli cause response rectification, clipping the downward phase of the response at 0 spikes/s (Figures 2C and 2D). Further, LGN responses do not increase linearly with contrast but instead saturate for contrasts above 32% (Figures 2E and 2F).

When the test and mask have identical contrasts, superimposing them results in a plaid pattern that moves up and to the right (Figure 2B, white arrow). Some LGN relay cells (e.g., Figure 2B, red) lie on a diagonal in the plaid stimulus where the dark bars from the two gratings superimpose, alternating with the locations where the bright bars superimpose. The result is a luminance modulation exactly twice as large as that generated by the test or mask stimuli alone (Figure 2D, red). The receptive fields of other LGN cells (e.g., Figure 2B, blue) lie at a location where the bright bars of one grating superimpose on the dark bars of the other and vice versa. These LGN cells see no modulation of luminance, and so their responses fall to zero (Figure 2D, blue).

Because the red curve has doubled in size while the blue one has fallen to zero, the sum of the two curves in Figure 2C is the

same as the sum of those in Figure 2D. For a purely linear system, in which the response of the LGN cells is proportional to the stimulus, and the simple cell sums its LGN inputs, the same would be true of the simple cell's V_m responses: the depolarization evoked by test + mask would equal the depolarization evoked by the test grating alone. It is the underlying assumption of linearity in many feedforward models, then, that leads to the conclusion that inhibition is required to explain cross-orientation suppression.

Contrast saturation and response rectification, however, are highly nonlinear. For the test + mask stimulus, the responses of the LGN cells that see no contrast modulation necessarily still fall to zero. But the responses of those LGN cells that see twice the contrast modulation (e.g., Figure 2H, red neuron) do not double. Their response to the test stimulus itself was already near saturation, so doubling the stimulus contrast increases the cell's responses only slightly. As a result, the sum of the LGN responses—and therefore the input to the simple cell—falls in the presence of the mask (compare Figures 2I and 2J).

Introducing realistic contrast saturation and rectification to an otherwise linear feedforward model results in cross-orientation suppression that is almost identical to that observed in real simple cells. In the model, the depolarization in a simple cell was taken to be proportional to the summed responses of eight LGN cells whose receptive fields were aligned in space. Response saturation and rectification were inserted by drawing the LGN responses from a database of the recorded responses of cat LGN cells (Priebe and Ferster, 2006). Cross-orientation suppression in the model matched closely the suppression observed in the V_m responses of V1 simple cells: 9% for the high-contrast test gratings and 52% for low-contrast test gratings (Priebe and Ferster, 2006).

To calculate the corresponding cross-orientation suppression in the spike responses of the model cell, the depolarization evoked by each stimulus was raised to the third power, to simulate the expansive nonlinearity of threshold as smoothed by trial-to-trial variability. The resulting cross-orientation suppression of the model's spiking responses (29% and 89% for

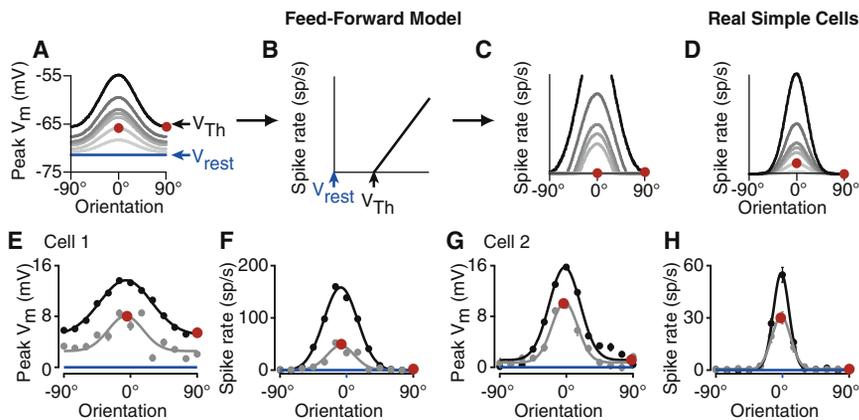


Figure 3. Contrast Invariance of Orientation Tuning

(A) Orientation tuning curves of a simple cell derived from a simple feedforward model for different stimulus contrasts. Red dots indicate the high-contrast null-oriented stimulus and the low-contrast preferred stimulus. (B) A threshold-linear relationship between V_m and spike rate. (C) The predicted orientation tuning curves for spike rate. (D) Tuning curves for spike rate in real recorded V1 simple cells show nearly identical width at all contrasts and zero response at the null orientation. (E and F) Orientation tuning curves for V_m (E) and spike rate (F) at 4% (gray) and 64% (black) contrast recorded intracellularly from a simple cell in cat V1. This cell probably received the bulk of its synaptic input from the LGN, as indicated by the significant depolarization at the null orientation (E).

(G and H) V_m (G) and spike (H) responses at 4% and 64% contrast for a simple cell that probably received the bulk of its synaptic input from other orientation-selective cortical cells, since it shows no depolarization at the null orientation (G). All response amplitudes are measured at the peak of the depolarization or spike rate increase evoked by a stimulus, which we derive from the mean response plus the amplitude of the response harmonic component at the stimulus frequency (DC + F1).

high- and low-contrast test stimuli) is consequently larger than what is predicted for V_m and is comparable to what has been observed experimentally.

While nonlinearities in relay cell responses account for the mask-induced reduction in the modulation component of simple cell V_m , these nonlinearities also predict a rise in the mean LGN input to V1 neurons, which is not observed experimentally. This discrepancy might arise in part from synaptic depression at the thalamocortical synapse (Carandini et al., 2002; Freeman et al., 2002) and because many simple cells receive less than half of their excitatory input from the LGN (Chung and Ferster, 1998; Ferster et al., 1996). In addition, some of the predictions of this model appear at odds with the interactions between low-contrast test + mask, for which relay cell contrast saturation should have little effect but nonetheless have been reported to interact in cortical complex cells (Busse et al., 2009; MacEvoy et al., 2009). Nevertheless, careful consideration of the properties of the thalamic input to cortical neurons reveals that a realistic feedforward model gives rise to cross-orientation suppression.

Contrast Invariance of Orientation Tuning

Natural stimuli are composed of a wide range of stimulus features. In order to extract these features properly, sensory systems must detect and respond selectively to stimulus features even in the face of large changes in signal strength. A primary method to address this problem is gain control, in which neurons adjust their responses on the basis of signal strength while maintaining the same relative feature selectivity. In this manner, the ratio of the responses of neurons with different stimulus preferences would be invariant to changes in stimulus strength and would therefore become a straightforward, strength-independent indicator of the stimulus parameter (Carandini and Heeger, 2012).

In V1, the width of orientation tuning of simple cells is invariant to stimulus contrast; the orientation tuning curve simply scales with contrast (Alitto and Usrey, 2004; Sclar and Freeman, 1982; Skottun et al., 1987). Contrast invariance, however, presents considerable difficulty for feedforward models of orientation

selectivity (Figure 3). A linear feedforward model predicts that the orientation tuning curve for the peak synaptic input from a row of LGN relay cells is approximately Gaussian in shape (Figure 3A). The curves ride on a nonzero offset because the LGN relay cells respond equally, although with different relative timing, at all orientations, including the null orientation. Thus, as relay cells' responses increase with contrast, both the offset and the amplitude of the simple cell's tuning curve increase proportionately. When these tuning curves for V_m are transformed by a simple threshold (Figure 3B), the predicted tuning curves for spike rate (Figure 3C), unlike in real simple cells (Figure 3D), are no longer contrast invariant. This apparent failure of the feedforward model is highlighted in Figure 3 by the red dots, which mark the responses to a high-contrast null stimulus and low-contrast preferred stimulus.

In some simple cells (Figures 3G and 3H), not even the V_m responses conform to the predictions of the feedforward model. Instead, they are themselves contrast invariant, with nearly identical tuning curve widths at different contrasts and virtually no depolarization at the null orientation at any contrast (Figure 3G). The spike-rate tuning curves are narrower than those for V_m (Figure 3H) but again do not narrow with contrast as would occur with a simple threshold nonlinearity.

Inhibition can easily account for the contrast invariance of real simple cells. Although the details vary, computational models have been developed that achieve contrast invariance using either cross-orientation inhibition (Troyer et al., 1998) or omniorientation inhibition (Ben-Yishai et al., 1995; Somers et al., 1995). Two experimental findings, however, suggest a different arrangement. First, simple cells receive little inhibition or excitation at the null orientation (Anderson et al., 2000). This observation would suggest that simple cells like those in Figure 3G, with contrast-invariant synaptic input, receive their dominant input from other cortical cells, rather than the LGN, since the spike output of most cortical cells is itself contrast invariant. In support of this proposal, in simple cells with invariant synaptic input, inactivation of the cortical circuit greatly reduces the size of visually evoked input (Finn et al., 2007). Conversely, in simple cells

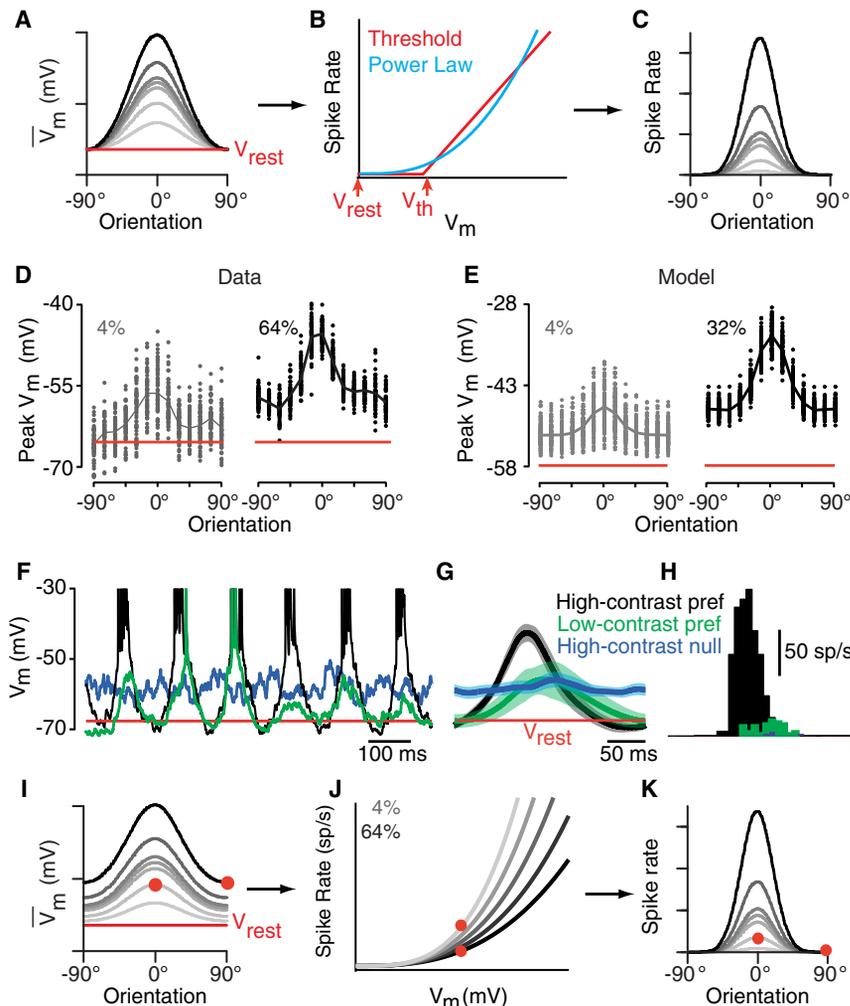


Figure 4. Trial-to-Trial Response Variability and the Origin of Contrast-Invariant Orientation Tuning in Simple Cells

(A–C) A power-law relationship between V_m and spike rate (B) will transform a set of Gaussian orientation tuning curves for V_m with identical widths (A) into a set of Gaussian tuning curves for spike rate, again with identical, but narrower, widths (C). Tuning curves with no offset from rest, as in (A), are typical of cells dominated by cortical input.

(D) Amplitudes of individual V_m responses (points) and mean response amplitude (curve) for low and high contrasts recorded intracellularly from a simple cell.

(E) Same as (D) but derived from a feedforward model as described in the text.

(F) Intracellularly recorded V_m responses to six cycles of a grating at three different combinations of orientation and contrast.

(G) Average and trial-to-trial standard deviation (shading) for the three stimuli.

(H) Average spike responses for the three stimuli. (I–K) As in (A)–(C), for a simple cell dominated by input from the LGN. The Gaussian-shaped tuning curves for V_m therefore ride on a contrast-dependent offset (I). In order to achieve contrast-invariant orientation tuning of the spike rate responses (K), the relationship between V_m and spike rate must be contrast dependent (J), as determined by the contrast dependence of trial-to-trial variability (D).

that receive most of their synaptic excitation from the LGN (Figure 3E), the V_m tuning curve rides on top of a contrast-dependent vertical offset as predicted by the feedforward model (Figure 3A).

Even if the feedforward model can explain the origins of the V_m tuning curves in Figures 3E and 3G, explaining how these curves are transformed into contrast-invariant spike-rate tuning curves in Figures 3F and 3H is more complex. The key lies in the ongoing cortical response variability, or noise, and how it affects the relationship between average V_m and spike rate (Anderson et al., 2000). Tuning curves for visual responses are generally derived from the averages of numerous trials. What is needed to understand contrast invariance, therefore, is not the familiar threshold-linear relationship between instantaneous V_m and spike rate but the relationship between mean V_m and mean spike rate. Mean spike rate, however, depends not only on average V_m but also on trial-to-trial variability (Carandini, 2004).

Consider, for example, a stimulus that evokes a mean depolarization that carries V_m , on average, half of the way toward threshold. One might expect such a stimulus to evoke, on average, no spikes. But because of trial-to-trial variability,

on some trials V_m actually exceeds threshold, whereas on others it stays near rest. The mean V_m is subthreshold, and yet the mean spike rate is no longer 0. Thus, for a given mean V_m , a higher variability leads to a higher mean spike rate. Similarly, for given variability, a higher mean V_m leads to a higher mean spike rate. Most importantly, the mean spike rate increases gradually, starting immediately from the resting potential, rather than remaining at 0 until threshold. This smoothing of the instantaneous threshold-linear relationship can be derived from a convolution of the threshold-linear curve with an approximately Gaussian distribution of trial-to-trial variability in V_m . The result approximates a power-law relationship, in which spike rate is proportional to $(V_m - V_{rest})^p$ (Hansel and van Vreeswijk, 2002; Troyer et al., 2002).

A power law is the one relationship between mean V_m and mean spike rate that can preserve contrast invariance of orientation tuning in spike rate for simple cells with invariant V_m , (i.e., in cells with predominantly cortical input, Figure 3G). Raising the Gaussians of the V_m responses, each with width σ , to a power p generates a new set of Gaussians for the spike responses, all with width $\sigma' = \sigma/\sqrt{p}$ (Figures 4A–4C). Since in most cells, p lies somewhere between 2 and 5 (Priebe et al., 2004), threshold generates a narrowing, or iceberg effect, of tuning width by a factor between 1.4 and 2.2.

Trial-to-trial variability also solves the problem of how the same mean depolarization for high-contrast preferred and low-contrast null stimuli (Figure 3E, red dots) can generate different

mean spike rates (Figure 3F, red dots) for simple cells dominated by input from the LGN (Finn et al., 2007). We know that mean spike rate depends on both mean V_m and trial-to-trial variability. Since mean V_m is the same for the two conditions, one of two things must change with contrast: either biophysical threshold or trial-to-trial variability. Biophysical threshold does vary somewhat in vivo (Azouz and Gray, 2000; Yu et al., 2008) in part because of moment-to-moment changes in dV_m/dt (Hodgkin and Huxley, 1952). But it does not change systematically with contrast. Trial-to-trial variability of the V_m responses, on the other hand, does. Figure 4D shows the average tuning curves at high and low contrast, along with the trial-to-trial variability (individual points), in which a trial is one cycle of a drifting grating. The larger vertical spread of points at low contrast leads to a systematic increase in the mean spike rate evoked by a given mean V_m .

The effects of a change in variability on the relationship between mean V_m response and mean spike rate are evident in raw membrane potential traces (Figure 4F). The V_m response to a high-contrast preferred stimulus (Figure 4F, black) is highly stereotyped across cycles and has a low standard deviation (Figure 4G, gray shading). V_m reaches threshold on every stimulus cycle and evokes significant numbers of spikes (Figure 4H, black). The V_m response to the high-contrast null stimulus (Figure 4F, blue) also varies little from trial to trial, has a low standard deviation (Figure 4G, blue), and because it is below threshold on nearly every trial, evokes few spikes (Figure 4H, blue). The response to a low-contrast preferred stimulus (Figure 4F, green) differs significantly in character. Its mean response (Figure 4G, green) peaks at exactly the same sub-threshold potential as the high-contrast null response (Figure 4G, blue) but has far greater trial-to-trial variability and standard deviation (Figure 4G, green shading). Because of the increased variability, on some trials the cell reaches threshold (Figure 4F, cycles 2 and 3) and the resulting mean spike rate is significant (Figure 4H, green).

We can now summarize the full transformation between V_m and spike rate for simple cells that receive their dominant input from the LGN. The V_m tuning curves in Figure 4I are transformed by a different power law for each contrast (Figure 4J) to give the spike-rate tuning curves in Figure 4C. As variability rises with decreasing stimulus contrast, the mean spike rate evoked by any given mean V_m rises as well. Thus, two stimulus conditions that evoke similar mean V_m responses evoke very different numbers of spikes (Figures 4I–4J, red dots).

If response variability and its contrast dependence contribute to the contrast invariance of orientation tuning, the next question becomes, “What is the source of the V_m response variability?” One possible source is trial-to-trial changes in cortical excitability. In this case, feedforward thalamic input would be stable from trial to trial, whereas amplification by the cortical circuit would vary from trial to trial. Intracortically generated shunting inhibition, for example, could modulate variability in a contrast-dependent manner (Monier et al., 2003), perhaps in association with the occurrence of cortical up and down states (Haider and McCormick, 2009; Stern et al., 1997). To determine the contribution of the cortical circuit to response variability of simple cells, Sadagopan and Ferster (2012) measured variability while the

cortical circuit was inactivated. As mentioned above, inhibition evoked by electrical stimulation of the cortex suppresses spike responses locally, without strongly affecting the LGN (Chung and Ferster, 1998). Even with the cortical circuit inactivated, at all orientations, V_m responses to flashing high-contrast stimuli still showed less variability than did responses to low-contrast stimuli, suggesting that intracortical circuitry neither generates nor amplifies variability in a contrast-dependent manner.

An alternate source of contrast-dependent changes in cortical response variability is the feedforward thalamic input. In this hypothesis, spontaneous fluctuations in the retina and the LGN are suppressed by visual stimulation in a manner that is dependent on the strength of the visual stimulation. To test this possibility, Sadagopan and Ferster (2012) made extracellular recordings from LGN cells under the same conditions as those Finn et al. (2007) used to make intracellular recordings from simple cells. As described previously (Hartveit and Heggelund, 1994; Sestokas and Lehmkuhle, 1988), for a given response, variance was lower at high contrast than at low contrast. Over the population, the average Fano factor (variance/mean) dropped nearly 45% (from 2.1 to 1.3) between 2% contrast and 32% contrast.

As suggestive as this change in variability is, however, it alone cannot explain the V_m response variability in simple cells. Cortical simple cells clearly pool the inputs from a number of LGN relay cells, and if the variability in each of those inputs were completely independent, then the variability in the simple cell would be lower than the variability in the individual inputs by \sqrt{N} , where N is the number of inputs. In that case, the V_m response variability of a simple cell with 10–20 independently varying LGN inputs would be three to four times lower than the variability of individual LGN relay cells, far lower than what is observed in simple cells. If, on the other hand, response variability in LGN relay cells were perfectly correlated, then variability in a simple cell’s V_m responses would be the same as in its presynaptic LGN cells. Simultaneous recording in groups of nearby LGN cells shows that the correlation coefficient for variability between pairs of cells falls in the range of 0.25, with little variation as a function of stimulus contrast (Sadagopan and Ferster, 2012).

With measurements of LGN response variability, its dependence on contrast, and its cell-to-cell correlation, it is possible to construct a highly constrained feedforward model of a V1 simple cell. Sadagopan and Ferster (2012) simulated simple cells with input from between 8 and 32 individual LGN cells arranged in two subfields with aspect ratios between 2 and 4. Each presynaptic LGN cell generated a change in conductance in the model simple cell in proportion to its spike rate, after application of rate-dependent synaptic depression (Boudreau and Ferster, 2005). The simple cell’s resting conductance and the peak conductance evoked by an optimal stimulus were taken from intracellular measurements (Anderson et al., 2000). Orientation tuning curves for the mean V_m response in the modeled simple cell are shown in Figure 4E at two different contrasts (solid curves), along with the responses on individual trials (points). The model’s response variability compares well to that of real simple cells (Figure 4D) in both the relative amplitude of the mean responses and the contrast-dependent—and relatively

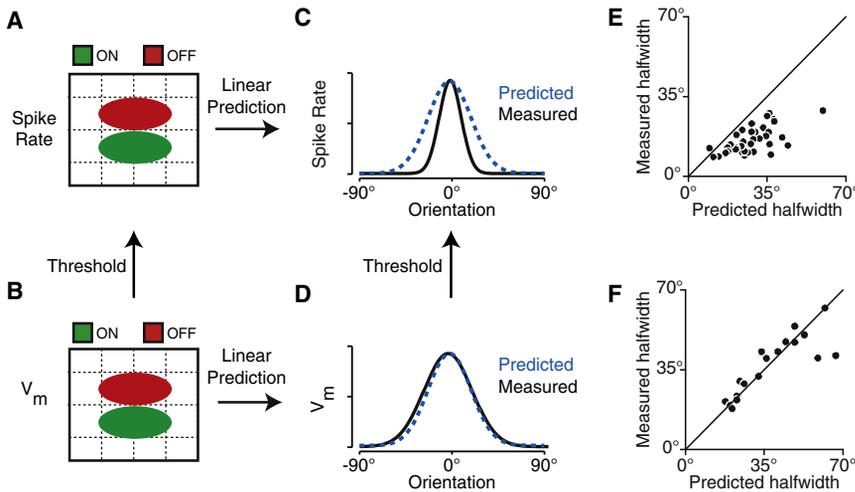


Figure 5. The Match between Measured Width of Orientation Tuning and that Predicted by Receptive Field Maps

(A and B) Simple cell receptive field maps based on spike rate (A) and membrane potential (B) are generally matched. Red indicates spatial locations with preference for OFF and green indicates preference for ON.

(C) Orientation tuning curves measured from peak spike rate (black) are narrower than those predicted from receptive field maps (blue).

(D) Measured and predicted tuning curves are matched for membrane potential.

(E) Population data showing the mismatch between measured and predicted tuning width in spike rate data (replotted from Gardner et al., 1999).

(F) Population data showing the match between measured and predicted tuning width in membrane potential data (replotted from Lampl et al., 2001).

orientation-independent—change in trial-to-trial variability. These features of the responses are relatively robust to changes in the two free parameters of the model, the number of LGN inputs, and the aspect ratio of the simple cell receptive field.

One surprising aspect of the model is that the match with the data requires the nonlinearities of synaptic depression and of the relationship between conductance and V_m (“driving force nonlinearity”). When these are removed from the model, the trial-to-trial variability becomes less dependent on contrast and more dependent on orientation. In other words, it is a combination of biophysical mechanisms that contribute to the contrast invariance of V1 simple cell responses.

Mismatch of Receptive Field Maps and Orientation Tuning

If orientation tuning were derived solely from the spatial organization of LGN input, it should be possible to predict the orientation tuning curve of a simple cell from a detailed map of its receptive field. That is, if both the orientation tuning curve and the receptive field map arise from the spatial organization of the thalamic input, there should be a direct correspondence between the two. Indeed, predictions derived from the receptive field map and measured orientation tuning curves match closely in orientation preference. When assayed from spike rate, however, there is a strong mismatch between predicted and measured width of tuning (Gardner et al., 1999; Jones and Palmer, 1987). The measured tuning is far narrower than the predictions (Figures 5A–5C).

This difference in selectivity has often been interpreted as evidence for intracortical cross-orientation inhibition. Lateral inhibition—particularly shunting inhibition—could selectively antagonize the feedforward excitatory input at orientations to either side of the preferred. Predicted tuning curves would reflect only the broadly tuned thalamocortical input, whereas measured tuning curves would include the sharpening effects of intracortical inhibition. As noted above, however, direct evidence for cross-orientation inhibition is not consistently observed.

An alternative mechanism that can account equally well for the tuning mismatch, and is present in all neurons, is spike

threshold. Threshold allows only the largest membrane potential deflections—those evoked by orientations close to the preferred orientation—to evoke spikes. This iceberg effect narrows the orientation tuning measured from spike rate about 3-fold, relative to the tuning for V_m responses (Carandini and Ferster, 2000; Volgushev et al., 2000). If threshold were responsible for the selectivity mismatch between receptive field maps and tuning curves, then a number of consequences follow. First, the mismatch between measured and predicted tuning width for spike rate responses should be comparable to the 3-fold narrowing of the iceberg effect. Second, the mismatch should disappear if threshold were taken out of the equation. And indeed it has been found (Lampl et al., 2001) that the measurements of tuning width match closely with predictions drawn from receptive field maps when both are drawn from V_m responses (Figures 5D–5F).

This match at the membrane potential level constrains the locus of the mismatch between tuning curves and receptive field maps to a point after the integration of synaptic inputs into membrane potential in the cortical simple cell. If synaptic inhibition were the mechanism underlying the mismatch between receptive field maps and tuning curves, then the mismatch would be evident in membrane potential as well.

If threshold so clearly narrows the orientation tuning curves, one question that remains is why there is no commensurate effect on the receptive field map? Why do the maps derived from spike rate and membrane potential match closely (Figures 5A and 5D)? The answer lies in the nature of stimuli employed to measure the receptive field maps. Receptive field maps are generally derived from a noise stimulus in which spots of light are flashed randomly (and independently) at each location in the receptive field simultaneously. At any one moment, several excitatory locations are likely to be on, and the membrane potential fluctuates near threshold. In this case, locations that on their own evoke only subthreshold membrane potential responses (were they to start from rest) can still influence the ongoing spike rate and be detected in spike rate recordings. Thus, a noise stimulus circumvents the threshold nonlinearity, resulting in a spiking receptive field map that is

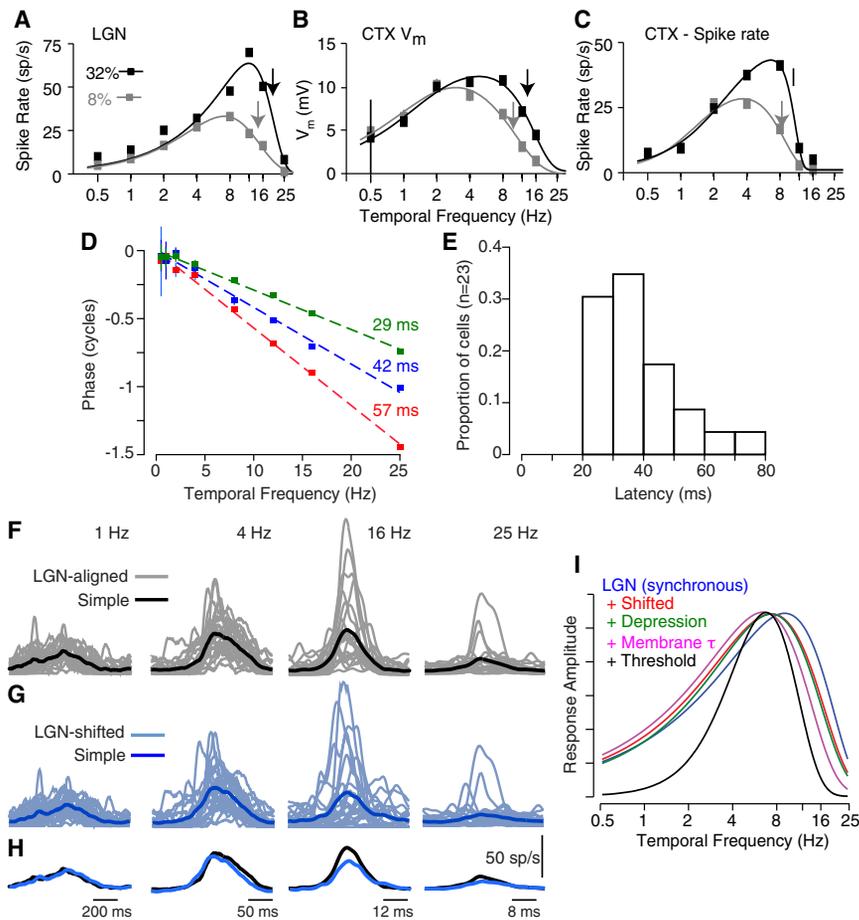


Figure 6. The Change in Temporal Frequency Tuning between LGN and Cortex (A–C) Temporal frequency tuning for peak response at two different contrasts for an LGN cell (A), and the V_m responses (B), and spike rate responses (C) of a V1 simple cell. (D) Response phase versus temporal frequency for three different LGN cells. (E) A histogram of the visual latency for 23 LGN cells, where visual latency is the slope of the relationship between phase and TF, as in (D). (F) Averaged single-cycle responses of 23 LGN cells to drifting gratings at four different TFs (gray). The responses have been shifted to have identical temporal phases at the stimulus frequency. The responses of all 23 cells are then averaged to model the input to a simple cell (black). (G) As in (F), but here the LGN responses are shifted relative to one another according to their visual latencies (D and E). (H) A comparison of the model simple cell responses from (F) and (G). (I) Temporal frequency tuning curves from a model simple cell. Moving from right to left, for each curve one additional mechanism has been added to the model as indicated by the color code.

comparable to that recorded directly from V_m responses (Mohanty et al., 2012).

Threshold is also likely to provide the explanation for why pharmacological blockade of GABA_A-mediated inhibition broadens orientation tuning in cortical cells (Sillito, 1975). Blocking inhibition appears to increase the overall excitability of cortical neurons such that previously ineffective stimuli on the edges of the spike-rate tuning curve become suprathreshold (Katzner et al., 2011).

Temporal Nonlinearities in Simple Cell Responses

Up to now, we have considered receptive field properties in the spatial domain—that two stimuli of different orientations suppress one another, that orientation tuning is contrast invariant, and that the width of orientation tuning is narrower than predictions based on the receptive field map. Here we consider three temporal aspects of simple cell responses that also fail to emerge from the simplest forms of the feedforward model.

First, simple cells do not respond well to rapidly changing stimuli. Compared to LGN cells, the preferred temporal frequencies (TFs) of simple cells are lower by a factor of 2 (Hawken et al., 1996; Orban et al., 1985). Here, temporal frequency refers to the number of bars of the drifting grating that pass over the receptive field in each second. Compare, for example, the TF tuning of the

LGN and cortex does not represent a nonlinearity; a linear, low-pass RC filter could shift the peak frequency of a simple cell's output relative to its input.

The second temporal feature of simple cells is that the preferred TF in simple cells decreases almost 2-fold with decreasing stimulus contrast (Albrecht, 1995; Carandini et al., 1997; Hawken et al., 1996; Holub and Morton-Gibson, 1981; Reid et al., 1992). Compare, for example, the black and gray tuning curves in Figure 6C. This property does represent a nonlinearity: the transformation between stimulus and response changes with stimulus strength (contrast).

One element that surely contributes to the mismatch in preferred TF between simple cells and their synaptic input from the LGN is the membrane time constant, τ . Together, the membrane input resistance R and membrane capacitance C form a linear low-pass filter with a time constant $\tau = RC$, which lies near 15 ms for most simple cells (Anderson et al., 2000). The frequency at which such a filter attenuates its input by a factor of 2 ($f_{3dB} = 1/2\pi\tau$) is about 11 Hz. The presence of this filter can explain about one-half the difference between preferred TFs in the V_m responses of simple cells and in the spike responses of LGN cells.

Much of the remaining difference probably arises when multiple LGN cells with slightly different visual latencies

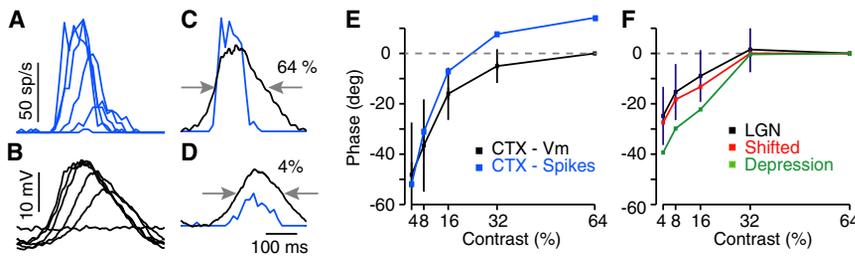


Figure 7. The Contrast-Dependent Change in Response Timing

(A and B) Spike rate (A) and V_m (B) responses of a simple cell to drifting gratings (TF = 2 Hz) at six different contrasts (0%, 4%, 8%, 16%, 32%, and 64%). Amplitude increases and phase advances with increasing contrast.

(C and D) V_m and spike rate responses superimposed for 64% (C) and 4% (D) contrast to show the relative phase shift between the two.

(E) Response phase (relative to V_m response at 64%) as a function of contrast for the records in (A) and (B).

(F) Response phase as a function of contrast for three different simple cell models: LGN responses synchronized in phase and averaged, LGN responses shifted according to visual latency and averaged, and LGN responses shifted and then passed through a model of synaptic depression.

converge on a single simple cell. Here, visual latency is defined as the slope of the relationship between response phase (relative to stimulus phase) and temporal frequency for a flickering grating (Saul and Humphrey, 1990). This relationship is shown for three different cells in Figure 6D, and a histogram of the slopes for 23 cells is shown in Figure 6E.

To understand how the spread of LGN latencies affects the feedforward model of a simple cell, we first created a model in which a number of LGN cells with identical latencies converge on a simple cell. We therefore superimposed the responses of the 23 recorded simple cells after aligning their responses to have identical temporal phases at four different TFs (Figure 6F, gray). The depolarization in the simple cell was taken to be proportional to the mean of the 23 input waveforms (black).

The LGN latencies are not identical (Figure 6E), however, but vary from one another by as much as 60 ms. As a result, even though receptive fields of the presynaptic LGN cells might be perfectly aligned in space, their responses will be misaligned in time. In a more realistic model, then, each response waveform in Figure 6F must be shifted by the visual latency of the corresponding LGN cell (Figure 6G). This creates a subtle dispersion of the peaks of the LGN responses. At low TFs (1–4 Hz), this dispersion is barely noticeable; the temporal shift, 60 ms, constitutes only one-sixteenth to one-fourth of a cycle. At these TFs, therefore, the latency shifts change the summed input to the simple cell hardly at all (blue traces). At higher TFs (8–16 Hz), however, 60 ms translates to a large proportion of a cycle (Figure 6H). The dispersion of the peaks of the individual LGN traces is easily visible and has a significant effect on the amplitude of the summed input to the simple cell. In other words, the temporal dispersion of the LGN inputs acts like a low-pass filter, selectively attenuating the peak of the visually evoked conductance change at high TFs (Figure 6I, compare synchronized LGN inputs, blue, and latency-shifted LGN inputs, red).

To make the model somewhat more realistic, we further added short-term synaptic depression (green), a membrane time constant of 15 ms (magenta), and finally a power-law relationship between V_m and spike rate (black). The overall effect is to shift the tuning curve of the model simple cell about two octaves to the left, as is observed in records from simple cells (Figures 6A and 6C, black). Repeated simulations of a simple cell, in which we selected a subset of cells from a population of 23 recorded LGN cells, showed shifts in the preferred TF of the model simple cell on average by 4.5 Hz and shifts in the TF_{50} of 8 Hz.

Note that this version of the feedforward model also exhibits the contrast-dependent shift in the preferred TF and TF_{50} seen in simple cells (Figures 6B and 6C, compare black and gray). The shift arises in part from the LGN responses, which themselves show such a shift (Figure 6A). In addition, at the preferred orientation, high-contrast stimuli decrease the simple cell's input resistance and therefore the membrane time constant (τ) about 2-fold (Anderson et al., 2000; Douglas et al., 1995). A 2-fold decrease in τ raises the cutoff frequency (f_{3dB}) of the membrane low-pass filter by a factor of 2 and therefore should raise the preferred temporal frequency and TF_{50} of the membrane potential responses.

The third temporal nonlinearity in simple cell responses is a phase advance with increasing contrast (Albrecht, 1995; Carandini and Heeger, 1994; Dean and Tolhurst, 1986). That is, the timing of spike responses shifts earlier and earlier as stimulus contrast increases (Figure 7A). One unexpected finding from intracellular records is that simple cell spike responses are consistently phase advanced relative to the underlying V_m responses (Figures 7C and 7D). Some mean membrane potentials evoke significant spike rates in the rising phase of the response (Figures 7C and 7D, right arrows) and yet no spikes on the falling phase (left arrows). A stationary threshold or power-law relationship between V_m and spike rate will not capture this behavior. Some aspect of the V_m -to-spike relationship is probably changing during the response. For example, trial-to-trial variability might change during the course of the response, or spike adaptation might occur. The maximum effect occurs at high contrasts (Figure 7E), in which the phase shift between V_m and spike rate is almost 20° .

We also noted that the contrast-dependent phase advance is smaller in V_m than it is in spike rate (Figure 7E). About half of the 48° phase shift in V_m between low and high contrast (Figure 7E, black) can be attributed to the responses of LGN cells (Figure 7F, black), which have a 25° phase shift of themselves. Adding a realistic dispersion in visual latency (as we did for the preferred TF shift above) has only a very small effect on the phase shifts of V_m responses in a model simple cell (Figure 7F, red). Adding synaptic depression (from Boudreau and Ferster, 2005) brings the total phase shift of the model to 40° . Depression, like spike adaptation, has the effect of reducing the depolarization evoked on the falling phase of the response relative to the rising phase, since the falling phase is preceded by a period of high activity and the rising phase is preceded by a period of low activity. Thus, it appears that the contrast-dependent phase advance is

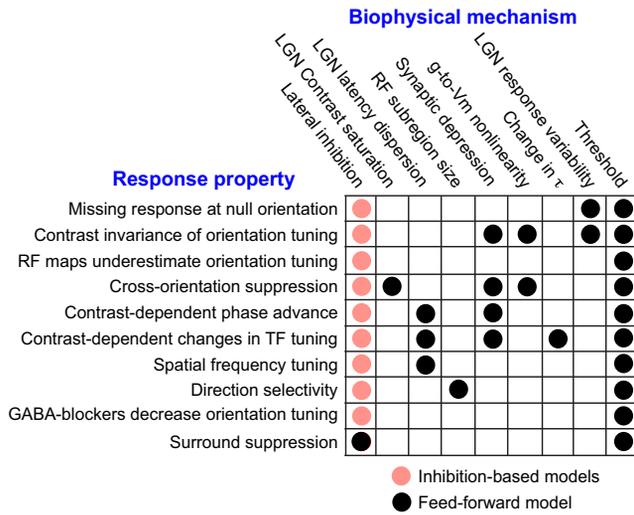


Figure 8. The Biophysical Mechanisms Underlying the Response Properties of V1 Simple Cells

Inhibitory models propose that all of a simple cell's nonlinear properties arise from intracortical synaptic inhibition that is either unselective for orientation or selective for the orthogonal orientation to the feedforward input (red points). Alternatively, a simple feedforward model incorporating experimentally determined nonlinear processes can account for most of the behavior of simple cells. Black points indicate which mechanisms contribute to which simple cell properties.

primarily accounted for by the responses of the LGN relay cells in combination with their known synaptic dynamics.

Conclusions

On the basis of their remarkable discovery of orientation preference, Hubel and Wiesel put forward a simple yet powerful model of how selectivity could emerge from nonselective thalamocortical inputs: that cortical neurons integrate input from LGN afferents with aligned receptive fields. This concise idea has become a central frame of reference for understanding cortical computation. Yet, it stands in contrast to many models of sensory processing. Since Hartline first described lateral inhibition in the retina (Hartline, 1949), lateral inhibition has been either found experimentally or proposed on theoretical grounds to operate in almost every sensory modality, and at every level of the brain, from the sensory periphery to cognitive and perceptual processing. It has been invoked to sculpt the crude selectivity of excitatory inputs for everything from sound frequency, to odorants, to phonemes. Hubel and Wiesel's model, by its omission, raises the question of whether, and how, inhibition contributes to generating the quintessential feature of cortical receptive fields.

A number of cortical receptive field properties have seemed at odds with the simple account provided by Hubel and Wiesel. These response properties have challenged the essence of the feedforward model and forced a critical evaluation of the mechanisms underlying cortical computation. Most of the nonlinear response properties discussed here can be described quantitatively within a theoretical framework in which the feedforward synaptic drive is normalized by a signal related to stimulus contrast (Carandini and Heeger, 2012; Carandini et al., 1997;

Geisler and Albrecht, 1992; Heeger, 1992). Formally, the response, R , of a cortical neuron can be described as:

$$R = R_{max} \left[h \frac{c}{\sqrt{c_{50}^2 + c^2}} \right]^n$$

where h is the linear, orientation-selective, feedforward drive, c is stimulus contrast, and c_{50} is the contrast at which R reaches half its maximal value (R_{max}). With proper selection of parameters, this one equation can fit the complete array of simple cell behaviors, including contrast saturation, cross-orientation inhibition, and surround suppression.

The equation itself is agnostic regarding the mechanism underlying contrast-dependent normalization; the normalization computation fits simple cell behavior well regardless of the origin of the contrast-dependent normalization signal (Carandini and Heeger, 1994). One widely discussed mechanism is shunting inhibition, in which contrast-dependent changes in input resistance scale the depolarization generated by the feedforward drive. Inhibition could arise either from pooling the input from orientation-specific interneurons with a range of preferred orientations or from interneurons that are unselective for orientation (Azouz et al., 1997; Cardin et al., 2007; Hirsch et al., 2003). In addition, the change in input conductance, through its effect on the membrane time constant, τ , could account for the temporal nonlinearities of simple cells (contrast-dependent changes in preferred temporal frequency and response phase). With shunting inhibition, then, a single underlying mechanism could give rise to the complete spectrum of simple cell properties (Figure 8, red dots).

And yet the inhibition most often found in cortical cells has neither the magnitude nor the orientation independence required to support the normalization framework. It is this observation that prompted our reexamination of the feedforward model. We find that when a series of biophysical properties common to nearly all neurons is incorporated into a feedforward model, all of the observed nonlinear properties of simple cells emerge (Figure 8, black points). None of these mechanisms is orientation specific and many are not even specific to the visual system. Driving force nonlinearity on synaptic currents, spike threshold, and synaptic depression are found throughout the brain; trial-to-trial response variability (Churchland et al., 2010) and response saturation are found across many sensory and motor systems.

Although the modified feedforward model accounts for much of the behavior of simple cells, it has only two free parameters: the number of presynaptic LGN cells and the aspect ratio of the simple cell's subregions. Even these two parameters have a wide range of permissible values. All of the other properties of the model are experimentally constrained, including thalamocortical synaptic depression, the relationship between V_m and spike rate, latency dispersion and contrast saturation in LGN cells, the driving-force nonlinearity on synaptic currents, and the membrane time constant. Thus, when the feedforward model is made realistic by the addition of very basic and well-characterized neuronal mechanisms, the known properties of simple cells emerge per force.

Among the biophysical mechanisms that contribute to cortical receptive fields, threshold has by far the most influence.

Simple cells rest well below threshold and have very little spontaneous activity. The resulting iceberg effect narrows orientation tuning for spikes relative to V_m by as much as 3-fold or more, increases direction selectivity by 4-fold or more (Carandini and Ferster, 2000; Lampl et al., 2001), increases spatial frequency selectivity (Lampl et al., 2001), enhances the distinction between simple and complex cells (Priebe et al., 2004), and increases ocular dominance (Priebe, 2008). Because of the iceberg effect, cortical connections need not be nearly as specific as they appear to be in measurements derived from spike responses; the V_m responses at the periphery of the tuning curve are hidden by threshold. Threshold might also have important implications for plasticity and development. The dramatic changes seen, for example, in ocular dominance plasticity are most often measured from spike responses. Changes in spike responses, however, might be generated by smaller shifts in the ocular dominance of V_m responses and therefore by relatively smaller changes in connectivity (Priebe, 2008).

Though inhibition may not sculpt orientation selectivity in cat V1, it is nonetheless a fundamental component of the cortical circuitry. Estimates put the proportion of inhibitory neurons in layer 4 at 25%. Inhibition and excitation share selectivity: those stimuli that elicit excitation also elicit inhibition onto cortical neurons (Douglas et al., 1988; Ferster, 1986). One possible function of such shared selectivity is to maintain the stability of the cortical circuitry. Inhibition allows a circuit to have strong excitatory recurrent connections to amplify small signals without risking runaway feedback in the excitatory network (Douglas and Martin, 1991). Strong excitatory recurrence in turn increases the dynamic range of cortical neurons, increases their information-carrying capacity, increases the ability of the cortex to perform complex computations (Hansel and Sompolinsky, 1996; Latham and Nirenberg, 2004; Tsodyks et al., 1997; van Vreeswijk and Sompolinsky, 1998), and may underlie surround suppression (Ozeki et al., 2009).

Surround suppression is one receptive field property that probably requires strong lateral inhibition (Figure 8, black dot in column 1). But here, the underlying inhibition has the same preferred orientation as excitation: surround suppression is greatly reduced when the surround stimulus is presented at the cross-orientation (Hubel and Wiesel, 1965; DeAngelis et al., 1994). Thus, the inhibition is "lateral" in the spatial domain, rather than in the orientation domain. The effects of even this inhibition, however, may be weak in simple cells. Among simple cells that are dominated by excitation from the LGN, few exhibit strong surround suppression (Ozeki et al., 2009).

Much effort has been directed recently into uncovering the mechanisms underlying orientation selectivity in rodents. The mouse provides opportunities to exploit recent advances in genetic labeling of specific neuronal subsets, in optogenetics, and in imaging. These techniques promise an even more detailed and fine-grained understanding of the cortical circuit than has so far been possible in the cat. Reports that inhibitory neurons are more broadly orientation selective than excitatory neurons (Kerlin et al., 2010; Runyan et al., 2010) and that the tuning width of inhibition recorded intracellularly is broader than that for excitation (Atallah et al., 2012; Li et al., 2012) raise the possibility of cross-orientation inhibition in the mouse. Not all results are in

agreement, however (Tan et al., 2011), and some experiments suggest that threshold is as important or more so in shaping neuronal responses (Jia et al., 2010). Whether or not mouse V1 uses identical mechanisms to cat V1, the following differences exist between the two in overall organization: mouse receptive fields are almost ten times larger than those in the cat, as is preferred stimulus size; mice have no orientation columns; it appears that the cortico-cortical excitatory inputs in the mouse come from cells of widely different orientation preference (Jia et al., 2010; Ko et al., 2011); and in some reports (Kerlin et al., 2010; Kuhlman et al., 2011; Sohya et al., 2007; Tan et al., 2011), though not others (Niell and Stryker, 2008; Wang et al., 2010), orientation tuning in the mouse is somewhat weaker than in the cat and in primates. We note, however, that most of the mechanisms that operate in concert with the feedforward model in the cat, including threshold, synaptic depression, response variability, and the conductance nonlinearity, will almost certainly be present in the mouse as well.

Hubel and Wiesel's original feedforward model contained two hierarchical stages, one to explain the emergence of V1 simple cells from LGN afferents and a second stage to explain the emergence of V1 complex cells (characterized by overlapping ON and OFF responses) from simple cells within V1. The model posits that V1 complex cells integrate excitatory inputs from a subset of simple cells of similar orientation preference but with different receptive field positions. Several lines of evidence support this aspect of the feedforward model: (1) spike-triggered averaging of simple- and complex-cell pairs show excitatory connections from the former to the latter (Alonso and Martinez, 1998); (2) anatomical studies show a strong projection from layer 4, which is dominated by simple cells, to the superficial layers, which is dominated by complex cells (Gilbert and Kelly, 1975); and (3) silencing simple cells generally silences complex cells (Martinez and Alonso, 2001).

One aspect of the original hierarchical feedforward model that has been open to question is whether the shift from simple cells to complex cells is made in one step, or whether multiple steps are required to generate completely overlapping ON and OFF subfields (Chance et al., 1999). The observed diversity in subfield overlap suggests that the generation of complex cells with completely overlapping ON and OFF subfields may emerge imperfectly (Priebe et al., 2004; Rust et al., 2005; though see Martinez et al., 2005). Nonetheless, the data are generally consistent with the hierarchy proposed by Hubel and Wiesel.

Orientation Selectivity as a Model for Cortical Computation

Orientation selectivity was originally identified in cat V1 and has since been identified in every mammalian species examined. The degree of orientation selectivity, the exact layer in which it emerges in the cortex, and whether cells of similar orientation preference are organized into columns varies between species, but orientation selectivity still appears to be a fundamental component of the image that V1 extracts. This raises the question of how well a computation performed in V1 represents the computations performed throughout the many areas of the cerebral cortex. Does V1 contain highly specialized and unique machinery for the computation of orientation from the

retinal image? Or do other areas of cortex perform a similar feedforward computation on inputs carrying different types of information?

The anatomical (Brodman, 1909) and emerging molecular (Bernard et al., 2012) diversity of cortical areas might suggest that the computation in V1 is specific, if not to V1 itself, then to other primary sensory areas that share a similar laminar structure. On the functional level, however, there are indeed reasons to believe that diverse cortical areas share common computational mechanisms. First, the normalization framework, which is a prominent feature of the V1 computation, is not limited to V1 but appears in many parts of the sensory system (Carandini and Heeger, 2012). Even high-level processes such as response modulation related to attention or behavioral state can be described as a normalization-like shift in gain (Reynolds and Heeger, 2009). Second, the feature selectivity for excitation and inhibition are often matched in other cortical areas as they are in V1. Third, the neuronal mechanisms underlying V1 feature selectivity are not limited to neurons in V1. Threshold, response variability, driving-force nonlinearities, response saturation, low-pass filtering, response diversity, and synaptic depression are mechanisms inherent to all neurons that support action potentials. Whether neurons in other areas of the cortex take advantage of them, and, if so, whether they use them in ways analogous to V1, is an open question.

ACKNOWLEDGMENTS

This work was supported by NIH grants R01 EY04726 (D.F.) and R01 EY019288 (N.J.P.) and by a grant from the Pew Charitable Trusts (N.J.P.).

REFERENCES

- Albrecht, D.G. (1995). Visual cortex neurons in monkey and cat: effect of contrast on the spatial and temporal phase transfer functions. *Vis. Neurosci.* *12*, 1191–1210.
- Allitto, H.J., and Usrey, W.M. (2004). Influence of contrast on orientation and temporal frequency tuning in ferret primary visual cortex. *J. Neurophysiol.* *91*, 2797–2808.
- Alonso, J.M., and Martinez, L.M. (1998). Functional connectivity between simple cells and complex cells in cat striate cortex. *Nat. Neurosci.* *1*, 395–403.
- Anderson, J.S., Carandini, M., and Ferster, D. (2000). Orientation tuning of input conductance, excitation, and inhibition in cat primary visual cortex. *J. Neurophysiol.* *84*, 909–926.
- Atallah, B.V., Bruns, W., Carandini, M., and Scanziani, M. (2012). Parvalbumin-expressing interneurons linearly transform cortical responses to visual stimuli. *Neuron* *73*, 159–170.
- Azouz, R., and Gray, C.M. (2000). Dynamic spike threshold reveals a mechanism for synaptic coincidence detection in cortical neurons in vivo. *Proc. Natl. Acad. Sci. USA* *97*, 8110–8115.
- Azouz, R., Gray, C.M., Nowak, L.G., and McCormick, D.A. (1997). Physiological properties of inhibitory interneurons in cat striate cortex. *Cereb. Cortex* *7*, 534–545.
- Ben-Yishai, R., Bar-Or, R.L., and Sompolinsky, H. (1995). Theory of orientation tuning in visual cortex. *Proc. Natl. Acad. Sci. USA* *92*, 3844–3848.
- Bernard, A., Lubbers, L.S., Tanis, K.Q., Luo, R., Podteleznikov, A.A., Finney, E.M., McWhorter, M.M., Serikawa, K., Lemon, T., Morgan, R., et al. (2012). Transcriptional architecture of the primate neocortex. *Neuron* *73*, 1083–1099.
- Blakemore, C., and Tobin, E.A. (1972). Lateral inhibition between orientation detectors in the cat's visual cortex. *Exp. Brain Res.* *15*, 439–440.
- Boudreau, C.E., and Ferster, D. (2005). Short-term depression in thalamocortical synapses of cat primary visual cortex. *J. Neurosci.* *25*, 7179–7190.
- Brodman, K. (1909). *Vergleichende Lokalisationlehre der Großhirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues* (Leipzig, Germany: Barth).
- Busse, L., Wade, A.R., and Carandini, M. (2009). Representation of concurrent stimuli by population activity in visual cortex. *Neuron* *64*, 931–942.
- Carandini, M. (2004). Amplification of trial-to-trial response variability by neurons in visual cortex. *PLoS Biol.* *2*, E264.
- Carandini, M., and Ferster, D. (2000). Membrane potential and firing rate in cat primary visual cortex. *J. Neurosci.* *20*, 470–484.
- Carandini, M., and Heeger, D.J. (1994). Summation and division by neurons in primate visual cortex. *Science* *264*, 1333–1336.
- Carandini, M., and Heeger, D.J. (2012). Normalization as a canonical neural computation. *Nat. Rev. Neurosci.* *13*, 51–62.
- Carandini, M., Heeger, D.J., and Movshon, J.A. (1997). Linearity and normalization in simple cells of the macaque primary visual cortex. *J. Neurosci.* *17*, 8621–8644.
- Carandini, M., Heeger, D.J., and Senn, W. (2002). A synaptic explanation of suppression in visual cortex. *J. Neurosci.* *22*, 10053–10065.
- Cardin, J.A., Palmer, L.A., and Contreras, D. (2007). Stimulus feature selectivity in excitatory and inhibitory neurons in primary visual cortex. *J. Neurosci.* *27*, 10333–10344.
- Chance, F.S., Nelson, S.B., and Abbott, L.F. (1999). Complex cells as cortically amplified simple cells. *Nat. Neurosci.* *2*, 277–282.
- Chapman, B., Zahs, K.R., and Stryker, M.P. (1991). 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.
- Chung, S., and Ferster, D. (1998). Strength and orientation tuning of the thalamic input to simple cells revealed by electrically evoked cortical suppression. *Neuron* *20*, 1177–1189.
- Churchland, M.M., Yu, B.M., Cunningham, J.P., Sugrue, L.P., Cohen, M.R., Corrado, G.S., Newsome, W.T., Clark, A.M., Hosseini, P., Scott, B.B., et al. (2010). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat. Neurosci.* *13*, 369–378.
- Dean, A.F., and Tolhurst, D.J. (1986). Factors influencing the temporal phase of response to bar and grating stimuli for simple cells in the cat striate cortex. *Exp. Brain Res.* *62*, 143–151.
- DeAngelis, G.C., Freeman, R.D., and Ohzawa, I. (1994). Length and width tuning of neurons in the cat's primary visual cortex. *J. Neurophysiol.* *71*, 347–374.
- DeAngelis, G.C., Robson, J.G., Ohzawa, I., and Freeman, R.D. (1992). Organization of suppression in receptive fields of neurons in cat visual cortex. *J. Neurophysiol.* *68*, 144–163.
- Douglas, R.J., and Martin, K.A.C. (1991). A functional microcircuit for cat visual cortex. *J. Physiol.* *440*, 735–769.
- Douglas, R.J., Martin, K.A.C., and Whitteridge, D. (1988). Selective responses of visual cortical cells do not depend on shunting inhibition. *Nature* *332*, 642–644.
- Douglas, R.J., Koch, C., Mahowald, M., Martin, K.A., and Suarez, H.H. (1995). Recurrent excitation in neocortical circuits. *Science* *269*, 981–985.
- Ferster, D. (1981). A comparison of binocular depth mechanisms in areas 17 and 18 of the cat visual cortex. *J. Physiol.* *311*, 623–655.
- Ferster, D. (1986). Orientation selectivity of synaptic potentials in neurons of cat primary visual cortex. *J. Neurosci.* *6*, 1284–1301.

- Ferster, D., Chung, S., and Wheat, H. (1996). Orientation selectivity of thalamic input to simple cells of cat visual cortex. *Nature* 380, 249–252.
- Finn, I.M., Priebe, N.J., and Ferster, D. (2007). The emergence of contrast-invariant orientation tuning in simple cells of cat visual cortex. *Neuron* 54, 137–152.
- Freeman, T.C., Durand, S., Kiper, D.C., and Carandini, M. (2002). Suppression without inhibition in visual cortex. *Neuron* 35, 759–771.
- Gardner, J.L., Anzai, A., Ohzawa, I., and Freeman, R.D. (1999). Linear and nonlinear contributions to orientation tuning of simple cells in the cat's striate cortex. *Vis. Neurosci.* 16, 1115–1121.
- Geisler, W.S., and Albrecht, D.G. (1992). Cortical neurons: isolation of contrast gain control. *Vision Res.* 32, 1409–1410.
- Gilbert, C.D., and Kelly, J.P. (1975). The projections of cells in different layers of the cat's visual cortex. *J. Comp. Neurol.* 163, 81–105.
- Gilbert, C.D., and Li, W. (2012). Adult visual cortical plasticity. *Neuron* 75, this issue, 250–264.
- Haider, B., and McCormick, D.A. (2009). Rapid neocortical dynamics: cellular and network mechanisms. *Neuron* 62, 171–189.
- Hansel, D., and Sompolinsky, H. (1996). Chaos and synchrony in a model of a hypercolumn in visual cortex. *J. Comput. Neurosci.* 3, 7–34.
- Hansel, D., and van Vreeswijk, C. (2002). How noise contributes to contrast invariance of orientation tuning in cat visual cortex. *J. Neurosci.* 22, 5118–5128.
- Hartline, H.K. (1949). Inhibition of activity of visual receptors by illuminating nearby retinal areas in the Limulus eye. *Fed. Proc.* 8, 69.
- Hartveit, E., and Heggelund, P. (1994). Response variability of single cells in the dorsal lateral geniculate nucleus of the cat. Comparison with retinal input and effect of brain stem stimulation. *J. Neurophysiol.* 72, 1278–1289.
- Hawken, M.J., Shapley, R.M., and Grosf, D.H. (1996). Temporal-frequency selectivity in monkey visual cortex. *Vis. Neurosci.* 13, 477–492.
- Heeger, D.J. (1992). Normalization of cell responses in cat striate cortex. *Vis. Neurosci.* 9, 181–197.
- Hirsch, J.A., Martinez, L.M., Pillai, C., Alonso, J.M., Wang, Q., and Sommer, F.T. (2003). Functionally distinct inhibitory neurons at the first stage of visual cortical processing. *Nat. Neurosci.* 6, 1300–1308.
- Hodgkin, A.L., and Huxley, A.F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol.* 117, 500–544.
- Holub, R.A., and Morton-Gibson, M. (1981). Response of visual cortical neurons of the cat to moving sinusoidal gratings: response-contrast functions and spatiotemporal interactions. *J. Neurophysiol.* 46, 1244–1259.
- Hubel, D.H., and Wiesel, T.N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.* 160, 106–154.
- Hubel, D.H., and Wiesel, T.N. (1965). Receptive fields and functional architecture in two non-striate visual areas (18 and 19) of the cat. *J. Neurophysiol.* 28, 229–289.
- Jia, H., Rochefort, N.L., Chen, X., and Konnerth, A. (2010). Dendritic organization of sensory input to cortical neurons in vivo. *Nature* 464, 1307–1312.
- Jin, J., Wang, Y., Swadlow, H.A., and Alonso, J.M. (2011). Population receptive fields of ON and OFF thalamic inputs to an orientation column in visual cortex. *Nat. Neurosci.* 14, 232–238.
- Jones, J.P., and Palmer, L.A. (1987). An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. *J. Neurophysiol.* 58, 1233–1258.
- Katzner, S., Busse, L., and Carandini, M. (2011). GABAA inhibition controls response gain in visual cortex. *J. Neurosci.* 31, 5931–5941.
- Kerlin, A.M., Andermann, M.L., Berezovskii, V.K., and Reid, R.C. (2010). Broadly tuned response properties of diverse inhibitory neuron subtypes in mouse visual cortex. *Neuron* 67, 858–871.
- Ko, H., Hofer, S.B., Pichler, B., Buchanan, K.A., Sjöström, P.J., and Mrsic-Flogel, T.D. (2011). Functional specificity of local synaptic connections in neocortical networks. *Nature* 473, 87–91.
- Kuffler, S.W. (1953). Discharge patterns and functional organization of mammalian retina. *J. Neurophysiol.* 16, 37–68.
- Kuhlman, S.J., Tring, E., and Trachtenberg, J.T. (2011). Fast-spiking interneurons have an initial orientation bias that is lost with vision. *Nat. Neurosci.* 14, 1121–1123.
- Lampl, I., Anderson, J.S., Gillespie, D.C., and Ferster, D. (2001). Prediction of orientation selectivity from receptive field architecture in simple cells of cat visual cortex. *Neuron* 30, 263–274.
- Latham, P.E., and Nirenberg, S. (2004). Computing and stability in cortical networks. *Neural Comput.* 16, 1385–1412.
- Li, B., Thompson, J.K., Duong, T., Peterson, M.R., and Freeman, R.D. (2006). Origins of cross-orientation suppression in the visual cortex. *J. Neurophysiol.* 96, 1755–1764.
- Li, Y.T., Ma, W.P., Pan, C.J., Zhang, L.I., and Tao, H.W. (2012). Broadening of cortical inhibition mediates developmental sharpening of orientation selectivity. *J. Neurosci.* 32, 3981–3991.
- MacEvoy, S.P., Tucker, T.R., and Fitzpatrick, D. (2009). A precise form of divisive suppression supports population coding in the primary visual cortex. *Nat. Neurosci.* 12, 637–645.
- Martinez, L.M., and Alonso, J.M. (2001). Construction of complex receptive fields in cat primary visual cortex. *Neuron* 32, 515–525.
- Martinez, L.M., Alonso, J.M., Reid, R.C., and Hirsch, J.A. (2002). Laminar processing of stimulus orientation in cat visual cortex. *J. Physiol.* 540, 321–333.
- Martinez, L.M., Wang, Q., Reid, R.C., Pillai, C., Alonso, J.M., Sommer, F.T., and Hirsch, J.A. (2005). Receptive field structure varies with layer in the primary visual cortex. *Nat. Neurosci.* 8, 372–379.
- Mohanty, D., Scholl, B., and Priebe, N.J. (2012). The accuracy of membrane potential reconstruction based on spiking receptive fields. *J. Neurophysiol.* 107, 2143–2153.
- Monier, C., Chavane, F., Baudot, P., Graham, L.J., and Frégnac, Y. (2003). Orientation and direction selectivity of synaptic inputs in visual cortical neurons: a diversity of combinations produces spike tuning. *Neuron* 37, 663–680.
- Morrone, M.C., Burr, D.C., and Maffei, L. (1982). Functional implications of cross-orientation inhibition of cortical visual cells. I. Neurophysiological evidence. *Proc. R. Soc. Lond. B Biol. Sci.* 216, 335–354.
- Morrone, M.C., Burr, D.C., and Speed, H.D. (1987). Cross-orientation inhibition in cat is GABA mediated. *Exp. Brain Res.* 67, 635–644.
- Niell, C.M., and Stryker, M.P. (2008). Highly selective receptive fields in mouse visual cortex. *J. Neurosci.* 28, 7520–7536.
- Orban, G.A., Hoffmann, K.P., and Duysens, J. (1985). Velocity selectivity in the cat visual system. I. Responses of LGN cells to moving bar stimuli: a comparison with cortical areas 17 and 18. *J. Neurophysiol.* 54, 1026–1049.
- Ozeki, H., Finn, I.M., Schaffer, E.S., Miller, K.D., and Ferster, D. (2009). Inhibitory stabilization of the cortical network underlies visual surround suppression. *Neuron* 62, 578–592.
- Priebe, N.J. (2008). The relationship between subthreshold and suprathreshold ocular dominance in cat primary visual cortex. *J. Neurosci.* 28, 8553–8559.
- Priebe, N.J., and Ferster, D. (2006). Mechanisms underlying cross-orientation suppression in cat visual cortex. *Nat. Neurosci.* 9, 552–561.
- Priebe, N.J., Mechler, F., Carandini, M., and Ferster, D. (2004). The contribution of spike threshold to the dichotomy of cortical simple and complex cells. *Nat. Neurosci.* 7, 1113–1122.
- Reid, R.C., and Alonso, J.M. (1995). Specificity of monosynaptic connections from thalamus to visual cortex. *Nature* 378, 281–284.

- Reid, R.C., Victor, J.D., and Shapley, R.M. (1992). Broadband temporal stimuli decrease the integration time of neurons in cat striate cortex. *Vis. Neurosci.* *9*, 39–45.
- Reig, R., Gallego, R., Nowak, L.G., and Sanchez-Vives, M.V. (2006). Impact of cortical network activity on short-term synaptic depression. *Cereb. Cortex* *16*, 688–695.
- Reynolds, J.H., and Heeger, D.J. (2009). The normalization model of attention. *Neuron* *67*, 168–185.
- Runyan, C.A., Schummers, J., Van Wart, A., Kuhlman, S.J., Wilson, N.R., Huang, Z.J., and Sur, M. (2010). Response features of parvalbumin-expressing interneurons suggest precise roles for subtypes of inhibition in visual cortex. *Neuron* *67*, 847–857.
- Rust, N.C., Schwartz, O., Movshon, J.A., and Simoncelli, E.P. (2005). Spatio-temporal elements of macaque v1 receptive fields. *Neuron* *46*, 945–956.
- Sadagopan, S., and Ferster, D. (2012). Feedforward origins of response variability underlying contrast invariant orientation tuning in cat visual cortex. *Neuron* *74*, 911–923.
- Saul, A.B., and Humphrey, A.L. (1990). Spatial and temporal response properties of lagged and nonlagged cells in cat lateral geniculate nucleus. *J. Neurophysiol.* *64*, 206–224.
- Sclar, G., and Freeman, R.D. (1982). Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. *Exp. Brain Res.* *46*, 457–461.
- Sestokas, A.K., and Lehmkuhle, S. (1988). Response variability of X- and Y-cells in the dorsal lateral geniculate nucleus of the cat. *J. Neurophysiol.* *59*, 317–325.
- Shadlen, M.N., and Newsome, W.T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.* *18*, 3870–3896.
- Sillito, A.M. (1975). The contribution of inhibitory mechanisms to the receptive field properties of neurones in the striate cortex of the cat. *J. Physiol.* *250*, 305–329.
- Skottun, B.C., Bradley, A., Sclar, G., Ohzawa, I., and Freeman, R.D. (1987). The effects of contrast on visual orientation and spatial frequency discrimination: a comparison of single cells and behavior. *J. Neurophysiol.* *57*, 773–786.
- Smith, M.A., Bair, W., and Movshon, J.A. (2006). Dynamics of suppression in macaque primary visual cortex. *J. Neurosci.* *26*, 4826–4834.
- Sohya, K., Kameyama, K., Yanagawa, Y., Obata, K., and Tsumoto, T. (2007). GABAergic neurons are less selective to stimulus orientation than excitatory neurons in layer II/III of visual cortex, as revealed by in vivo functional Ca²⁺-imaging in transgenic mice. *J. Neurosci.* *27*, 2145–2149.
- Somers, D.C., Nelson, S.B., and Sur, M. (1995). An emergent model of orientation selectivity in cat visual cortical simple cells. *J. Neurosci.* *15*, 5448–5465.
- Stern, E.A., Kincaid, A.E., and Wilson, C.J. (1997). Spontaneous subthreshold membrane potential fluctuations and action potential variability of rat corticostriatal and striatal neurons in vivo. *J. Neurophysiol.* *77*, 1697–1715.
- Tan, A.Y., Brown, B.D., Scholl, B., Mohanty, D., and Priebe, N.J. (2011). Orientation selectivity of synaptic input to neurons in mouse and cat primary visual cortex. *J. Neurosci.* *31*, 12339–12350.
- Tanaka, K. (1983). Cross-correlation analysis of geniculostriate neuronal relationships in cats. *J. Neurophysiol.* *49*, 1303–1318.
- Troyer, T.W., Krukowski, A.E., Priebe, N.J., and Miller, K.D. (1998). Contrast-invariant orientation tuning in cat visual cortex: thalamocortical input tuning and correlation-based intracortical connectivity. *J. Neurosci.* *18*, 5908–5927.
- Troyer, T.W., Krukowski, A.E., and Miller, K.D. (2002). LGN input to simple cells and contrast-invariant orientation tuning: an analysis. *J. Neurophysiol.* *87*, 2741–2752.
- Tsodyks, M.V., Skaggs, W.E., Sejnowski, T.J., and McNaughton, B.L. (1997). Paradoxical effects of external modulation of inhibitory interneurons. *J. Neurosci.* *17*, 4382–4388.
- van Vreeswijk, C., and Sompolinsky, H. (1998). Chaotic balanced state in a model of cortical circuits. *Neural Comput.* *10*, 1321–1371.
- Volgushev, M., Pernberg, J., and Eysel, U.T. (2000). Comparison of the selectivity of postsynaptic potentials and spike responses in cat visual cortex. *Eur. J. Neurosci.* *12*, 257–263.
- Walker, G.A., Ohzawa, I., and Freeman, R.D. (1998). Binocular cross-orientation suppression in the cat's striate cortex. *J. Neurophysiol.* *79*, 227–239.
- Wang, B.S., Sarnaik, R., and Cang, J. (2010). Critical period plasticity matches binocular orientation preference in the visual cortex. *Neuron* *65*, 246–256.
- Yu, Y., Shu, Y., and McCormick, D.A. (2008). Cortical action potential backpropagation explains spike threshold variability and rapid-onset kinetics. *J. Neurosci.* *28*, 7260–7272.