Contrast-Invariant Orientation Tuning in Cat Visual Cortex: Thalamocortical Input Tuning and Correlation-Based Intracortical Connectivity

Todd W. Troyer[*,2,6], Anton E. Krukowski[*,5,6], Nicholas J. Priebe[4,6] and Kenneth D. Miller[1,3–7]

[1]Depts. of Physiology, [2]Psychiatry and [3]Otolaryngology
[4]Neuroscience and [5]Biophysics Graduate Programs
[6]W.M. Keck Center for Integrative Neuroscience
[7]Sloan Center for Theoretical Neurobiology at UCSF
University of California, San Francisco, CA 94143-0444

*: **The first two authors contributed equally to this work.**

In Press, Journal of Neuroscience, May 14, 1998
This paper, with very slight changes in proof, appeared as
Journal of Neuroscience **18**, 5908-5927.

Address Correspondence To:
Kenneth Miller
Dept. of Physiology, UCSF
513 Parnassus
SF, CA 94143-0444
Email: ken@phy.ucsf.edu
WWW: http://www.keck.ucsf.edu/∼ken

Acknowledgements: We thank P. Bush, A.J. Doupe, M. Crair, S.G. Lisberger, T.J. Sejnowski, and D. Buonomono for helpful comments on the manuscript, and P. Bush for many helpful discussions. Supported by a Howard Hughes Medical Institute predoctoral fellowship (AEK), the McDonnell-Pew Program in Cognitive Neuroscience (TWT), an NEI predoctoral training grant (NJP), and by a Whitaker Foundation Biomedical Engineering Research Grant, the Searles Scholars' Program and the Lucille P. Markey Charitable Trust (KDM).

Abstract

The origin of orientation selectivity in visual cortical responses is a central problem for understanding cerebral cortical circuitry. In cats, many experiments suggest that orientation selectivity arises from the arrangement of LGN afferents to layer 4 simple cells. However, this explanation is not sufficient to account for the contrast invariance of orientation tuning.

To understand contrast invariance, we first characterize the input to cat simple cells generated by the oriented arrangement of LGN afferents. We demonstrate that it has two components: a spatial-phase-specific component (*i.e.*, one that depends on receptive field spatial phase), which is tuned for orientation; and a phase-nonspecific component, which is untuned. Both components grow with contrast.

Second, we show that a correlation-based intracortical circuit – in which connectivity between cell pairs is determined by the correlation of their LGN inputs – is sufficient to achieve well-tuned, contrast-invariant orientation tuning. This circuit generates both spatially opponent, "anti-phase" inhibition ("push-pull"), and spatially matched, "same-phase" excitation. The inhibition, if sufficiently strong, suppresses the untuned input component and sharpens responses to the tuned component at all contrasts. The excitation amplifies tuned responses. This circuit agrees with experimental evidence showing spatial opponency between, and similar orientation tuning of, the excitatory and inhibitory inputs received by a simple cell. Orientation tuning is primarily input driven, accounting for the observed invariance of tuning width after removal of intracortical synaptic input, as well as for the dependence of orientation tuning on stimulus spatial frequency.

The model differs from previous push-pull models in requiring dominant rather than balanced inhibition, and in predicting that a population of layer 4 inhibitory neurons should respond in a contrast-dependent manner to stimuli of all orientations, although their tuning width may be similar to that of excitatory neurons. The model demonstrates that fundamental response properties of cortical layer 4 can be explained by circuitry expected to develop under correlation-based rules of synaptic plasticity, and shows how such circuitry allows the cortex to distinguish stimulus intensity from stimulus form.

Thirty-five years ago, Hubel and Wiesel discovered that cells in cat primary visual cortex (V1) are tuned for the orientation of light/dark borders (Hubel and Wiesel, 1962). The inputs to V1 come from the lateral geniculate nucleus (LGN), whose cells are not significantly orientation selective (Hubel and Wiesel, 1961). The origin of orientation selectivity in visual cortex has been one of the most thoroughly investigated questions in neuroscience, and serves as a model problem for understanding how the cortex processes and represents information.

In cats, orientation selective responses appear in cortical layer 4. Cat layer 4 is composed of *simple cells* (Hubel and Wiesel, 1962; Gilbert, 1977; Bullier and Henry, 1979): cells with receptive fields (RFs) composed of oriented subregions, each giving exclusively ON or OFF responses (response to light onset/dark offset or light offset/dark onset). Hubel and Wiesel proposed that the orientation selectivity of these cells derives from an oriented arrangement of inputs from the LGN: ON-center LGN inputs have RF centers aligned over the simple cell's ON-subregions, and similarly for OFF-center inputs (Hubel and Wiesel, 1962). Such an input arrangement has been confirmed experimentally (Reid and Alonso, 1995; see also Tanaka, 1983). Because the total LGN input grows with increasing contrast for stimuli of all orientations, this model by itself is insufficient to explain the invariance of orientation tuning under change in stimulus contrast (Sclar and Freeman, 1982; Skottun et al., 1987). A threshold for spiking responses might narrow the tuning at any one contrast, but higher contrast would require a higher threshold to prevent broadening of tuning.

Two major approaches to achieving contrast invariance have been proposed. Many authors have suggested that responses in simple cells are approximately linear, *i.e.* the response can be predicted by linear summation of stimulus luminance (relative to background), weighted by the cell's RF (Movshon et al., 1978; Glezer et al., 1982; Tolhurst and Dean, 1990; Albrecht and Geisler, 1991; Heeger, 1992; Carandini and Heeger, 1994; Carandini et al., 1997a,b). Contrast change in such a model simply multiplies responses by a constant; contrast-invariant tuning follows automatically. It has been proposed that linear responses might be achieved by a balanced "push-pull" arrangement of inputs, in which an ON-subregion shows equal excitation ("push") to light stimuli as inhibition ("pull") to dark stimuli, and conversely for OFF-subregions (Glezer et al., 1982; Tolhurst and Dean, 1990; Carandini and Heeger, 1994; Carandini et al., 1997a,b). However, there are two problems with achieving linear response in an actual neural circuit. First, spike thresholds are nonzero, and therefore oriented stimuli that at low contrast give positive but sub-threshold input would yield spike responses at higher contrast. Second, at contrasts above about 5%, LGN responses increase more than they decrease, since spike rates cannot decrease below zero (*i.e.* responses "rectify"). This input nonlinearity alters the balance between push and pull.

Other authors have proposed that orientation tuning emerges from orientation-specific short-range excitation and longer-range inhibition in cortex (Ben-Yishai et al., 1995; Somers et al., 1995), despite evidence that in cat layer 4, excitation and inhibition show similar orientation tuning (Ferster, 1986). The width of orientation tuning in these models is an emergent property of intracortical circuitry, and so doesn't depend on the parameters of the stimulus, including stimulus contrast. These proposals appear inconsistent with the fact that orientation tuning widths in cats do depend on at least one stimulus parameter, the spatial frequency of sinusoidal grating stimuli (Vidyasagar and Sigüenza, 1985; Webster and De Valois, 1985; Jones et al., 1987; Hammond and Pomfrett, 1990).

We propose a new model for cat layer 4 cortical circuitry that yields contrast-invariant orientation tuning. Our model addresses two basic questions. First, what is the nature of the thalamocortical input to cortical simple cells? We assume thalamocortical connectivity can be modeled by a Gabor function: a two-dimensional Gaussian multiplied by a sinusoid (Jones et al., 1987; Reid and Alonso, 1995). The *spatial phase* of the sinusoid determines the location of ON and OFF subregions within the thalamocortical RF. Using a simple model of LGN responses, we show that the total LGN input has two components: a spatial-phase-specific component (a component that varies with the spatial phase of a cell's RF) that is tuned for orientation, and a phase-nonspecific component that is entirely untuned. Both components grow with contrast. Separating these input components helps clarify the debate over whether the LGN input to simple cells is well or poorly tuned. In response to drifting gratings, the phase-specific component corresponds to the temporally modulated input component, which Ferster et al. (1996) recently demonstrated to be tuned. However, the total input includes the phase-nonspecific, temporally unmodulated component; this should be untuned, and was not measured by Ferster et al. (1996). Separating the input components also clarifies the problem that cortical circuitry must solve to achieve contrast-invariant orientation tuning: eliminate the untuned component of the LGN input in a contrast-dependent manner, while extracting and sharpening the tuned component.

Second, what patterns of intracortical connectivity are sufficient to yield contrast-invariant orientation tuning? We arrive at a surprisingly simple answer: "correlation-based" connectivity yields contrast invariance. By correlation-based connectivity, we mean that intracortical connection strengths between two cells are fixed based on the correlation in their thalamocortical RFs. Thus, inhibitory connections occur between cells with anticorrelated RFs, while excitatory connections occur between cells with correlated RFs. The "anti-phase" inhibition eliminates the untuned input component and sharpens responses to the tuned component, while "same-phase" intracortical excitation amplifies the tuned response. As a result, our model achieves contrast-invariant tuning in the presence of positive thresholds and LGN rectification.

Our model uses a form of "push-pull" circuitry, but differs from other such models in that inhibition dominates rather than balances excitation and responses are not linear. Furthermore, we predict that a population of inhibitory neurons in cat layer 4 should respond in a contrast-dependent manner to stimuli of all orientations, although they may be tuned for orientation. The model has both developmental and functional implications for understanding the layer 4 cortical circuit, and suggests a general means of separating stimulus intensity (here represented by contrast) from stimulus form (represented by orientation).

An abstract of this work has appeared previously (Krukowski et al., 1996).

Methods

We study both a very simple ("conceptual") model and a more realistic ("computational") model. We first present the elements common to both, then present each model.

Elements Common to Both Conceptual and Computational Models

LGN Model. Our model was based on cat V1 at approximately 5^o eccentricity. LGN spatial RFs were center-surround difference of Gaussians, with cells responding either to light onset (ON cells) or light offset (OFF cells) in their RF centers. LGN spatial filter parameters ($\frac{17}{\sigma^2_{\text{center}}}e^{-x^2/\sigma^2_{\text{center}}} - \frac{16}{\sigma^2_{\text{surround}}}e^{-x^2/\sigma^2_{\text{surround}}}$; $\sigma_{\text{center}} = 15'$, $\sigma_{\text{surround}} = 1^o$) were taken from Peichl and Wassle (1979) and Linsenmeier et al. (1982). Firing rates in response to sinusoidal gratings were calculated on the assumption of linear rectified responses (unrectified firing rate was a sinusoid of the same temporal frequency as the stimulus; negative rates were then set to zero), using contrast response curves from Cheng et al. (1995) (Fig. 1). Assuming background firing rates of 10 Hz (ON cells) and 15 Hz (OFF cells) (modified from Kaplan et al. 1987, considering the lower mean luminance of 20 cd/m^2 used in Cheng et al. 1995), we calculated the sinusoidal amplitude that would lead to the reported values of the first harmonic (F1) after rectification. (Throughout, we will use F1 to denote the amplitude of the sinusoidal component at the frequency of the grating stimulus, although this value is twice as large as the value obtained using the Fourier transform normalized so that the F0 or DC component is the mean level (Skottun et al., 1991).) The amplitudes were then fit to $R = R_{max}(\frac{C^n}{C^n_{50}+C^n})$, where R is response amplitude and C is contrast (ON cells: $R_{max} = 53.0$ Hz, $n = 1.20$, $C_{50} = 13.3\%$; OFF cells: $R_{max} = 48.6$ Hz, $n = 1.29$, $C_{50} = 7.18\%$). LGN responses for gratings of non-optimal spatial frequencies were calculated by reducing modulation amplitudes by the factor predicted from the application of LGN spatial filters. ON and OFF cells had temporal phases offset by 180^o. To calculate the firing rates in response to moving bars, LGN cell spatiotemporal RFs were used. Temporal filters were taken from the central RF pixel in reverse correlation data from 100% contrast M-sequences (supplied by R.C. Reid); center and surround temporal filters were assumed equal for simplicity.

Cortical Receptive Fields. Cat cortical layer 4 simple cell RFs were modeled as Gabor functions (Fig. 2A). A Gabor function is a 2D Gaussian, here with peak value 1, multiplied by a sinusoid. Positive regions of the Gabor correspond to ON subregions and yield connections from ON-center LGN cells, negative regions correspond to OFF subregions and yield OFF-center inputs; the strength of the connection depends on the magnitude of the Gabor. The number of subregions is defined as the ratio of the width of the Gaussian envelope (at 5% of peak) to the width of a half-cycle of the sinusoid. The aspect ratio of a single subfield is defined as the ratio of the Gaussian envelope length to the sinusoid half-cycle width.

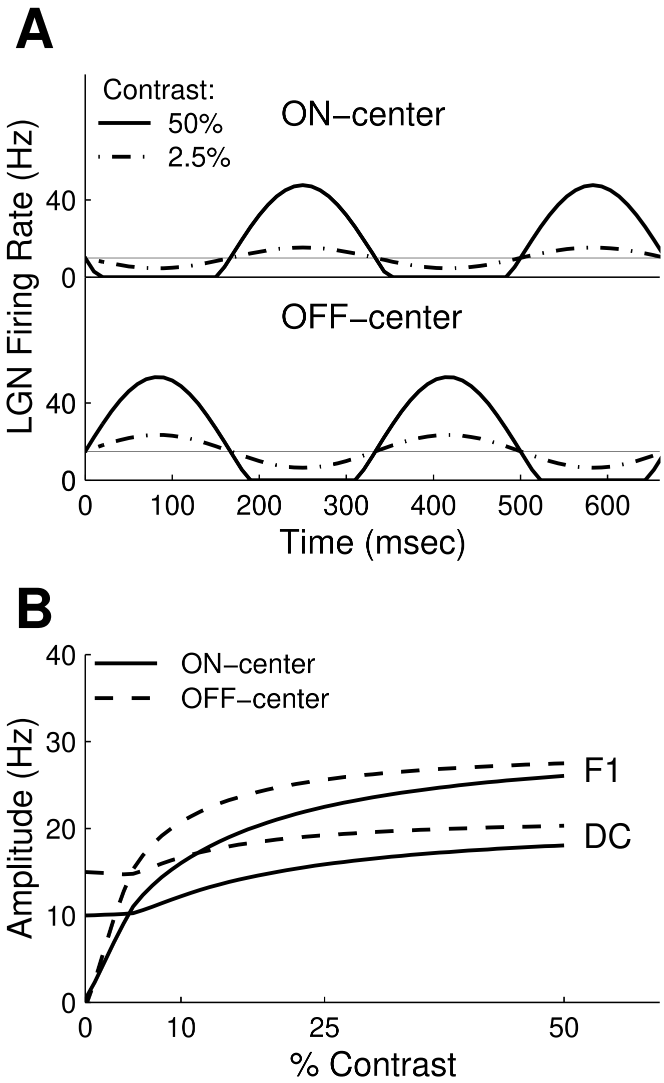

Figure 1. LGN cell responses to 3 Hz, 0.8 cycles/degree moving gratings. **A**. Instantaneous firing rate. Straight line: background. **B**. Contrast response functions. Upper: amplitude of first harmonic (F1); lower: mean (DC) firing rate. The mean rate increases at contrasts greater than 5%, due to rectification as seen in (A). Data modified from Cheng et al., (1995) (see Methods).

FIGURE 2 IS VERY DENSE AND CAN TAKE MANY MINUTES TO LOAD IN SOME POSTSCRIPT VIEWERS. THEREFORE, WE OMIT FIGURE 2 HERE, INSTEAD PLACING IT AND ITS LEGEND AT THE END OF THIS PAPER. WE ALSO INCLUDE THE LEGEND HERE:

Figure 2. Gabor-shaped cortical RFs. Lighter grays to white: positive values of Gabor function, corresponding to weights of ON-center LGN cells with centers at corresponding spatial positions; darker grays to black: negative values of Gabor function, corresponding to weights of OFF-center cells. **A**. A full Gabor function, used to determine LGN inputs to a cortical cell in the conceptual model. **B**. Typical LGN inputs to a cortical cell in the computational model, after probabilistic sampling from the full Gabor (see Methods). These receptive fields are 'typical'; different cortical cells may have different preferred orientations, spatial phase (relative locations of ON or OFF subregions), spatial location and, in the computational model, different outcomes of the probabilistic sampling. Spatial frequency of sinusoid in Gabor function is 0.8 cycles/degree.

Two sets of Gabor parameters were used. "Default" parameters were the mean values for simple cell physiological RFs reported in Jones and Palmer, (1987): 2.65 subregions and an aspect ratio of 4.54. (Care must be taken when comparing these numbers to other experimental estimates – *e.g.*, using a 10% cutoff for the Gaussian reduces these numbers by nearly 1/4). All RFs have 0.625^o half-cycle width, corresponding to a preferred spatial frequency of 0.8 cycles/degree, the approximate mean preferred spatial frequency of cortical cells at 5^o eccentricity (Movshon et al., 1978). Gaussian 5% envelope length and width are equal to 2.84^o and 1.65^o respectively. The measurements of Ferster et al. (1996) suggest that the net LGN input to a simple cell has broader orientation tuning than results from the default parameters (see Results). To model this broader tuning, we used a second set of Gabor parameters, identical to those above except that the Gaussian envelope was compressed by a factor of .7 in both length and width. This yields 1.85 subfields, a subfield aspect ratio of 3.18 and a 5% envelope length and width of 1.99^o and 1.15^o respectively.

Conceptual Model

To explore the basic concepts underlying our results, we constructed a "conceptual model" designed to be as simple as possible. The model contains two "rate-coded" cortical neurons, one excitatory and one inhibitory; the inhibitory cell inhibits the excitatory cell. The activity of each cell is represented by a scalar value corresponding to average firing rate. The LGN was modeled as a uniform sheet of cells, approximated as a dense lattice (lattice spacing $= .05^o$). The two cortical RFs were determined by Gabor RFs with identical Gaussian shape and location but having sinusoids of opposite spatial phase (thus, the inhibitory cell provides "anti-phase" inhibition).

For computational convenience in obtaining orientation tuning curves, rather than showing many gratings to one pair of cells, we showed one grating to many independent cell pairs. Thus, we constructed multiple pairs of cortical RFs with identical retinotopic positions, and with orientation and spatial phases spaced at 10^o and 20^o intervals respectively.

For each time step, we first calculated the LGN input to each RF by summing LGN firing rates, weighted by the Gabor function, to give the excitatory input $A(\theta, \phi)$ to the cell of orientation θ and phase ϕ. The net input to an excitatory cell with parameters (θ, ϕ) was the weighted sum $A(\theta, \phi) - wA(\theta, \phi + 180^o)$; $A(\theta, \phi + 180^o)$ is the LGN input to the (inhibitory cell) RF having the same orientation but opposite (180^o difference) spatial phase. The inhibitory gain factor, w, is unitless and represents the transformation from LGN excitatory current to inhibitory spike rate to inhibitory current in the excitatory cell. w is the only free cortical parameter in this model and controls the width of orientation tuning (Fig. 5). A match to experimental tuning widths of $\approx 20^o$ is given by $w = 1.5$ for default Gabor parameters (Figs. 4,7), and $w = 4.5$ for broadly tuned Gabor parameters.

The output rate of an excitatory cell was obtained by thresholding the net input, *i.e.* spike rate is proportional to $[A(\theta, \phi) - wA(\theta, \phi + 180^o) - \xi]^+$. For each set of Gabor parameters, the threshold ξ was set automatically according to the following algorithm (thus, ξ is not a free parameter). For a given level of inhibition w, orientation tuning curves were constructed by determining the peak input over a stimulus cycle for cells of each orientation preference, averaged over cells of all spatial phases. Such curves were obtained for gratings of 5%, 10%, 25% and 50% contrast. Linear interpolation was used to sample these tuning curves at 0.1^o intervals, and the orientation that gave the smallest variance in peak input across contrasts was determined (see Fig. 7). The threshold $\xi(w)$ was then set to the average across contrasts of the peak input for that orientation and level of inhibition. The excitatory cell's total response was determined by integrating its activity (calculated every 10 *msec*) over the course of one cycle. A single stimulus cycle was sufficient since the conceptual model is completely deterministic.

The inhibition level w_{best} that gave a best match to experimental tuning widths ($w = 1.5$ or $w = 4.5$ depending on Gabor parameters, as just described) was determined by constructing tuning curves for a range of w. Note that, by the procedure just described, each value of w yields a different threshold $\xi(w)$. To test the robustness of the model to variations in w (Fig. 5), for each set of Gabor parameters, we fixed ξ to the level appropriate for w_{best}, and calculated all responses using this fixed threshold.

Computational Model

The majority of simulations were carried out in a "computational model" incorporating details of cortical cells and maps.

Computational LGN Model. For the computational model, a realistically dense lattice of LGN cells was used. We restricted our attention to LGN X-cells, which dominate central cat V1 physiology (Ferster, 1990). At 5^o of eccentricity, $1mm^2 = 5^o \times 5^o$ of visual field (Bishop et al., 1962), and retinal ganglion X-cells (X-RGCs) have density $1000/mm^2$ (Peichl and Wassle, 1979), including both ON and OFF cells. We assume each X-LGN cell receives input from a single X-RGC, and each X-RGC projects to 4 X-LGN cells (as in Worgotter and Koch, 1991; this value is intermediate between values from Sherman, 1985 and Peters and Yilmaz, 1993). We thus use 7200 LGN cells to cover 6.8^o x 6.8^o of the visual field, arranged in 4 overlying sheets of ON cells (30 x 30 cells each) and 4 sheets of OFF cells (30 x 30), with ON and OFF lattices offset by 1/2 lattice spacing. After calculating LGN spike rates as above, spikes were produced in a random (Poisson) fashion: firing rates were converted into the probability of producing a spike in each simulated time step (.25 $msec$). To match data showing correlations among LGN cells with overlapping RFs (Alonso et al., 1996), overlaying cells had 25% correlations in their spike trains (each of 4 overlaying cells picked spikes with probability 1/4 from a common set of 4 Poisson processes). These correlations made no detectable difference in model behavior.

The connection strength to a given cortical cell from each LGN cell was determined by a repeated probabilistic sampling of the Gabor function describing the cortical RF (Fig. 2B). LGN synaptic strengths were equal to $(\bar{g}^{\rm ff}_{\rm ex}/n^{\rm ff}_{\rm pick}) \sum_{i=1}^{n^{\rm ff}_{\rm pick}} p_i$ where $n^{\rm ff}_{\rm pick} = 3$, $\bar{g}^{\rm ff}_{\rm ex} = 0.89$ nS, and $p_i = 1$ with probability determined by the absolute value of the Gabor function and equals 0 otherwise. The number of picks, $n^{\rm ff}_{\rm pick}$, determines the degree of sampling of the Gabor function: for $n^{\rm ff}_{\rm pick} \to \infty$, the RF becomes a perfect Gabor function. A typical sampled RF is shown in Fig. 2B. With this sampling, cortical cells received input from 125 ± 8 (mean $\pm$ standard deviation) LGN cells using the default Gabor. Using the more broadly tuned Gabor, cortical cells received input from 61 ± 5 LGN cells.

Cortical Model. Cortical cells were modeled as simple integrate-and-fire neurons as described in Troyer and Miller (1997a,1997b), with parameters matched to experimental data from McCormick et al. (1985). Excitatory cells were fitted to responses from regular spiking cells, and inhibitory cells to responses from fast spiking neurons. Briefly, each cell is a single compartment with a capacitance C, leak conductance $g_{\rm leak}$, resting potential $V_{\rm leak}$, and two synaptic conductances: fast (AMPA) excitation, $g_{\rm ex}$ (reversal potential $V_{\rm ex} = 0$ mV); and fast (GABA-A) inhibition, $g_{\rm in}$ ($V_{\rm in} = -70$ mV). Excitatory cells also have a spike-triggered adaptation conductance $g_{\rm adapt}$ ($V_{\rm adapt} = -90$ mV). Each time varying conductance, g, is modeled as a difference of exponentials: $g(t) = \sum_{t_j<t} \bar{g} \left(e^{-(t-t_j)/\tau^{\rm fall}} - e^{-(t-t_j)/\tau^{\rm rise}} \right)$, where the sum is over spike times t_j (presynaptic spike times for $g_{\rm ex}, g_{\rm in}$; postsynaptic for $g_{\rm adapt}$). When V crosses threshold, $V_{\rm thresh} = -52.5$ mV, synaptic events are triggered after a delay (randomly chosen for each spike from a uniform distribution, 0.25 $msec \leq t_{\rm delay} \leq 2.25$ $msec$),

adaptation is triggered (exc. cells only), and V is set to V_{reset} and held there for t_{refract}. V_{reset} was fit to the experimentally measured DC gain of cortical cells (the curve of firing rate vs. level of DC injected current; Troyer and Miller, 1997a, 1997b). All cells receive non-thalamocortical background excitatory input (Poisson with a mean rate of 5800 Hz and synaptic conductances equal to $\bar{g}_{\text{ex}}^{\text{bg}}$). The magnitude of this input was set to give low mean background firing rates for excitatory cells (.16 Hz) at default values of the parameters; identical background input was given to inhibitory cells and resulted in mean background firing rates of 12.2 Hz. Parameters are, for excitatory cells: $C = 500\ pF$, $g_{\text{leak}} = 25\ nS$, $V_{\text{leak}} = -73.6\ mV$, $V_{\text{reset}} = -56.5\ mV, t_{\text{refract}} = 1.5\ msec$; for inhibitory cells: $C = 214\ pF$, $g_{\text{leak}} = 18.0\ nS$, $V_{\text{leak}} = -81.6\ mV$, $V_{\text{reset}} = -57.8\ mV$, $t_{\text{refract}} = 1.0\ msec$; for conductances: $\tau_{\text{ex}}^{\text{rise}} = 0.25\ msec$, $\tau_{\text{ex}}^{\text{fall}} = 1.75\ msec$, $\tau_{\text{in}}^{\text{rise}} = 0.75\ msec$, $\tau_{\text{in}}^{\text{fall}} = 5.25\ msec$, $\tau_{\text{adapt}}^{\text{rise}} = 1\ msec$, $\tau_{\text{adapt}}^{\text{fall}} = 83.3\ msec$, $\bar{g}_{\text{adapt}} = 3\ nS$, $\bar{g}_{\text{ex}}^{\text{bg}} = 0.89\ nS$. $\bar{g}_{\text{ex}}^{\text{ctx}}$, $\bar{g}_{\text{ex}}^{\text{ff}}$, and $\bar{g}_{\text{in}}$ were free parameters and set as described below.

The model contains 1600 excitatory and 400 inhibitory layer 4 simple cells, representing a 2/3mm x 2/3mm patch of cortex and $.75^o x .75^o$ in visual angle ($0.9mm = 1^o$ of visual field at 5^o eccentricity (Tusa et al., 1978)). A 20 x 20 grid of inhibitory cells was interspersed within a 40 x 40 grid of excitatory neurons, with each inhibitory RF center aligned with every other excitatory cell. Gabor-shaped RFs were defined by three parameters in addition to those described above: preferred orientation, determined by an optically measured cortical map from cat V1 (provided by Michael Crair and Michael Stryker – shown in Fig. 8A); retinotopic position, progressing uniformly across the sheet; and spatial phase, assigned randomly to each cell (DeAngelis et al., 1992; Ghose et al., 1993).

The probability that any two cortical cells were connected depended on the correlation between their RFs. The following scheme was used for both excitatory and inhibitory connections. Raw correlation $c'(a, b)$ between RFs of cortical cells a, b is $c'(a, b) = \sum_{i,j \in \text{LGN}} w(a, i)w(b, j)c(i, j)$. Here, i, j are LGN cells, $w(a, i)$ and $w(b, j)$ are the thalamocortical weights from i to a and j to b, and $c(i, j)$ is the cross-correlation of the spatial RFs of i and j, where OFF spatial RFs are negative of ON. Correlation is then $c(a, b) = c'(a, b)/\sqrt{c'(a, a)c'(b, b)}$. A connectivity function $C(a, b)$ – roughly, the probability of a connection from a to b – is defined as $C(a, b) = [\text{sgn}(a)c(a, b)^{n_{\text{pow}}}]^+$ where $\text{sgn}(a) = 1$ if a is excitatory, -1 if a is inhibitory; $[x]^+ = x$, $x > 0$, $[x]^+ = 0$ otherwise. n_{pow} is a parameter that determines connectivity strength as a function of correlation. Smaller values of n_{pow} lead to broader connectivity and more intracortical connections per cell; larger values have the opposite effect (see Fig. 8B). At the default value, $n_{\text{pow}} = 6$, a cortical cell receives connections from 132 ± 38 (mean $\pm$ standard deviation) other cortical cells (80% from excitatory cells, 20% from inhibitory cells, on average). Just as the thalamocortical connections were sampled from the Gabor function, the intracortical connections were sampled from $C(a, b)$:

the strength of intracortical connection from a to b, $g(a,b)$, is $g(a,b) = (\bar{g}/n^{\mathrm{ctx}}_{\mathrm{pick}}) \sum_{i=1}^{n^{\mathrm{ctx}}_{\mathrm{pick}}} p_i$, where $p_i = 1$ with probability $C(a,b)$ ($\bar{g} = \bar{g}^{\mathrm{ctx}}_{\mathrm{ex}}$ or $\bar{g} = \bar{g}_{\mathrm{in}}$; $n^{\mathrm{ctx}}_{\mathrm{pick}} = 10$). As $n^{\mathrm{ctx}}_{\mathrm{pick}} \to \infty$, the connectivity becomes exactly $\bar{g}C(a,b)$.

The main parameters controlling model behavior were the total synaptic strength for each type of connection: thalamocortical (LGN); intracortical excitation (e$\rightarrow${e,i}), and intracortical inhibition onto excitatory cells (i$\rightarrow$e). The total synaptic strength is obtained by (i) assuming the cell is voltage clamped at threshold; (ii) for each synapse, integrating over time the synaptic current induced by one presynaptic spike; and (iii) summing over all synapses of the given type. Thus, total synaptic strength is expressed in units of nA $msec$. The parameters were chosen to satisfy various experimental constraints such as orientation tuning width. We used two different parameter sets: the "feedforward" set with LGN and intracortical inhibitory connections only, and the "full circuit" set which also included feedback intracortical excitation. For simplicity, inhibitory cells received only excitation; we have not yet explored the influence of inhibitory-to-inhibitory connections. For most simulations, the total intracortical excitatory synaptic strength onto each excitatory cell (e$\rightarrow$e connections) and onto each inhibitory cell (e$\rightarrow$i connections) was identical. Some simulations were run with intracortical excitatory connections onto excitatory cells only (e$\rightarrow$e, but no e$\rightarrow$i). After determining the pattern of synaptic strengths by probabilistic sampling, synaptic conductances were multiplicatively scaled so that the total conductance from each synaptic type received by each cell was set to its respective mean across cells. This avoids large differences in the amount of input to different cells resulting from the unequal representation of orientations in our spatially limited sample of an orientation map. For the "feedforward" parameter set, (Figs. 3,4,7), total synaptic strengths received by a cell from each type of connection were 10 nA $msec$ (LGN), and 3.75 nA $msec$ (i$\rightarrow$e), yielding mean values for unitary conductances of $\bar{g}^{\mathrm{ff}}_{\mathrm{ex}} = 2.1$ nS, $\bar{g}_{\mathrm{in}} = 8.3$ nS. For the "full circuit" parameter set, (Figs. 8-12), total synaptic strengths received by a cell from each type of connection were 5 nA $msec$ (LGN), 4.25 nA $msec$ (e$\rightarrow${e,i}), and 7.5 nA $msec$ (i$\rightarrow$e), yielding mean values for unitary conductances of $\bar{g}^{\mathrm{ctx}}_{\mathrm{ex}} = 2.0$ nS, $\bar{g}^{\mathrm{ff}}_{\mathrm{ex}} = 1.0$ nS, and $\bar{g}_{\mathrm{in}} = 16.6$ nS. Effects of varying these values were also explored (see Fig. 13). Note that we have realistic numbers of LGN cells but unrealistically small numbers of cortical cells, therefore intracortical connections are unrealistically strong relative to thalamocortical.

Simulations. A typical simulation consisted of 3 cycles of a 3 Hz sinusoidal grating. During each time step (.25 $msec$), values for time varying conductances were updated, and the membrane time constant and the equilibrium voltage for each cell were then calculated from the cell's conductances. Each cell's voltage was then adjusted according to an exponential decay. Finally, threshold crossings were detected and subsequent synaptic, adaptation and refractory events were registered. Simulations were written as C subroutines (mex files) in the MATLAB simulation environment. Initial conditions were determined by simulating 1 sec of model behavior at default parameter values and with LGN cells at background firing rates.

All orientation preferences are represented in the cortical network. Orientation tuning curves were constructed from the presentation of a single stimulus, by binning responses from all cells in the network according to their preferred orientation in 10^o bins. Most results used as a stimulus a grating oriented at 128^o. This orientation was chosen to avoid artifacts that might result from alignment of the stimulus with the axes of the LGN grid, but we saw no evidence of such behavior.

When displaying synaptic conductances and currents, we show "stimulus induced" curves in which we have subtracted the mean values of these conductances and currents at background. These mean values were determined by running "blank stimulus" trials in which LGN firing rates were unmodulated.

To reproduce the results of Nelson et al. (1994), we ran simulations in which the inhibition and adaptation currents were blocked in a single cell (Fig. 12). To accomplish this in a computationally convenient way, we ran a single simulation without any blockade, but monitored the behavior of an additional "blocked cell" for each cell in the network. The blocked cell made no connections. It received identical excitatory input as its unblocked partner cell, but had no inhibitory or adaptation current and was injected with sufficient hyperpolarizing current to bring the background firing rates back to normal. Thus each blocked cell received input from a network in which all other cells were normal (unblocked), but did not itself affect any other cells in the network. Under the assumption that altering a single cell does not affect network behavior, this method allows us to simulate numerous experiments in which one cell undergoes intracellular inhibitory blockade.

Results

Modeling Approach

We pursued two parallel approaches to modeling contrast-invariant orientation tuning. In order to explore the basic ideas underlying such tuning, we constructed a "conceptual" model, designed to be as simple as possible. This model considered two cortical simple cells, one excitatory and one inhibitory, with a monosynaptic connection from inhibitory to excitatory. The RFs of the two cells had identical position and preferred orientation, but opposite spatial phase (see Methods). The neurons were "rate-coded": the average firing rate of each cell was determined by a linear thresholding operation applied to the weighted sum of input cell firing rates. For simplicity, the inhibitory threshold was set to zero (*i.e.* the inhibitory cell's response was a linear function of its input). The excitatory cell's threshold was set automatically to the level that best produced contrast-invariant tuning for contrasts of 5% and above (see Methods). Therefore, after determining the structure of the cortical receptive fields, the conceptual model had only a single free parameter: the strength of intracortical inhibition relative to the strength of thalamocortical excitation.

In order to study robustness of our ideas to the complexity of real cortical circuits, we also constructed a "computational" model that incorporated known details of cortical cells and maps. The cortical component of this model consisted of 1600 excitatory and 400 inhibitory layer 4 simple cells, arranged in a 2/3mm x 2/3mm cortical sheet. Preferred orientations were determined by a measured V1 map, and intrinsic connectivity was determined probabilistically based on correlations in input RFs. Excitatory and inhibitory cells were modeled as conductance-based integrate-and-fire neurons, with parameters matched to those measured in cortical regular-spiking and fast-spiking cells respectively, including a spike-rate adaptation current in the excitatory cells (McCormick et al., 1985; Troyer and Miller, 1997a, 1997b) (details in Methods). We considered only the effects of fast synaptic conductances (AMPA and GABA-A); the role of slow conductances (NMDA and GABA-B) will be explored in future work.

LGN Input

We focused our research on the response to full-field sinusoidal gratings, as these are the only stimuli for which contrast dependence of orientation tuning has been studied (Sclar and Freeman, 1982; Skottun et al., 1987). Our model was based on cat V1 at approximately 5^o eccentricity. Circularly symmetric, center-surround LGN spatial receptive fields were used (Peichl and Wassle, 1979; Linsenmeier et al., 1982), and LGN firing rates were determined as rectified linear filterings of the input luminance using experimentally measured contrast gain curves (see Methods and Fig. 1B) (Peichl and Wassle, 1979; Cheng et al., 1995). To determine whether our model would yield well-tuned responses to transient stimuli, we also modeled responses to moving bars.

LGN cells responded to sinusoidal grating stimuli with a sinusoidal modulation in firing rate (Fig. 1A). The temporal responses of ON-center and OFF-center cells with spatially overlapping RFs were 180^o out of phase. Increasing stimulus contrast resulted in a larger modulation of firing rate. At contrasts above about 5%, the spike rate modulation exceeds the background firing rate. For these contrasts, responses are no longer purely sinusoidal, because spike rate cannot be negative (Fig. 1A, solid lines); that is, LGN responses rectify. Once responses rectify, mean ("DC") firing rates increase with increasing contrast (DC curves in Fig. 1B), because peak firing rates continue to increase while minimal firing rates cannot decrease below zero. This contrast-dependent increase in mean LGN firing rates has important consequences for contrast-invariant orientation tuning that will be discussed in detail below.

The oriented arrangement of LGN inputs to simple cell RF subregions was modeled using a Gabor function – a two dimensional Gaussian multiplied by a sinusoid (Fig. 2A). In the conceptual model, the Gabor function directly determined the weights of geniculocortical connections: positive values corresponded to the weights of ON-center inputs, negative values to the weights of OFF-center inputs (Fig. 2A). In the computational model, geniculocortical synaptic strengths were determined by probabilistic sampling of the Gabor function from a realistically dense lattice of LGN cells (Fig. 2B).

We considered two different sets of Gabor parameters to describe geniculocortical connections. The first set was matched to RF parameters taken from physiological measurements of cat simple cells (Jones and Palmer, 1987). The use of the Jones and Palmer parameters as a measure of LGN connectivity in simple cells is based on the experiments of Reid and Alonso (1995), which show that physiological RF parameters at least roughly correspond to the pattern of geniculocortical connections in cat layer 4. These will be used as our "default" parameters. We also consider a second set of parameters representing more broadly tuned LGN input, for several reasons. If cortical circuitry plays a significant role in sharpening simple cell orientation tuning, then the LGN input to a cell would have broader tuning than the cell's responses. Furthermore, the parameters of Jones and Palmer (1987) represent an average of simple cells from all layers, whereas layer 4 cells may on average be more broadly tuned for orientation than other layers (Tolhurst and Thompson, 1981). We base our more broadly tuned parameter set on the experiments of Ferster et al. (1996), who cooled the cortex to largely eliminate cortical inputs. Using intracellular electrodes, they then measured the direct LGN input for gratings presented at 30^o intervals. The tuning of this input was quantified by measuring the first harmonic (F1) of the voltage response, as a function of stimulus orientation. Although orientation was sampled only coarsely, the figures presented in Ferster et al. (1996) show average orientation tuning halfwidth at half height (HWHH) of approximately 35^o. This is significantly broader than the input F1 tuning under our default Gabor parameters, which we find to be 24^o. To mimic the broader tuning observed by Ferster et al., we artificially shrunk the default RFs by a factor of 0.7, leaving the width of each subregion unchanged. This resulted in an input F1 tuning width of 34.8^o.

In the conceptual model, the excitatory and inhibitory cells had identical Gabor RFs except that their sinusoids were 180^o out of phase. In the computational model, a distribution of receptive fields was obtained from variations in three parameters: preferred orientation, determined by a measured cortical map (Fig. 8A); retinotopic position, progressing uniformly across the sheet; and spatial phase, assigned randomly to each cell (DeAngelis et al., 1992).

Tuning of the LGN Input to a Simple Cell

At the preferred orientation, the bright and dark portions of a sinusoidal grating stimulus align with the cortical cell's ON and OFF subregions simultaneously. Thus, all of the cortical cell's LGN inputs fire relatively synchronously and the temporal modulation of this input is large (Fig. 3A). At the null orientation, the inputs are stimulated asynchronously, so the temporal modulation of the total input is small. Note that the mean rate of LGN input does not depend on stimulus orientation. This follows from the assumption that LGN cells are untuned for orientation: since the mean LGN input received by a simple cell is the (weighted) sum of the mean rates of the LGN cells projecting to it, this mean input must also be untuned for orientation (Ferster, 1987). Therefore, only the temporally modulated component of the LGN input is orientation tuned.

The untuned mean input presents the primary problem for a purely thalamocortical explanation of contrast-invariant orientation tuning. Mean LGN firing rates increase with increasing contrast due to LGN rectification (Fig. 1). This contrast-dependent increase in firing rate is sufficiently large that the *mean* LGN input at the null orientation at high contrasts exceeds the *peak* LGN input at the preferred orientation at low contrast (Fig. 3). No single spiking threshold level can yield well-tuned responses for stimuli of all contrasts.

Therefore, to achieve contrast-invariant orientation tuning in response to sinusoidal gratings, the cortex must cancel the untuned, mean input component in a contrast-dependent manner, while extracting the tuned, modulated component. We will show that this decomposition of the input into a tuned and an untuned component generalizes to stimuli such as flashed and moving bars.

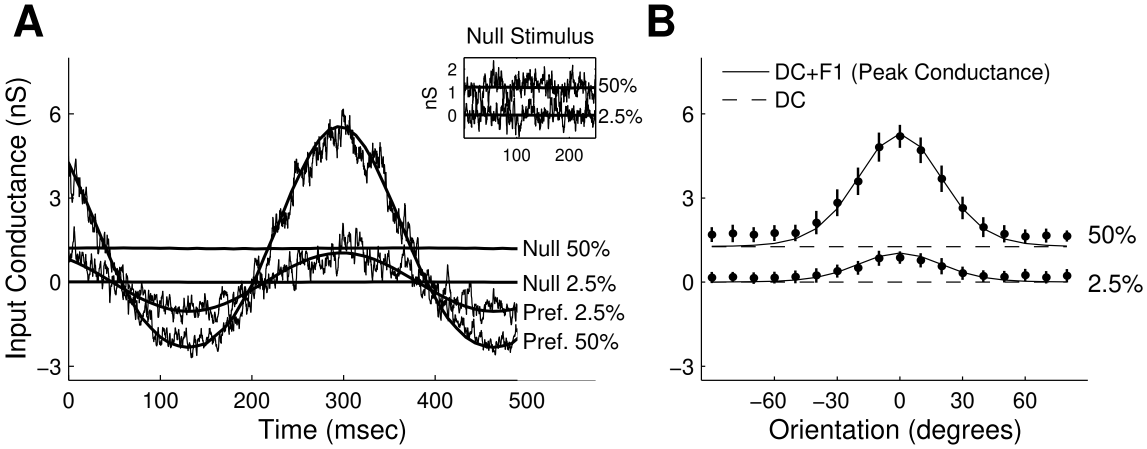

Figure 3. Tuning of total LGN input. **A**. Input to cortical cells in response to gratings at the preferred and null (orthogonal to preferred) orientations. Upper traces: high (50%) contrast; lower traces: low (2.5%) contrast. Curved traces show input in response to preferred orientation: black traces: average input (40 presentations) from computational model, using a sampled Gabor RF (as in 2B); gray curves: input for conceptual model, using connections from the full Gabor function (2A). Gray straight lines show response in the conceptual model to a stimulus at the null orientation; in inset, these lines are repeated and compared to average input to null stimuli in computational model (black traces). Note that input to null stimulus at 50% contrast typically exceeds peak input to preferred stimulus at 2.5% contrast. Agreement of the two models for both preferred and null stimuli indicates RF sampling and Poisson firing of LGN inputs have little effect. **B**. Tuning of mean (dashed lines) and mean plus first harmonic (solid lines) of thalamic input conductance. Lines show results from the conceptual model; solid circles show results from the computational model; error bars represent ± 1 std. dev. Sum of mean plus first harmonic represents peak input during a cycle of the grating stimulus. Note that mean input is untuned for orientation, and mean input at high contrasts exceeds peak input to preferred orientation at low contrasts. Thus, although the first harmonic is well tuned, no single spike threshold can give tuned responses at both high and low contrasts. In this and subsequent figures showing orientation tunings, cells are grouped by preferred orientation in 10^o bins, and orientation axis represents difference of stimulus orientation from preferred. The plot shows line connecting means for each bin; error bars are over cells within a bin. See Methods for more details.

Anti-Phase Inhibition Can Achieve Contrast-Invariant Orientation Tuning

The main result of this paper is to demonstrate that correlation-based intracortical inhibition can achieve contrast-invariant orientation tuning (the effects of correlation-based intracortical excitation will be considered below). By correlation-based inhibition, we mean that the probability of a connection from an inhibitory cell to an excitatory cell is an increasing function of the degree of anticorrelation between their RFs, *i.e.* the strongest inhibitory connections are made between cells with the most anti-correlated RFs (see Methods). This implies that an excitatory cell receives the strongest inhibition from inhibitory cells with identical Hubel-Wiesel RFs but of opposite spatial phase. We will call such an inhibitory neuron the cell's "anti-phase partner". (By "spatial phase" of an RF, we refer to absolute position in visual space of the ON or OFF subregions, rather than to their position relative to each cell's Gabor function; thus, two RFs have "opposite spatial phase" if the ON-subregions of one tend to overlap the OFF-subregions of the other in visual space). The existence of such "spatially opponent" or "anti-phase" inhibition in cat layer 4 is well-supported experimentally: at ON locations, where a light stimulus evokes excitation (EPSPs), dark stimuli evoke inhibition (IPSPs), and vice versa for OFF locations (Palmer and Davis, 1981; Ferster, 1988; Hirsch et al., 1995). Note that, since simple cells with orthogonal orientation preference have weakly- or un-correlated RFs, correlation-based connectivity results in little or no inhibition from cells with orthogonal tuning. Instead, inhibition comes from cells with similar preferred orientations.

Our model is not a developmental model: we first determined the pattern of LGN input to cortical cells, and then fixed the pattern of intracortical connections according to the above correlation-based rule. However, this pattern of inhibition would be expected to arise from a Hebb-type synaptic modification rule, generalized to apply to inhibitory synapses. Such a rule states that synaptic strengths grow more negative (more strongly inhibitory) when pre- and postsynaptic firings are anticorrelated; or equivalently, that synapses strengthen when they are effective, *i.e.* when the inhibitory presynaptic cell is active while the postsynaptic cell is inactive. Such generalization of Hebb-type learning rules to inhibitory synapses is only a hypothesis; plasticity of inhibitory synapses is not well understood (but see Komatsu, 1996). (This intracortical connectivity could also emerge without inhibitory synaptic plasticity: in models in which only thalamocortical synapses undergo correlation-based plasticity, presence of a fixed inhibitory connection from one cortical cell to another tends to cause the two to develop anticorrelated thalamocortical RFs (Miller, 1994).)

The sufficiency of correlation-based inhibition for contrast-invariant tuning is demonstrated in Fig. 4, which shows tuning curves for both the computational and conceptual models, for gratings of 2.5%, 5%, 10%, 25%, and 50% contrast. Both models display

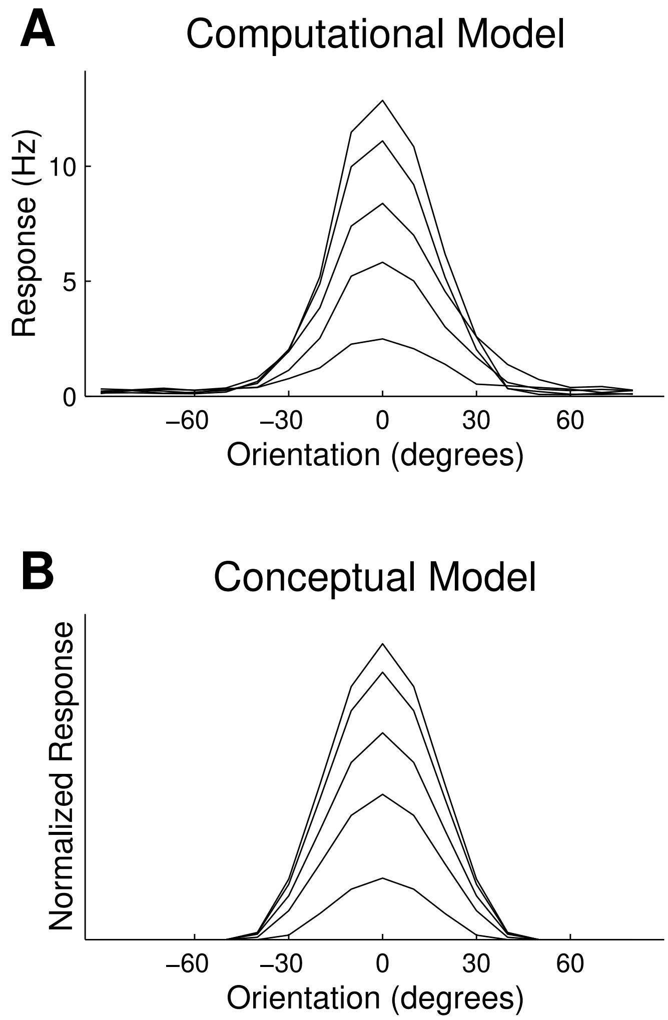

Figure 4. Contrast-invariant tuning. Response vs. orientation for gratings of 2.5%, 5%, 10%, 25% and 50% contrast. **A**. Computational model. **B**. Conceptual model. Both models yield contrast-invariant tuning at 5% contrast and above.

contrast-invariant orientation tuning above 5% contrast. By choosing the appropriate level of inhibition, both models were able to match experimental estimates of mean orientation tuning width for simple cells. For example, Heggelund and Albus (1978) report that simple cells have mean tuning width (HWHH) of 19.5^o. Model tuning widths (HWHH) above 5% contrast were between 18.7^o and 20.8^o for both the computational and conceptual models.

The width of the tuning is largely determined by the strength of the inhibition and the tuning of the LGN input. Fig. 5 shows cortical tuning halfwidths, using either narrowly or broadly tuned LGN inputs, for a variety of levels of inhibition. Tuning narrows with stronger inhibition, but remains contrast-invariant above 5% contrast. Tuning to a long moving bar (width 0.62^o, velocity $3.75^o/sec$) is slightly broader but shows identical sharpening with increasing levels of inhibition: filled circles show bar tuning at 50% contrast for the narrowly tuned LGN inputs.

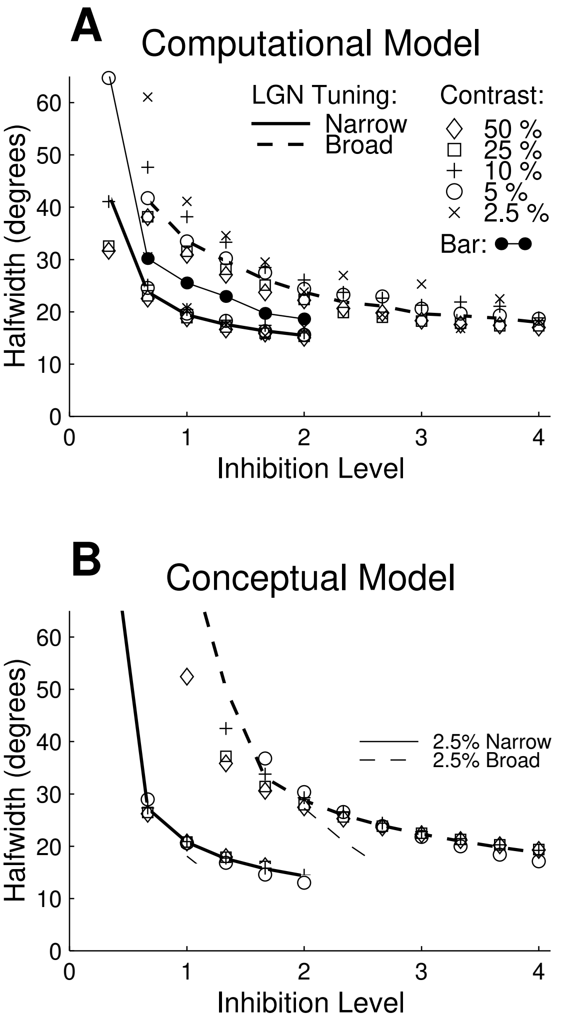

Figure 5. Increasing inhibition leads to sharper tuning. Tuning halfwidth at half height (HWHH) vs. level of inhibition for gratings of 2.5%, 5%, 10%, 25% and 50% contrast. Thick solid (lower) curve: mean tuning HWHH above 5% for RFs with large subfield aspect ratios and narrow LGN tuning (matched to data from Jones and Palmer, 1987); Thick dashed (upper) curve: mean tuning HWHH for RFs with small subfield aspect ratios and broad LGN tuning (matched to data from Ferster et al. 1996). Level of inhibition is normalized so that 1 is the level that produces physiological halfwidths for narrow LGN input (Fig. 4). Overlapping symbols indicate contrast- invariance. Tuning gradually sharpens with increased levels of inhibition. **A**. Computational model. **B**. Conceptual model. In conceptual model, tuning narrows slightly at 5% contrast for large levels of inhibition due to the fact that spike threshold is optimized for default parameters, *i.e.* inhibition level of 1 (see Methods). Responses to 2.5% contrast gratings at high inhibition levels for both narrow (solid) and broad (dashed) LGN tuning are shown using thin lines. At very low contrast, conceptual model predicts much narrower tuning.

For higher levels of inhibition and broadly tuned input, tuning at 5% contrast narrows slightly in the conceptual model. This is due to the fact that spike threshold was optimized for the default level of inhibition (see Methods), and could be corrected if spike threshold were separately optimized for each set of parameters. At 2.5% contrast and high levels of inhibition (Fig. 5B, thin lines), the conceptual model predicts much narrower tuning, for reasons that are more general (see below).

The conceptual and computational models yield qualitatively similar results. Simple additions to the conceptual model led to progressively closer quantitative matches to computational model behavior. A significantly improved match was obtained by adding inhibitory thresholds and using correlation-based inhibitory connectivity from cells with a range of RF properties (rather than from only the single cell with precisely opposite spatial phase). Using simulated synaptic noise (and hence changing the threshold linear function to a smoother function near spike threshold) led to an even closer match between the models, and nearly eliminates the difference in responses to 2.5% contrast gratings (see below). However, incorporation of these features required additional unconstrained parameters and we began to lose the simplicity that was the strength of the conceptual model. Therefore, the results of these investigations are not further reported.

The behavior of our correlation-based model is presented below, in three steps. First, we analyze the reasons why anti-phase inhibition achieves contrast-invariant tuning, using the simple conceptual version of the model. Second, we incorporate correlation-based intracortical excitation into the computational model, and present results from this completed model. Finally, we explore the robustness of this computational model to variations in the key parameters controlling model behavior. A schematic representing the behavior of both models is shown in Fig. 6, and will be referred to throughout the text.

Conceptual Model: Anti-Phase Inhibition Cancels the Untuned Component of the Input

Recall that the main obstacle to achieving contrast-invariant tuning is the untuned component of the LGN input, which increases with contrast due to the rectification of LGN responses at higher contrasts. The ability of anti-phase inhibition to overcome this problem is most easily demonstrated in the context of the two-cell conceptual model. Here we introduce the term "feedforward", by which we mean input from LGN to cortical cells not mediated by cortical excitatory cells. Thus, the geniculocortical input represents feedforward excitation, while the pathway from LGN to cortical inhibitory cell to cortical excitatory cell represents feedforward inhibition.

Suppose an excitatory simple cell receives total input A^e from the LGN, while its in-

FIGURE 6 IS VERY DENSE AND CAN TAKE MANY MINUTES TO LOAD IN SOME POSTSCRIPT VIEWERS. THEREFORE, WE OMIT FIGURE 6 HERE, INSTEAD PLACING IT AND ITS LEGEND AT THE END OF THIS PAPER. WE ALSO INCLUDE THE LEGEND HERE:

Figure 6. Behavior of model using correlation-based connectivity. Schematic representing behavior of the model in response to preferred (**A**) and null (**B**) stimuli. The excitatory cell described in the Results is in the upper left; its inhibitory "anti-phase partner" is in the lower right. E: excitatory cells; I: inhibitory cells. Solid lines: excitation and depolarization; open lines: inhibition and hyperpolarization. Line thickness and size of RF icon represent magnitude of activity. Dashed lines represent correlation-based excitation, which is included in the complete computational model only (Figs. 8-11). Some simulations were performed without cortical excitatory projections onto inhibitory neurons (gray dashed lines), but this did not substantially affect network behavior (see Fig. 13B).

hibitory anti-phase partner receives LGN input A^i. Assuming for simplicity that inhibitory cell response is linear, the total feedforward input to the excitatory cell is $A^e - wA^i$, where $w > 1$ is the total strength of the inhibitory synaptic connection multiplied by the gain of the inhibitory cells. During the peak response to the preferred orientation, LGN excitation A^e is large, while the anti-phase inhibition wA^i is weak (Figs. 6A,7A-top). Thus, the cell gives a strong response. At the null orientation, cells at all spatial phases are receiving an intermediate level of feedforward excitation $A^e \approx A^i$, and the inhibition $wA^i > A^e$ is sufficient to prevent excitatory cell spiking (Figs. 6B,7A-bottom). Since A^i and A^e both rise with contrast at the same rate, the dominance of inhibition over excitation is maintained for null stimuli of all contrasts.

The ability of anti-phase inhibition to achieve contrast-invariant tuning for a wide variety of stimuli can be best understood by dividing the LGN input into two components: the phase-nonspecific component $A_{non} = (A^e + A^i)/2$ – the average of the input to the cell and to its anti-phase partner – and the remaining phase-specific component, $A_{spec} = (A^e - A^i)/2$. The total input to the cell, $A^e - wA^i$, can then be rewritten $(1 - w)A_{non} + (1 + w)A_{spec}$. Thus, anti-phase inhibition acts to eliminate the phase-nonspecific component of the LGN input, while amplifying the phase-specific component. For all of the commonly presented oriented stimuli (moving or flashed bars, flashed or counterphased gratings), Hubel-Wiesel RFs yield a phase-specific component tuned for orientation, and a phase-nonspecific component that is nearly or completely untuned. Thus, the effectiveness of the anti-phase model in achieving contrast-invariant tuning generalizes across stimuli.

This can by summarized by noting that the schematic circuit (Fig. 6) acts as a "differential phase filter": with inhibition sufficiently large, any stimulus that gives similar excitation to

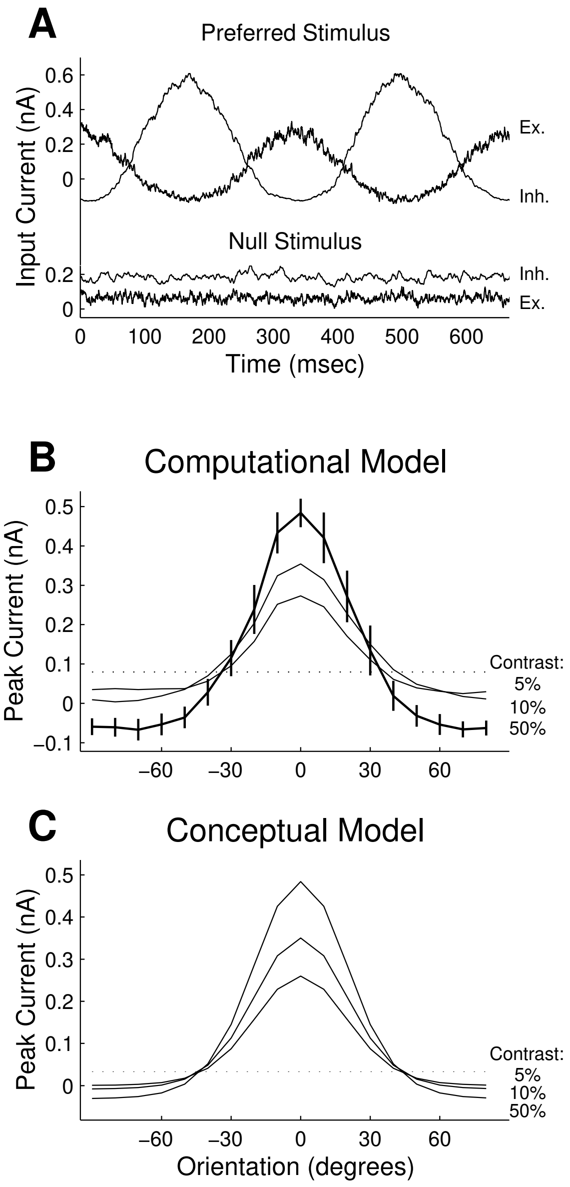

Figure 7. Inputs to a cortical cell given anti-phase inhibition (inputs shown relative to background). **A**. Averaged computational model responses (40 presentations) to 50% contrast gratings. Excitatory LGN input marked 'Ex'; intracortical inhibitory input marked 'Inh'. To compare excitatory and inhibitory inputs, synaptic conductances were converted to currents obtained if the cell was voltage clamped at threshold. **B**. Peak synaptic current vs. orientation for computational model. Responses are to single presentations of 50%, 10% and 5% contrast gratings at 128^o. Peak current is the first harmonic (F1) plus the mean (DC) of the stimulus-induced current (including excitation and inhibition). Error bars for 50% contrast are $\pm$ 1 std. dev. Dotted line shows approximate threshold level that would lead to contrast-invariant tuning; actual threshold in computational model is determined independently from *in vitro* data (see Methods). **C**. Peak synaptic current vs. orientation for conceptual model. Since there is no noise, true peak current is shown. Thin line shows automatically selected threshold (see Methods). For both models, mean input decreases and modulation increases with contrast. Thresholds near the crossover point of net input tuning curves result in sharp, contrast-invariant tuning.

each of two opposite phases will cause more inhibition than excitation in excitatory cells and hence will be "filtered out". Only stimuli that predominantly excite one phase and not its opposite can "pass" through this "filter" and cause the excitatory cells to fire. The only stimuli that can accomplish this are stimuli near the preferred orientation; stimuli far from the preferred will give similar input to both phases. This argument applies to any type of oriented stimulus.

Conceptual Model: Dominant Anti-Phase Inhibition Provides a Contrast-Dependent Effective Threshold

Although the most important effect of anti-phase inhibition is to eliminate the phase-nonspecific component of the LGN input, this is not sufficient to achieve contrast-invariant tuning. This can be seen by setting $w = 1$, thereby causing $(1 - w)A_{non} = 0$. In this case, contrast invariance can only be achieved if spike threshold is negligible, *i.e.* if any positive input leads to spiking. Otherwise, orientations that at low contrast give positive but sub-threshold phase-specific input would yield spike responses at higher contrast, because A_{spec} grows with contrast; thus, orientation tuning would broaden with contrast.

This problem is remedied by including relatively strong inhibition, ($w > 1$). Then the phase-nonspecific component $(1 - w)A_{non}$ has a net inhibitory influence that increases with contrast. Since the phase-nonspecific input is untuned for orientation, it serves as a "plateau" – an input identical for stimuli of all orientations – to which the orientation-tuned, phase-specific component is added. The distance from this plateau to the cell's spike threshold can be thought of as a contrast-dependent effective threshold for the tuned input component (Bonds, 1989; Ben-Yishai et al., 1995). With $w > 1$, this plateau is inhibitory and moves further from spike threshold with increasing contrast (Fig. 7B,C). By "pulling down" the tuned component, so that only a portion of it is above the spike threshold, this inhibition serves to sharpen the spiking orientation tuning relative to the tuning of the phase-specific input. If spike threshold falls near the crossover point of the net input tuning curves for varying contrasts (Dotted lines in Fig. 7B,C), this inhibition sharpens the feedforward input in a contrast-invariant manner.

In the conceptual model, spike threshold for excitatory cells was automatically set at this crossover point in the input current (see Methods). Somewhat surprisingly, we have found that in the computational model, simply using a physiologically-based spiking neuron model (Troyer and Miller, 1997a, 1997b) was adequate to robustly attain contrast-invariant tuning; no parameter adjustments were required. One possible explanation is that synaptic noise "smears out" spike threshold, making it relatively easy to match "threshold" with the crossover. Also, with inhibition dominant, the orientation tuning curves cross one another where input changes rapidly as a function of orientation, so moderate changes in threshold should make little difference in tuning. Simulations with the conceptual model show that moving threshold by as much as 10% of the peak-to-peak variation in the input driven by 5% gratings changes tuning by less than 1.5^o.

At very low contrasts, the conceptual model predicts that orientation tuning will narrow. Below about 5% contrast, LGN responses do not rectify and therefore the plateau, $(1 - w)A_{non}$, does not change with contrast. Orientation tuning narrows with further decreases in contrast (Fig. 5B), because the tuned input component is reduced while the nonzero "effective threshold" is left unchanged. It is unclear whether one could expect to see narrower tuning in the experimental data. As mentioned above, synaptic noise eliminates sharp thresholds, and the effect may be lost in the noise. Computational model results bear this out: tuning for 2.5% contrast has HWHH similar to that at higher contrast (Fig. 5C). This conclusion is further supported by simulations in which synaptic noise was added to the conceptual model. As mentioned above, in this case the conceptual model behavior matched the broader tuning of the computational model, even at 2.5% contrast (not shown).

Computational Model: Adding Correlation-Based Excitation

To this point, we have not considered the effect of intracortical excitation. We have seen that correlation-based inhibition is sufficient to achieve sharp, contrast-invariant tuning. Here we show that the addition of correlation-based excitation "amplifies" these contrast-invariant responses, without altering their tuning. The conceptual model, which contains only two cortical neurons, is too simple to explore the effects of intracortical excitation in any meaningful way. Hence, the remainder of the paper will present results from the computational model only.

Intracortical excitation was incorporated using a correlation-based rule analogous to that used for intracortical inhibition: excitatory connections were determined probabilistically, such that the strongest connections are found between cells whose RFs are most strongly correlated, *i.e.* those with similar preferred orientation and *similar* spatial phase. This is illustrated schematically by the dashed lines in Fig. 6. That intracortical excitation comes primarily from cells of similar orientation preference and similar spatial phase is supported by the fact that EPSP's are evoked only by stimuli of appropriate position and phase, with opposite phase to the stimuli that evoke IPSP's (Ferster, 1988; Hirsch et al., 1995). More direct support is provided by Freeman et al. (1997) who recorded from pairs of cat V1 simple cells isolated on a single electrode. Cell pairs had similar preferred orientations but randomly varying spatial phases. However, cross-correlations indicative of a monosynaptic excitatory connection were found only when the cells had similar absolute spatial phase (G. Ghose, private communication).

Since the dependence on correlation of intracortical inhibition and excitation differs only in sign, excitatory and inhibitory connections in our model have *precisely* the same average distribution in terms of orientation preference; they differ only in spatial phase. An example is shown in Fig. 8A, which illustrates the experimental V1 orientation map used to assign preferred orientations to cortical cells in the computational model. In this figure, white squares show the locations of cells making excitatory connections to the excitatory cell at the X, while black squares show the locations of cells making inhibitory connections. Excitatory and inhibitory connections to this cell have similar distributions as a function of orientation. Fig. 8B shows the theoretical average distribution of connections for retinotopically identical RFs, as a function of orientation difference (top) and spatial phase difference (bottom). The tightness of tuning as a function of correlation is determined by the parameter n_{pow} (see Methods). Large values of n_{pow} lead to tighter connectivity as a function of correlation, whereas smaller values of n_{pow} lead to broader connectivity. Increasing and decreasing n_{pow} had only minor effects on the behavior of the model.

For the full computational model, we reduced the synaptic strengths of the LGN input by a factor of 2, relative to the previous simulations without intracortical excitation. We thus relied on the positive feedback from same-phase intracortical excitation to amplify the response to suprathreshold stimuli. We also increased the default level of inhibition (i→e conductances) by a factor of 2, leaving the level of feedforward inhibition (LGN→i→e) roughly constant. The reduction of LGN input strength was important when comparing model behavior to experimental estimates of stimulus-induced conductance changes (discussed below). The intracortical excitation increased average firing rates to a 50% contrast stimulus of the preferred orientation by a factor of 2.1, relative to an identical circuit with intracortical excitation removed.

Intracortical amplification of an effective stimulus by feedback excitation has been used in many other models (Douglas et al., 1989, 1995; Ben-Yishai et al., 1995; Somers et al., 1995; Suarez et al., 1995). Our use of excitatory feedback differs in that our "amplifier" is localized in spatial phase as well as in orientation. It also differs from Ben-Yishai et al. and Somers et al. in two important respects. First, our intracortical inhibitory connections spread no further than excitatory connections – both are equally localized in orientation. Second, the resulting amplified responses have F1 tuning identical to that of the thalamocortical input alone, as observed by Ferster et al. (1996) (shown in Fig. 10).

Computational Model: Tuning

For the full circuit, typical currents and voltages from excitatory cells at various orientations are shown in Fig. 9. The behavior of a cell tuned to the stimulus is shown at the top left. At the anti-preferred temporal phase, large inhibitory currents hyperpolarize the cell

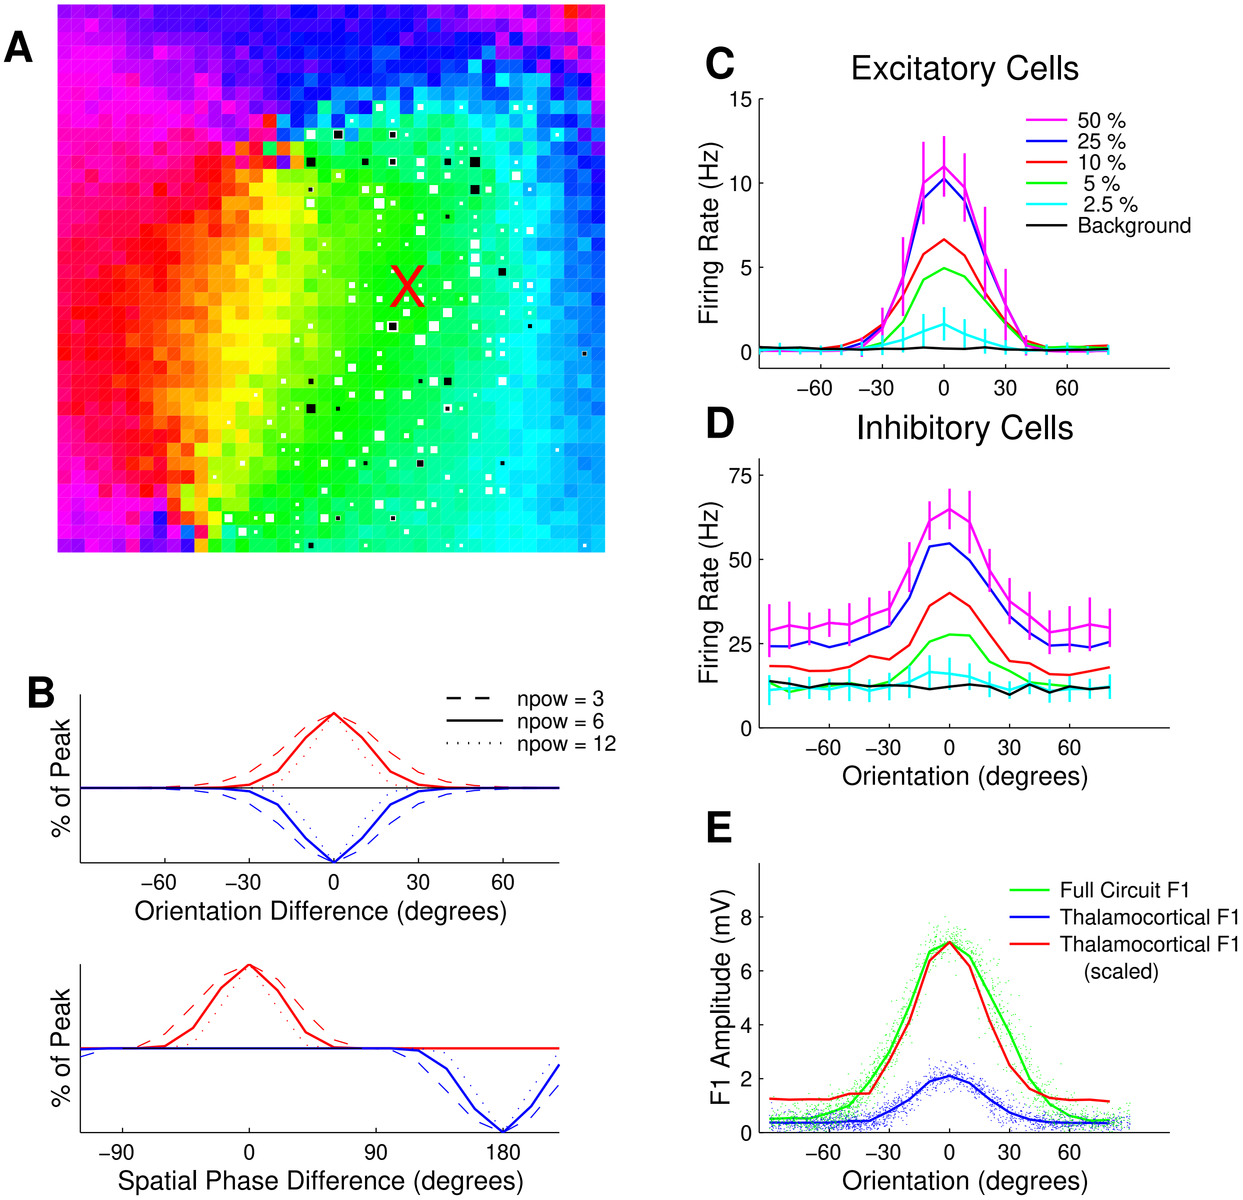

Figure 8. Behavior of the full computational model. **A**. Orientation map used, and typical pattern of intracortical connections. There is one excitatory cell in every location of the cortical map (40×40 lattice) and one inhibitory cell in every fourth location (20×20 lattice). Cells were assigned preferred orientation according to illustrated 40×40 color map (red-green-blue-red representing $0^o - 60^o - 120^o - 180^o$) (map is 2/3 mm x 2/3 mm from measurement in cat V1, provided by Michael Crair and Michael Stryker); RF spatial phases were assigned randomly to each cell, while retinotopic centers progress continuously across the map (described in Methods). Intracortical connections were assigned probabilistically according to RF correlations (excitatory connections, yielding roughly same-phase excitation) or anti-correlations (inhibitory connections, yielding roughly anti-phase inhibition). A typical connectivity pattern is shown by the black and white squares, which illustrate locations of cells making inhibitory or excitatory intracortical synaptic connections, respectively,

Fig. 8 Caption, continued:

to the excitatory cell at the red X. Area of squares is proportional to connection strength. The distributions of excitatory and inhibitory connections across orientations are similar; on average, these distributions are identical. **B**. Theoretical distribution of connectivity as a function of the difference in preferred orientation (top) and the difference in spatial phase (bottom) between two cortical neurons with overlapping RF centers. Probability of excitatory connections shown in red; inhibitory probabilities shown inverted and in blue. All values shown as % of maximal connection probability. The parameter n_{pow} controls the width of tuning as a function of correlation (see Methods); $n_{\text{pow}} = 6$ (solid line) is the default value. Excitation and inhibition have identical spreads as a function of orientation difference, but have opposite preferences for spatial phase. Distribution vs. preferred orientation is averaged over cells of all spatial phases; distribution vs. spatial phase averaged over cells of all preferred orientations, with spatial phase measured with respect to the center of the RF for all orientations. **C-E**. All responses are to 3 Hz, 0.8 cycle/^{o} sinusoidal grating. **C,D**. Firing rates of of excitatory and inhibitory cells, vs. orientation, as function of contrast (indicated by key in C). Error bars for the 50% contrast and 2.5% contrast responses are ± 1 std. dev. **E**. Amplitude (F1) of excitatory cell voltage modulation, with and without the intracortical circuitry, vs. difference of stimulus orientation from preferred. Dots: F1 for all 1600 excitatory cells; traces: means in 10^o orientation bins, as in C,D. Blue: F1 for thalamocortical inputs alone; green: F1 with the full cortical circuit. Red trace is the thalamocortical response scaled to the peak response of the full cortical circuit; note that thalamocortical and full circuit have same tuning, as in Ferster et al. (1996).

to near the chloride reversal potential. As the preferred phase is approached, the inhibitory currents drop and the excitatory currents rise, depolarizing the membrane to spike threshold. Spikes in turn evoke adaptation currents, which help to shut off the cell's response and also, by acting on other cells across the network, lead to lowering of the intracortically evoked excitatory current. For a cell tuned to stimuli perpendicular to that presented (top right in Fig. 9), inhibition dominates at all phases and the cell is prevented from firing.

The full model, including same-phase excitation, achieves contrast invariant tuning to sinusoidal gratings (Fig. 8C, colored lines). Orientation tuning halfwidths at half height of the mean firing rates of excitatory cells were between $19^o - 21^o$ for contrasts ranging from 2.5% to 50%. The model was also well tuned to a moving oriented bar. Tuning width for the default parameters was 27.5, slightly broader than for gratings.

Our model also achieves nearly identical orientation tuning of the excitation and inhibition received by a simple cell, as observed in Ferster (1986). Fig. 10 demonstrates that the F1 tuning curves of excitatory and inhibitory conductances onto cortical cells have nearly identical shapes. However, the spiking responses of excitatory and inhibitory cells differ. Like excitatory cells, inhibitory interneurons have a tuned peak in their firing rates near the preferred orientation; however, unlike the excitatory cells, at the null orientation their firing rates rise above background with increasing contrast (Fig. 8D). This untuned component of the inhibitory response counters the untuned component of the LGN input and is necessary to prevent excitatory spiking for null-oriented stimuli. Since the untuned response component increases with increasing contrast, inhibitory cell tuning measured as HWHH (with background subtracted) broadens with contrast from 32.3^o at 5% contrast to 41.6^o at 50% contrast. If, however, one measures the width of tuning after subtracting the response to null stimuli, inhibitory cells have HWHH across contrasts of 18.6^o to 20.7^o, similar to excitatory cell tuning.

Orientation Tuning is Driven by LGN Input

The tuning in our model is driven by the LGN inputs, despite the strong role of intracortical inhibition. In agreement with the cortical cooling experiments of Ferster et al. (1996), the tuning of the first harmonic (F1) of the membrane potential response to moving gratings was identical for the full circuit and when all intracortical synapses were set to zero strength (Fig. 8E). For our default parameters, we find that the cortical circuitry amplifies the purely thalamocortical F1 by a factor of 3.4, slightly higher than the estimate (2.7) reported by Ferster et al. (1996).

A more rigorous test of the hypothesis that cortical tuning is driven by the LGN can be obtained by systematically varying the tuning of the LGN input. This can be accomplished by varying the spatial frequency of a sinusoidal grating (Fig. 11). As spatial frequency

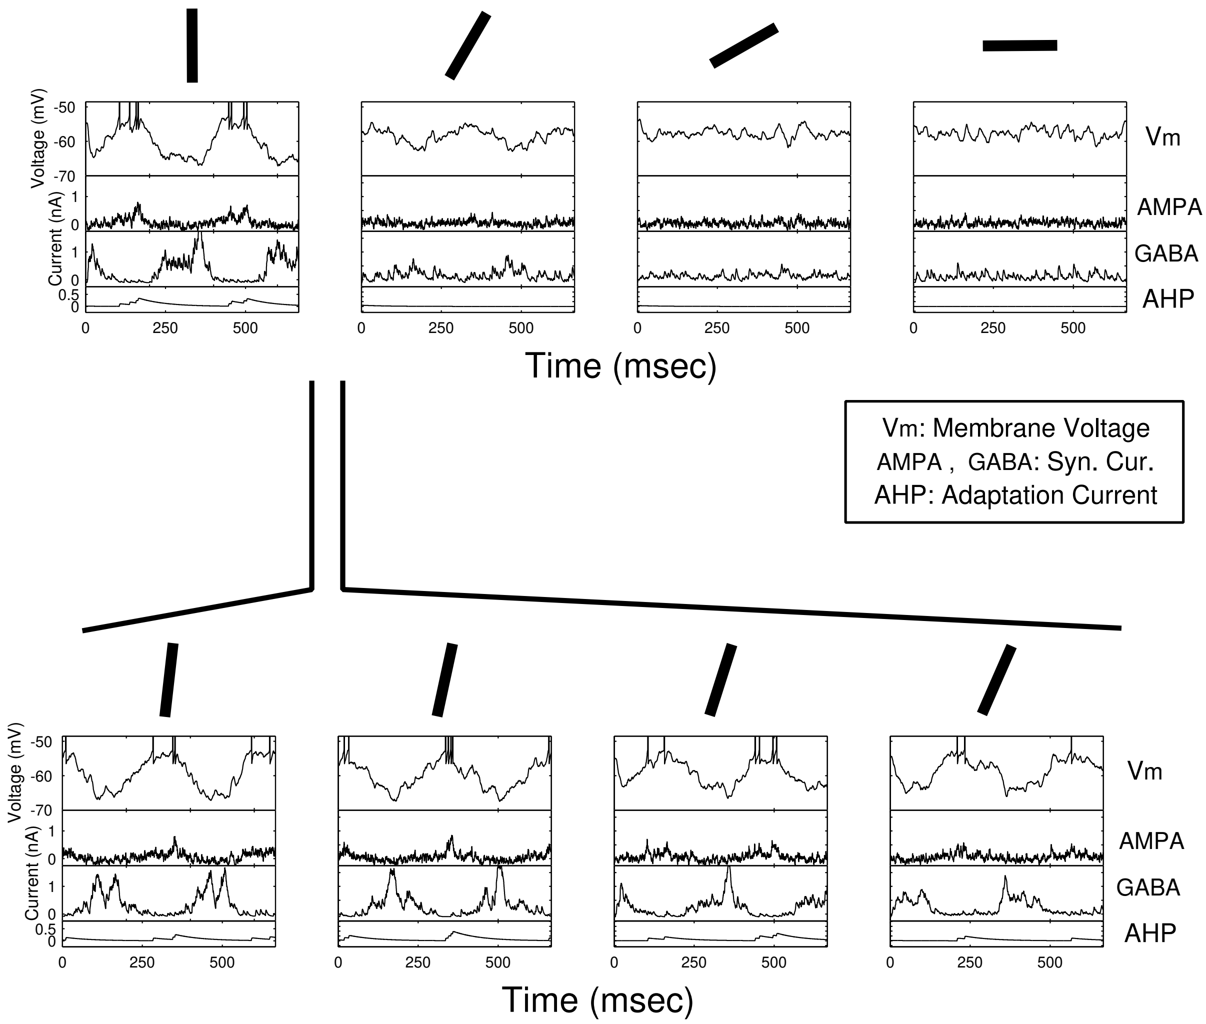

Figure 9. Example traces. Voltage (V_m) and conductance in response to a 3 Hz sinusoidal grating at 50% contrast. AMPA,GABA-A: Synaptic conductances with background subtracted (converted to currents at threshold as in Fig. 7A). AHP: spike triggered potassium conductance. Bars show stimulus orientation relative to preferred (vertical). Top row shows orientation differences spaced at 30^o intervals; bottom row shows model behavior at orientations between 0^o and 30^o. Excitation and inhibition arrive out of phase and have similar orientation tuning. Inhibition dominates at the null.

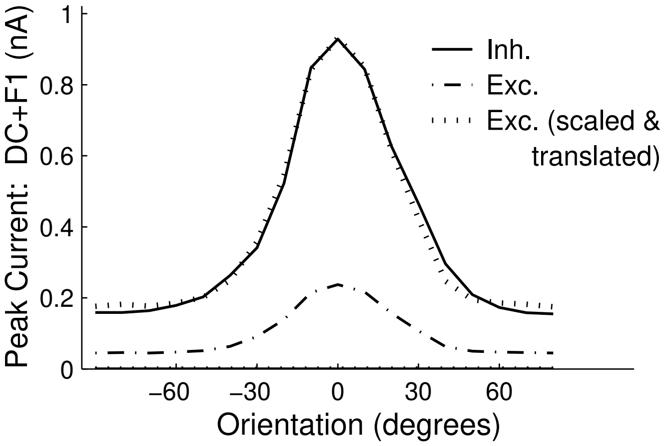

Figure 10. Orientation tuning of peak excitatory current (dot-dashed line) and inhibitory current (solid line) at 50% contrast (peak current equals DC+F1 - see Fig. 7B). Dotted line shows excitatory current scaled and translated to match maxima and minima of inhibitory currents. Peak inhibition is larger than excitation at all orientations, but the tuned component of excitation and inhibition have nearly identical shape.

is increased, the orientation tuning of the F1 of the LGN inputs becomes narrower (Fig. 11A, top). The F1 of the full circuit closely follows the LGN F1 across spatial frequencies (Fig. 11A, middle), yielding a narrowing of orientation tuning of the spiking response with increasing spatial frequency (Fig. 11B). Note that there is an optimal spatial frequency, at which spike response to stimuli of the preferred orientation is maximal (Fig. 11B, inset); thus, as spatial frequency increases, the orientation tuning width monotonically narrows, while spike response to the preferred orientation rises to a maximum and falls again. This is precisely the behavior observed in cat cortical cells (Vidyasagar and Sigüenza, 1985; Webster and De Valois, 1985; Jones et al., 1987; Hammond and Pomfrett, 1990) and is the behavior expected for a cell receiving inputs from a linear Gabor function RF. Similar behavior should occur for other stimulus manipulations that would vary the tuning of the LGN F1 but that have not yet been reported experimentally, such as use of moving or flashed ellipses of various eccentricities; the key point is that the full circuit F1 and the spiking tuning closely covary with the LGN tuning.

Inhibitory sharpening of orientation tuning is due exclusively to disynaptic inhibition driven by the LGN input, *i.e.* sharpening depends on feedforward inhibition. Feedback inhibition (resulting from cortical cell excitation of inhibitory cells – gray dashed lines in Fig. 6) has little effect on model behavior. In simulations run after setting these connections to zero, tuning curve halfwidths at the default parameters remain unchanged to within 1^o for 10% contrast and above (4^o change at 5% contrast), and peak spiking responses increased

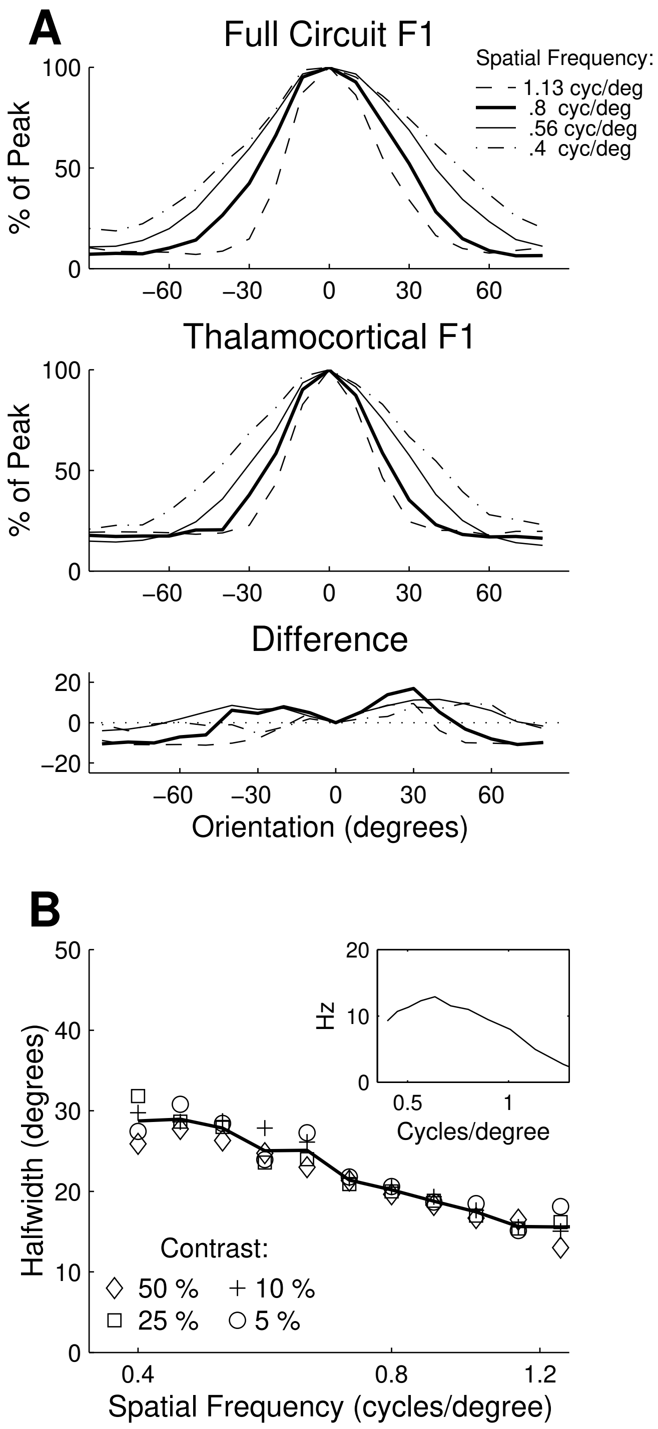

Figure 11. Orientation tuning narrows with increasing spatial frequency of a sinusoidal grating (3 Hz). **A**. Tuning of the F1 voltage modulation for the full circuit (top) and with LGN excitation only (middle). Tuning curves, from widest to narrowest, represent response to spatial frequencies 0.4 (dot-dashed), 0.56 (thin), 0.8 (thick), and 1.13 (dashed) cycles/degree respectively; each curve is normalized to its peak response. Thicker line is the spatial frequency used for simulations in other figures. The bottom shows the difference between the normalized tuning curves – LGN input F1 and full circuit F1 closely match. **B**. Halfwidth at half height vs. spatial frequency for 5%, 10%, 25% and 50% contrast gratings. Orientation tuning remains contrast invariant over a broad range of spatial frequencies. Inset: spatial frequency tuning curve at the preferred orientation and 50% contrast.

by less than 6.4% (see Fig. 13B). These small changes in output occur even though the mean inhibitory conductance in response to a 50% contrast grating at the preferred orientation is reduced by 38%.

This result can be understood by examining Fig. 6. With same-phase excitation and anti-phase inhibition, firing of excitatory cells of one phase increases the inhibition onto excitatory cells of the opposite phase. But excitatory cells with RFs of opposite *spatial* phase spike during opposite *temporal* phase of the stimulus. Therefore, feedback inhibition is directed onto cells that are already not spiking, and has little effect on model behavior.

Inhibitory Dominance and Inhibitory Blockade

The model operates in an inhibition dominated regime, as revealed by a variety of measures. For the "default" parameter settings used to measure tuning (Fig. 8), the total synaptic strength received by a cell (see Methods for definition) was 5 $nA\ msec$, 4.25 $nA\ msec$, and 7.5 $nA\ msec$ for the LGN inputs, cortical excitation and cortical inhibition respectively. A null oriented stimulus induces an inhibitory current that is approximately 3.9 times as large as the excitatory current for a cell clamped at threshold voltage, corresponding to an inhibitory conductance change that is 11.7 times greater than the change in excitatory conductance. Strong inhibition is consistent with the experimental result that non-specific stimulation of LGN afferents (or cortical white matter) yields a short EPSP followed by a large IPSP (e.g. Ferster and Jagadeesh, 1992). However, quantitative measurements of the balance of excitation and inhibition under physiological conditions have not been reported, and estimates based on indirect evidence are imprecise (Shadlen and Newsome, 1994).

Our model agrees with the experiments of Nelson et al. (1994), which demonstrated that intracellular blockade of inhibition in a single neuron, with DC current injection sufficient to suppress excess background firing, did not disrupt cortical orientation tuning. Figure 12 shows extracellular tuning curves for cells under inhibitory blockade with (solid line) and without (dot-dashed line) compensating DC current injection. Tuning curves without blockade are shown for comparison (dashed line). (Note that Nelson et al. (1994) used moving bar stimuli, while we use drifting gratings.) We find a small amount (1.9 Hz) of elevated spiking at the null, consistent with the report that some cells did show spiking at the null under inhibitory blockade (S.B. Nelson, private communication). This reflects the untuned mean DC component of the LGN input.

Effects of global blockade of inhibition are discussed below (Fig. 13A).

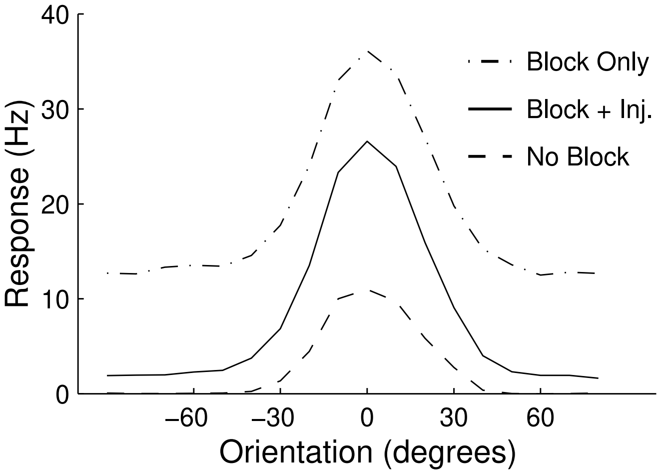

Figure 12. Orientation tuning after blocking inhibition in a single cell, as in Nelson et al., 1994. Tuning curves derived when inhibitory and adaptation currents are blocked within a single cell (see Methods). Dotted line shows tuning with inhibitory blockade only; solid line shows the tuning when negative current equal to the mean inhibitory synaptic current at background is injected into the cell. Dashed line shows tuning without blockade for comparison. Inhibitory blockade with current injection has little effect on tuning. Notice, however, a slight (1.9 Hz) rise in the response to null stimuli. Contrast equals 50%.

Robustness of Model Results

The effect of parameter variations on computational model behavior is shown in Fig. 13. Each of the subplots represents the effect of varying the strength of two variables out of the following four: the three types of synaptic connections in the model (thalamocortical excitation (LGN), intracortical excitation (e→{e,i}), and inhibition (i→e)), and spike-rate adaptation in excitatory cells. Magnitudes of a given type of connection or conductance were varied by multiplying all the corresponding unitary conductances by a constant factor. In Figures 13B and 13D, the intracortical excitation of inhibitory cells was removed (e→i = 0).

Simulations were run at four contrasts (5%, 10%, 25%, 50%) for each point in the 5×5 parameter grid. We monitored contrast-invariance of orientation tuning, average tuning width (HWHH), and firing rates at the preferred orientation and 50% contrast. Contrast-invariance is represented by color within the ovals; darker ovals represent loss of contrast invariance. Spike tuning HWHH is denoted by oval width, relative to oval height which represents 30^o. The light gray shaded regions indicate HWHH tuning of less than 22^o. Spike rate to a 50% contrast preferred stimulus is printed inside each oval. The lines show the relation of model behavior to various experimental estimates and are discussed below.

The model robustly achieves sharp, contrast-invariant tuning. Thin white ovals within

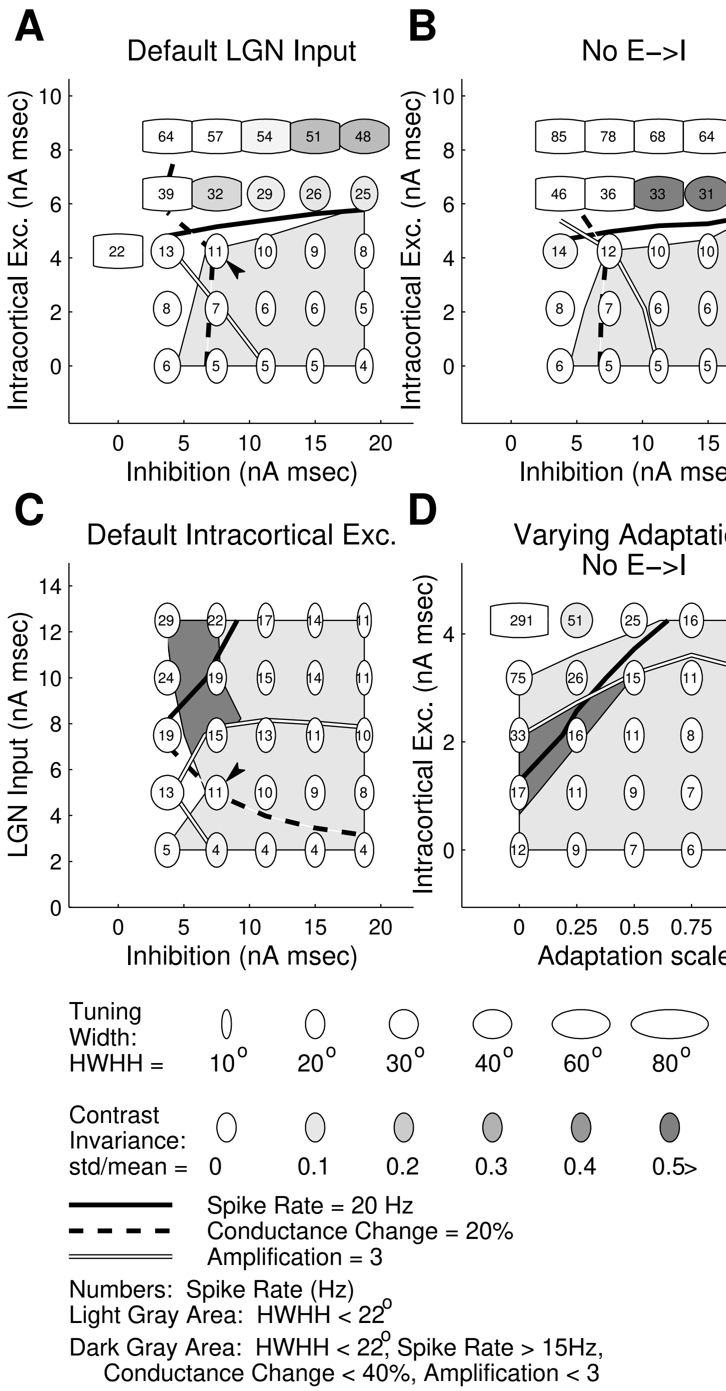

Figure 13. Robustness of model behavior to changes in various model parameters. Each subplot represents the effect of varying the strength of two variables out of the following four: the three types of synaptic connections (LGN input, intracortical excitation, and inhibition), and spike-rate adaptation in excitatory cells. Mean spike tuning half-width at half-height (HWHH) for 5%, 10%, 25%, and 50% contrast gratings is represented by oval width. Oval height represents 30^o. Mean spike rates (Hz) for preferred stimuli at 50% contrast are printed inside each oval. Darker ovals indicate a loss of contrast invariance, monitored as standard deviation divided by the mean of the HWHH over the four contrasts sampled. Points with extremely broad tuning are contrast invariant only because all contrasts give maximal HWHH. Lines show experimentally reported values for mean spike rate (bold line equals 20 Hz; Albrecht, 1995), maximal conductance change from background for a high contrast, null stimulus (dashed line equals 20%; Douglas et al., 1988), and ratio of voltage F1 with and without input from cortical circuitry ("amplification"; white line = 3; Ferster et al., 1996). Light gray area indicates areas of sharp tuning (HWHH $< 22^o$). Dark gray area in C and D indicates regions with HWHH $< 22^o$, amplification < 3, spike rate > 15 Hz, and conductance change $< 40\%$ (C) or $< 22\%$ (D). Arrows in A,C indicate default network parameters used in Figs.

Fig. 13 Caption, continued:
8-12. Note that (i) parameter values that lead to sharp tuning also yield contrast invariance; (ii) higher levels of inhibition sharpen tuning; (iii) high levels of excitation lead to instability – runaway feedback excitation – indicated by high spike rates, broad tuning and loss of contrast invariance; (iv) removing e→i connections causes little change in stable region except that amplification is reduced (also true for varying adaptation as in D, not shown). Within light gray areas, null conductance changes range from: (A) 20% –44%; (B) 21% –42%; (C) 4% – 60%; (D) 21% –23%; amplification ranges from (A) 2.7–3.8; (B) 2.7–3.5; (C) 2.4–3.9; (D) 2.2–3.2; CV of contrast invariance ranges from (A) .02–.06; (B) .03–.11; (C) .01–.15; (D) .01–.06. Gray area and line interpolations obtained with MATLAB "contour" command. See Results for detailed discussion.

the gray shaded regions indicate that, in the large regions of parameter space in which tuning is sharp, it also remains contrast invariant. Increasing inhibition leads to sharper tuning (narrower ovals) while having only moderate effects on spike rate, except with very strong levels of intracortical excitation.

This qualitative behavior is quite robust to changes in other model parameters. Increasing and decreasing the tightness of the correlation-based connectivity by changing the parameter n_{pow} (see Fig. 8B) by a factor of 2 had only moderate effects on model behavior. Using the more broadly tuned geniculocortical connectivity matched to the data of Ferster et al. (1996), tuning widened substantially. However, large levels of inhibition still produced sharp tuning (see also Fig. 5), and the level of intracortical excitation that led to feedback instability remained relatively unchanged.

The Roles of Feedforward and Feedback Inhibition

High levels of intracortical excitation (upper portions of 13A,B,D) leads to unstable feedback excitation, indicated by sharply increased spike rates, broadening of orientation tuning and a loss of contrast invariance. Strong excitation is difficult to control, even for high levels of inhibition, because the model lacks the center-surround ("Mexican hat") intracortical connectivity commonly used to stabilize feedback excitation (Ben-Yishai et al., 1995; Somers et al., 1995). A set of activated excitatory cells will excite cells of nearby orientation and similar spatial phase, while driving feedback inhibition (e→i→e) onto cells of nearby orientation but nearly opposite spatial phase. Excitation can thus spread in orientation along a series of cells linked by similar spatial phase, unchecked by feedback inhibition.

The most important parameter determining the presence of runaway excitation is the strength of intracortical e→e connections. However, feedforward inhibition from the LGN (LGN→i→e) also plays a role. As discussed earlier, the phase-nonspecific component of feedforward input is inhibitory and acts like a contrast-dependent effective threshold, increasing the depolarizing current necessary to reach threshold with increasing contrast (Fig. 7); intracortical excitation must overcome this inhibition in order to spread. At low contrast, the effective threshold is small and presents a weak barrier to the spread of excitation. Consequently, whenever feedback excitation resulted in a loss of contrast invariance (darker ovals), this loss was due to a widening of tuning at lower contrasts. That is, feedback instability poses the greatest problem at lower contrasts, even though peak firing rates are lowest at these contrasts.

Removing feedback inhibition by removing intracortical excitatory input onto inhibitory cells has little effect on the model behavior. This is demonstrated in Fig. 13B, which shows the results of using parameters identical to those of Fig. 13A except that e→i connections are set to 0 strength. Even though mean inhibitory conductance is substantially reduced (by 38% at the preferred orientation, 50% contrast, for default parameters), model behavior is virtually unchanged.

In agreement with experiments using large amounts of GABA-A antagonists (Sillito, 1975; Tsumoto et al., 1979), strong reduction of inhibition leads to a loss of orientation tuning (a single run with zero inhibition is shown at left of parameter grid in Fig. 13A). This loss of tuning is due primarily to the fact that, at very low inhibition levels, the untuned phase-nonspecific component of the LGN input remains uncancelled by feedforward inhibition (though unchecked feedback excitation also contributes). Under less severe blockade, our model predicts that, much as in intracellular GABAergeic blockade (Fig. 12, "block only" condition), excitatory cell tuning should resemble that of inhibitory cells: a well-tuned peak on an untuned platform whose height increases with contrast (see left portion of Fig. 5: orientation tuning gradually widens before being lost as inhibition decreases). Such an untuned platform at a single stimulus contrast has been observed after inhibitory blockade in cat simple cells (Sillito, 1975; Tsumoto et al., 1979); similar results can be seen in monkey visual cortex (Sato et al., 1996).

Comparisons to Experimental Estimates

In addition to the width and invariance of orientation tuning, we monitored the conductance change to a null stimulus, the magnitude of the full circuit voltage F1, and the mean spike rate for a 50% contrast preferred stimulus. The lines in Fig. 13 shows the contours through parameter space corresponding to experimental estimates of these variables. First, some results (Douglas et al., 1988, 1991; Koch et al., 1990) suggest conductance change at the null is small, although there is a great deal of uncertainty about this issue (see Discussion). The dashed lines represent conductance changes that are 20% of background. Smaller changes were achieved only for sufficiently small levels of excitation and inhibition (regions to left of dashed lines). Strong levels of inhibition can sharpen tuning, but at the expense of larger conductance change at the null. Second, Ferster et al. (1996) very roughly estimated a ratio of 2.7 for the magnitude of the voltage F1 with vs. without input from cortical circuitry (we refer to this ratio as the level of "cortical amplification"). A cortical amplification level of 3 is represented by the white lines. Amplifications of this size or smaller are obtained for cortical synaptic strengths that are sufficiently small compared to thalamic input (lower left corner in 13A and 13B; upper portion of 13C). Third, Albrecht (1995) reported that the F1 of the spike rate of simple cells to preferred stimuli at 50% contrast is $\approx$ 38 Hz. Assuming a ratio of F1 to DC of 1.57 (Skottun et al., 1991), this results in a mean spike rate of $\approx$ 20 Hz (bold line). For default levels of LGN input (13A and 13B), spike rates for parameters that achieve contrast-invariant tuning were relatively low.

Simply by varying synaptic strengths, we were not able to quantitatively match all three of these estimates simultaneously. The reason is as follows. Low conductance change at the null means that feedforward conductances must be small relative to the cell's background conductance. Given the responsiveness of our model neurons to these weak inputs, we find that this implies that firing rates will be small unless there is large cortical amplification of the feedforward input. However, such large cortical amplification both disagrees with the estimates of Ferster et al. (1996; see also Chung and Ferster, 1997) and, in our model, leads to instability and loss of tuning. By increasing thalamocortical conductances, we can achieve sharp, contrast-invariant tuning with realistic firing rates and amplification levels, but conductance changes to a null stimulus become larger. This is shown by the dark gray region in Fig. 13C, which shows points with spike rate > 15 Hz, tuning HWHH $< 22^o$, amplification ratio < 3, and conductance change $< 40\%$.

Varying Neuronal Gain

One obvious way out of this dilemma is that cortical cells may respond more strongly to weak thalamocortical input and moderate amplification than do our model neurons. We used extremely simple model neurons, including a simple, voltage-independent mechanism for spike rate adaptation. These were matched to experimental data on firing rate vs. constant somatic current level *in vitro* (McCormick et al., 1985; see Methods and Troyer and Miller, 1997a, 1997b). Our model may not accurately model responses to synaptic currents *in vivo*, where cells are subject to various neuromodulators, slow currents, and active conductances, receive fluctuating synaptic input over a spatially extended dendritic tree, and may show weaker spike rate adaptation. All of these may affect the "neuronal gain", *i.e.* the magnitude of a neuron's response to a given level of synaptic input (*e.g.* Fox et al., 1990; Mel, 1993; Nowak et al., 1997; Storm, 1990; Tang et al., 1997). A simple argument shows that neuronal gain and synaptic strength can trade off against one another: suppose a neuron's gain could be doubled, so that it fires at twice its previous rate in response to any stimulus pattern. If the strength of its synaptic conductances onto other cells were simultaneously cut in half, then each cell in the network would receive the same total input as before the change. By making these changes for all excitatory cells in the network, spike rates could be doubled without altering network behavior, including orientation tuning widths.

To investigate the changes in model behavior induced by increasing neuronal gain, we studied the effects of lowering spike-rate adaptation in model excitatory neurons (Fig. 13D). We also removed e$\rightarrow$i conductances as in 13B; this reduces cortical amplification without altering network behavior (compare 13A and 13B). Spike rates are increased while tuning remains sharp and contrast-invariant (13D). Conductance changes at the null are relatively unaffected by the change in neuronal gain (changes range from 20.6% to 22.5%, excluding the unstable point at upper left). Reducing the spike-triggered AHP current in our model by $\approx 75\%$ leads to a substantial region of parameter space satisfying all experimental estimates: the dark gray region shows points for which spike rates are > 15 Hz, amplification ratios < 3, and tuning HWHH $< 22^o$; in this region, conductance changes are $< 22\%$.

Discussion

Principal Findings and Predictions

We have constructed a simple model that accounts for contrast-invariant orientation tuning in layer 4 of cat visual cortex. First, we analyzed the LGN input to simple cells in response to drifting sinusoidal gratings. This input has two components: a temporally modulated, phase-specific component which is tuned for orientation, and an unmodulated, phase-nonspecific component which, assuming LGN cells are not orientation tuned, must be completely untuned. Because of LGN rectification, this untuned input increases significantly with increasing contrast, implying that no simple threshold mechanism could correct for it across contrasts. This problem is general: the phase-nonspecific LGN input component for any stimulus is poorly tuned and increases with contrast. To counteract this increase in untuned feedforward excitation, stimulation of an excitatory cell at its null orientation must evoke increasing intracortical inhibition and/or increasing withdrawal of intracortical excitation with increasing contrast. Thus, the central questions for understanding contrast-invariant orientation tuning in cat layer 4 are how the level of this contrast-dependent "pull" is computed, and which cells are its source.

Second, we have demonstrated that cortical circuitry relying on correlation-based connectivity is sufficient to robustly yield contrast-invariant orientation tuning. In our model, the contrast-dependent inhibition to layer 4 simple cells comes from inhibitory cells with similar orientation tuning and roughly opposite spatial phase (Fig. 8B). The tuned input component is amplified by excitatory connections between cells with similar orientation tuning and similar spatial phase. While some details might vary, our key predictions are that the layer 4 circuitry is strongly phase-specific, inhibition-dominated, and localized in orientation, and has orientation tuning that depends on the LGN input tuning.

Beyond the circuit itself, our most striking prediction is the existence of inhibitory simple cells that are tuned for orientation, but have a contrast-dependent response to null-oriented stimuli. We also predict that (i) the DC LGN input should increase with increasing contrast at all orientations, while (ii) the dominance of inhibition should cause the net DC feedforward input (LGN plus feedforward inhibition) to decrease with increasing contrast at all orientations. These predictions are most easily tested with null-oriented stimuli, for which feedback excitation should be negligible. Prediction (i) could be tested by blocking inhibition intracellularly in a single layer 4 cell. Note that, for (ii), the predicted decrease is in net current at threshold voltage; this need not imply voltage decrease from rest due to differences between excitatory and inhibitory reversal potentials. In addition, we predict that intracellular, or sufficiently weak extracellular, GABAergic blockade in layer 4 should reveal a tuned response component sitting on an untuned plateau response (*e.g.* Fig. 12).

Possible Sources of Contrast-Dependent Inhibition

Besides anti-phase inhibition, what are other possible sources of the contrast-dependent "pull" required to achieve contrast-invariant tuning? Since the LGN input is exclusively excitatory (Ferster and Lindstrom, 1983), inhibition must come from other cortical cells. One obvious possibility is inhibitory cells preferring stimuli of the orthogonal (null) orientation. A major contribution from such cells seems ruled out in cat layer 4 by the similar orientation tuning of excitation and inhibition (Ferster, 1986). and by the phase-specificity and spatial opponency of EPSPs and IPSPs (Ferster, 1988; Hirsch et al., 1995).

Another source of "pull" could be withdrawal of excitation from other cortical cells. Since layer 4 simple cells have very low spontaneous rates, significant withdrawal would most likely come from complex cells, which do show contrast-dependent inhibition to a null stimulus (Sclar and Freeman, 1982). This explanation would simply move the problem of the origin of the required inhibition onto complex cells.

One study has suggested that inhibition onto layer 4 simple cells with multiple ON/OFF subregions may come primarily from layer 4 simple cells with single-subregion RFs (Toyama et al., 1981), while multi-subregion inhibitory cells in layer 4 may primarily act on layer 3 complex cells. Thus, the population showing an untuned plateau response might be restricted to the single-subregion inhibitory cells; inhibition from these cells might block responses at the null orientation in both inhibitory and excitatory multi-subregion cells.

Carandini and Ferster (1997) have shown that a tonic hyperpolarization underlies contrast adaptation in cat simple cells. This contrast-dependent "pull" may make an important contribution toward contrast-invariant orientation tuning for steady-state stimuli. However, this mechanism would not play a role for transient stimuli such as moving or flashed bars. While contrast-invariance for these stimuli has not been explored in detail, our analysis shows that tight orientation tuning to high-contrast bars combined with robust responses to low contrast bars would require contrast-dependent cancellation of the phase-nonspecific component of the LGN input. Furthermore, contrast-dependent "pull" is required at all orientations, while contrast adaptation appears to not be induced by null-oriented stimuli (Allison and Martin, 1997).

Our analysis shows that the untuned component of the LGN input will cause simple cells to spike in response to a high-contrast, null-oriented stimulus, unless they are inhibited. Our model relies on the simplest explanation for the source of this contrast-dependent inhibition: inhibitory neurons (or a subset of them) are not strongly inhibited and *do* spike, providing the inhibition necessary to prevent responses in the remaining layer 4 neurons.

Experimental Results Related to Inhibitory Cell Tuning

Several studies have reported that cat layer 4 inhibitory interneurons have simple cell RF structure much like that of excitatory neurons, and are orientation tuned (Gilbert and Wiesel, 1979; Toyama et al., 1981; Martin, 1988, p. 665 and references therein; J. Hirsch, unpublished observations). However, only Azouz et al. (1997) reported details of orientation tuning in V1 inhibitory neurons. Consistent with our predictions (Krukowski et al., 1996), they found that 5 of 8 inhibitory cells identified in intracellular recording showed significant spiking response at the null; the strongest null response was found in the only layer 4 interneuron studied. However, they did not study contrast dependence of responses, and also reported that in their hands 70% of extracellularly recorded cells showed spiking at the null.

Our prediction for inhibitory cell tuning is also consistent with observations in other systems. In rabbit visual cortex, putative inhibitory neurons respond to all orientations, while other neurons respond only to a limited range of orientations (Swadlow, 1988); orientation tuning of these cells was not otherwise reported. More generally, putative inhibitory neurons are more broadly tuned than other cells across a variety of cortical systems (Simons and Carvell, 1989; Swadlow, 1989, 1990, 1991, 1994; Brumberg et al., 1996).

Input-Driven Tuning

Orientation tuning in our model is driven by the LGN inputs. The model yields voltage responses that have identical F1 tuning with and without intracortical synaptic input (Figs. 8,10; Ferster et al. 1996). In addition, orientation tuning sharpens with increasing spatial frequency of a sinusoidal grating, as is also observed in cat visual cortex (Vidyasagar and Sigüenza, 1985; Webster and De Valois, 1985; Jones et al., 1987; Hammond and Pomfrett, 1990)

Nonetheless, cortical circuitry plays a key role in our model. Inhibitory cells have an untuned component to their response that eliminates the untuned component of the LGN input and sharpens spike tuning. Studies in the rat whisker barrel system also indicate that the layer 4 computation is local, input-driven, and dependent on broad inhibitory tuning (Brumberg et al., 1996; Brumberg et al., 1997; Pinto et al., 1996; Simons and Carvell, 1989), suggesting that these may be general properties of layer 4 cortical circuitry.

Other models (Ben-Yishai et al., 1995; Somers et al., 1995; Adorjan et al., 1997) predict that orientation tuning is determined by cortical circuitry, independent of input tuning, and thus should not change with stimulus spatial frequency. Such models do not allow simultaneous representation of multiple orientations at a single retinotopic position (Carandini and Ringach, 1997), which may play an important role in visual processing, including figure/ground separation and object recognition (Zucker, 1986).

Model Robustness and Experimental Constraints

Our model uses correlation-based inhibition to robustly achieve sharp, contrast-invariant, input-driven orientation tuning over a wide parameter range. The ability of a simple, two-cell conceptual model to capture the essential behavior of the more realistic computational model argues strongly for the robustness of the underlying mechanisms.

We examined the conditions under which this tuning could be achieved while satisfying several other constraints suggested by a variety of experiments: (i) small ($< 20\%$) conductance changes to a null stimulus; (ii) high ($\approx$ 20Hz) steady-state firing rates to a high-contrast preferred stimulus; (iii) only moderate ($\approx$2.7 times) cortical amplification of the LGN input. We have found that these conditions can be simultaneously satisfied if the neuronal gain of our model excitatory neurons is increased, for example by reducing the magnitude of the spike-trigged adaptation (AHP) current. Alternately, by increasing feedforward synaptic strengths, all but the conductance condition can be robustly met.

The experimental support for these constraints in layer 4 is not entirely clear. With the exception of the amplification level, our numerical targets are based on average values across cortical simple cells from all layers, but properties of layer 4 cells may differ (*e.g.* see Tolhurst and Thompson, 1981, which suggests orientation tuning is broadest in layer 4). Indeed, if the layer 4 circuitry solves the problem of contrast-invariant tuning, it may be relatively easy for cells in other layers to display sharp tuning with high rates and small conductance changes. In addition, quantitative estimates of conductance change remain controversial. Different labs and techniques are producing widely varying measurements (Douglas et al., 1988, 1991; Hirsch et al., 1995 and unpublished observations; Carandini and Ferster, 1997; Monier et al., 1997). Furthermore, the relationship between net synaptic input and conductance changes measured at the soma is complicated by factors such as voltage-dependent dendritic conductances (*e.g.* Yuste and Tank, 1996) and large synaptic background conductances (Bernander et al., 1991), which were not considered in previous theoretical studies (Koch et al., 1990).

Withdrawal of tonic input could also have an important effect on stimulus-induced conductance changes. Withdrawal of excitation from complex cells, discussed above, would contribute to low conductance change in response to a null stimulus. Withdrawal of tonic anti-phase inhibition might allow a large spike response to a preferred stimulus without a large change in conductance. There are some indications from experiments that this change may be small during spike responses (Carandini and Ferster, 1997; Monier et al., 1997). Our model shows a large conductance change at the preferred orientation (mean increase 72%; peak 148%). Withdrawal of tonic inhibition would be enhanced in our model by raising the tonic background spike rate of our model inhibitory cells (currently $\approx$ 12 Hz, see Methods) to match values typically found experimentally ($\approx$ 20 Hz; Simons and Carvell, 1989; Swadlow, 1989, 1990, 1991, 1994; Brumberg et al., 1996). Anti-phase i→i connections, which have been observed in inhibitory simple cells in layer 4 (J. Hirsch, unpublished observations), could also contribute to withdrawal of inhibition at the preferred orientation, but at the cost of reducing inhibition at the null. This reduction in inhibition at the null could be compensated for by increasing inhibitory gain and/or i→e connection strengths.

Many factors not studied here might increase the stability of feedback excitation, enabling high firing rates with small increases in null conductance (at the cost of large amplification). These factors include frequency-dependent short-term synaptic depression (Abbott et al., 1997; Priebe et al., 1997; Thomson and Deuchars, 1994; Tsodyks and Markram, 1997) and slow GABA-B mediated inhibition (Allison et al., 1996; Bush and Priebe, 1998). In addition, weaker components of inhibitory connectivity that are phase-nonspecific or broadly tuned relative to excitation (*e.g.* Worgotter and Koch, 1991) could help to stabilize the amplifier without dominating orientation tuning.

Nonlinear Response Properties With Local Circuitry

Many groups have suggested that (i) important features of simple cell responses can be approximated using a linear function of stimulus contrast (Movshon et al., 1978; Glezer et al., 1982; Tolhurst and Dean, 1990; Albrecht and Geisler, 1991; Heeger, 1992; Carandini and Heeger, 1994; Carandini et al., 1997a,b); and (ii) such linear responses might be achieved by a circuit involving opponent or "push-pull" inhibition, in which inhibition *balances* excitation (Glezer et al., 1982; Tolhurst and Dean, 1990; Carandini and Heeger, 1994; Carandini et al., 1997a,b). Several groups further proposed that a variety of nonlinear contrast effects could be explained through addition to the linear model of a global, divisive normalization of activity levels (Albrecht and Geisler, 1991; Heeger, 1992; Carandini and Heeger, 1994; Carandini et al., 1997a,b); Carandini and colleagues proposed that this normalization could arise through global, orientation-nonspecific inhibitory connectivity.

Our model uses opponent inhibition, but assumes that inhibition dominates, rather than balances, excitation. The inhibition resulting from the phase-nonspecific LGN input depends strongly on contrast but not orientation, and can produce nonlinear effects similar to the normalizing inhibition of Carandini and colleagues, although the biological substrate is quite different. For example, preliminary simulations indicate that one such nonlinearity, cross-orientation inhibition (suppression of response to a stimulus of the preferred orientation by simultaneous presentation of a null-oriented stimulus), arises naturally in our model from the anti-phase inhibition (A. Hoffman and present authors, unpublished observations).

Additional nonlinearities include effects of contrast on temporal responses. Preliminary results (Priebe et al., 1997) indicate that the adaptation current contributes to contrast-dependent temporal phase advance, and that the addition of frequency-dependent short-term synaptic depression (Markram and Tsodyks, 1996; Stratford et al., 1996; Abbott et al., 1997; Tsodyks and Markram, 1997) to our model can largely account for contrast effects on both temporal phase and temporal frequency tuning in cat V1 (related observations have been made independently by Chance et al. (1997)). Thus, we hypothesize that local rather than global circuitry can account for contrast nonlinearities as well as for contrast invariance, although this will require further study.

Developmental and Functional Implications

The LGN input to simple cells assumed here – a "Hubel-Wiesel" model – has previously been shown to develop from correlation-based rules of synaptic development under simple assumptions (Miller, 1994). Our model uses cortical connectivity that should also arise from such rules, as generalized to include plasticity of inhibitory synapses. Thus our model circuit, in addition to explaining many properties of cortical functional responses, naturally suggests its own developmental origin: central features of the intracortical and thalamocortical connectivity in cortical layer 4 may arise together through similar mechanisms of correlation-based development.

The anti-phase inhibition that results from such development allows cortical circuitry local to a small number of iso-orientation columns to distinguish the form (orientation) from the intensity (contrast) of an oriented stimulus. Suppose that spatial phase is ignored, so that all cortical cells preferring the same orientation receive the same single component of LGN input, an instantaneous synaptic input rate. Then the two input variables of contrast and orientation are confounded. An intermediate input rate might correspond to the preferred orientation at low contrast, or a non-preferred orientation at higher contrast. To disambiguate these, a comparison across cells of all preferred orientations is needed, to determine for each cell whether it is receiving more or less input than cells of other preferred orientations. Recent models achieve this comparison through circuitries in which intracortical inhibition is broader in orientation than excitation (Ben-Yishai et al., 1995; Somers et al., 1995).

Our model instead takes into account spatial phase, and predicts that the computation in layer 4 is local. If the cell and its anti-phase partners are receiving similar input (Fig. 6B), the preferred orientation is *not* being seen, and the cell does not respond, regardless of stimulus contrast (intensity). If the cell is receiving input and its anti-phase partners are not (Fig. 6A), the preferred orientation (form) is present, and the cell responds; the strength of its response reflects the stimulus contrast.

This provides a specific example of a possible, more general principle of cortical organization that should arise from correlation-based development: a cortical cell should respond to the *difference* between its preferred stimulus (call it $\mathcal{P}$) and its "anti-preferred" stimulus (call it $\bar{\mathcal{P}}$), rather than to the preferred stimulus alone. By anti-preferred stimulus, we mean the stimulus that evokes an input pattern most anti-correlated with that of the preferred. A stimulus that is *un*correlated with the preferred stimulus pattern (*i.e.* evokes an uncorrelated input pattern), while partially stimulating the cell, should also stimulate the cell's "anti-preferred partners" (cells with preferred stimulus $\bar{\mathcal{P}}$). Correlation-based inhibition prevents the cell from responding to this inappropriate stimulus, regardless of stimulus intensity. This hypothetical principle, that a cell selective for $\mathcal{P}$ actually responds to "$\mathcal{P}$ AND NOT $\bar{\mathcal{P}}$", might aptly be termed "dialectical," in recognition of the philosophical school that argues that objects exist and are known only in relation to their opposites (*e.g.* Merleau-Ponty, 1962).

Correlation-based connectivity also suggests a possible developmental explanation for columnar organization (the vertical invariance of RF properties). Why don't inhibitory and excitatory neurons in a given column take on opposite preferred orientations (or ocular dominance (OD))? We hypothesize that interconnected excitatory and inhibitory neurons share RF properties that are shared by both $\mathcal{P}$ and $\bar{\mathcal{P}}$, and differ in the RF properties that distinguish these. Thus, orientation (and OD) is invariant in a column, while spatial phase varies (Freeman et al., 1997), because the most anticorrelated stimulus pair has identical orientation (and OD) but opposite spatial phase. It will be of great interest to determine whether such a principle might apply to cortical representations of other sensory modalities.

References

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

Adorjan P, Lund J, and Obermayer K (1997) A model for generation and sharpening of orientation selectivity with the same cortical microcircuits in macaque striate cortex. *Soc Neurosci Abstr*, 23:2061.

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

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

Allison J, Kabara J, Snider R, Casagrande V, and Bonds A (1996) GABA-B-receptor-mediated inhibition reduces the orientation selectivity of the sustained response of striate cortical neurons in cats. *Visual Neurosci*, 13:559–566.

Allison J and Martin K (1997) Contrast adaptation produced by null direction and cross orientation stimulation of neurons in cat visual cortex. *Soc Neuro Abs*, 23:454.

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

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

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

Bernander O, Douglas R, Martin K, and Koch C (1991) Synaptic background activity influences spatiotemporal integration in single pyramidal cells. *Proc Nat Acad Sci USA*, 88:11569–11573.

Bishop P, Kozak W, and Vakkur G (1962) Some quantitative aspects of the cat's eye: axis and plane of reference, visual field co-ordinates and optics. *J Physiol*, 163:466–502.

Bonds A (1989) Role of inhibition in the specification of orientation selectivity of cells in the cat striate cortex. *Visual Neuroscience*, 2:41–55.

Brumberg J, Goldreich D, Kyriazi H, and Simons D (1997) Adjacent whisker input to a layer IV barrel does not arise from adjacent barrels. *Soc Neuro Abstr*, 23:2344.

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

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

Bush P and Priebe N (1998) GABAergic inhibitory control of the transient and sustained components of orientation selectivity in a model microcolumn in layer 4 of cat visual cortex. *Neural Comput*, 10:(in press).

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

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

Carandini M, Heeger D, and Movshon J (1997a) Linearity and gain control in V1 simple cells In Jones E and Peters A, editors, *Cerebral Cortex, Vol. X* Plenum, NY.

Carandini M, Heeger D, and Movshon J (1997b) Linearity and normalization in simple cells of the macaque primary visual cortex. *J Neurosci*, 17(12):8621–8644.

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

Chance F, Nelson S, and Abbott L (1997) Effects of synaptic depression of temporal nonlinearities in responses of model simple cells. *Soc Neurosci Abstr*, 23:1266.

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

Chung S and Ferster D (1997) The size of thalamic input to simple cells of the cat visual cortex. *Soc Neurosci Abstr*, 23:2059.

DeAngelis G, Ghose G, Ohzawa I, and Freeman R (1992) Spatiotemporal receptive field structure and phase relationships between adjacent simple cells in the cat's striate cortex. *Soc Neurosci Abstr*, 18:10.

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

Douglas R, Martin K A, and Whitteridge D (1989) A canonical microcircuit for neocortex. *Neural Comput*, 1:480–488.

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

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

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

Ferster D (1987) Origin of orientation-selective EPSPs in simple cells of cat visual cortex. *J Neurosci*, 7:1780–1791.

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

Ferster D (1990) X- and Y-mediated synaptic potentials in neurons of areas 17 and 18 of cat visual cortex. *Vis Neurosci*, 4:115–133.

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

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

Fox K, Sato H, and Daw N (1990) The effect of varying stimulus intensity on NMDA-receptor activity in cat visual cortex. *J Neurophysiol*, 64:1413–1428.

Freeman R, DeAngelis G, Ghose G, and Ohzawa I (1997) Clustering of response properties of neurons in the visual cortex. *Soc Neurosci Abstr*, 23:567.

Ghose G, DeAngelis G, Ohzawa I, and Freeman R (1993) Similarity between the spatiotemporal response profiles of nearby cells in the cat's and kitten's visual system. *Soc Neurosci Abstr*, 19:424.

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

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

Glezer V, Tsherbach T, Gauselman V, and Bondarko V (1982) Spatio-temporal organization of receptive fields of the cat striate cortex. the receptive fields as the grating filters. *Biol Cybern*, 43:35–49.

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

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

Heggelund P and Albus K (1978) Orientation selectivity of single cells in striate cortex of cat: the shape of orientation tuning curves. *Vision Res*, 18:1067–1071.

Hirsch J, Alonso J, and Reid R (1995) Two-dimensional spatial maps of excitation and inhibition in the simple receptive field. *Soc Neurosci Abstr*, 21:21.

Hubel D and Wiesel T (1961) Integrative action in the cat's lateral geniculate body. *J Physiol*, 155:385–398.

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

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

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

Kaplan E, Purpura K, and Shapley R (1987) Contrast affects the transmission of visual information trough the mammalian lateral geniculate nucleus. *J Physiol*, 391:267–288.

Koch C, Douglas R, and Wehmeier U (1990) Visibility of synaptically induced conductance changes: theory and simulations of anatomically characterized cortical pyramidal cells. *J Neurosci*, 10:1728–1744.

Komatsu Y (1996) $GABA_B$ receptors, monamine receptors, and postsynaptic inositol trisphosphate-induced Ca^{2+} release are involved in the induction of long-term potentiation at visual cortical inhibitory synapses. *J Neurosci*, 16:6342–6352.

Krukowski A, Priebe N, and Miller K (1996) A model of simple-cell orientation tuning: Feedforward tuning and correlation-based intracortical connectivity. *Soc Neurosci Abst*, 22:642.

Linsenmeier R, Frishman L, Jakiela H, and Enroth Cugell C (1982) Receptive field properties of X and Y cells in the cat retina derived from contrast sensitivity measurements. *Vision Res*, 22:1173–1183.

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

Martin K (1988) From single cells to simple circuits in the cerebral cortex. *J Exp Physiol*, 73:637–702.

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

Mel B (1993) Synaptic integration in an excitable dendritic tree. *J Neurophysiol*, 70:1086–1101.

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

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

Monier C, Borg Graham L, and Fregnac Y (1997) The dominant increase in synaptic conductance evoked by visual input in cat area 17 is due to shunting inhibition. *Soc Neurosci Abstr*, 23:2363.

Movshon J, Thompson I, and Tolhurst D (1978) Spatial and temporal contrast sensitivity of neurones in areas 17 and 18 of the cat visual cortex. *J Physiol*, 283:101–120.

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

Nowak L, Sanchez Vives M, and McCormick D (1997) Influence of low and high frequency inputs on spike timing in visual cortical neurons. *Cerebral Cortex*, 7:487–501.

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

Peichl L and Wassle H (1979) Size, scatter and coverage of ganglion cell receptive field centres in the cat retina. *J Physiol*, 291:117–141.

Peters A and Yilmaz E (1993) Neuronal organization in area-17 of cat visual cortex. *Cerebral Cortex*, 3:49–68.

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

Priebe N, Kayser A, Krukowski A, and Miller K (1997) A model of simple cell orientation tuning: The role of synaptic depression. *Soc Neurosci Abstr*, 23:2061.

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

Sato H, Katsuyama N, Tamura H, Hata Y, and Tsumoto T (1996) Mechanisms underlying orientation slectivity of neurons in the primary visual cortex of the macaque. *J Physiol*, 494:757–771.

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

Shadlen M and Newsome W (1994) Noise, neural codes and cortical organization. *Curr Opin Neurobiol*, 4:569–579.

Sherman S (1985) Functional organization of the W-, X-, and Y-cell pathways in the cat: a review and hypothesis. *Progr Psychobiol Physiol Psychol*, 11:233–314.

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.

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

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

Skottun B, De Valois R, Grosof D, Movshon J, Albrecht D, and Bonds A (1991) Classifying simple and complex cells on the basis of response modulation. *Vision Res*, 38:1079–86.

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

Storm J (1990) Potassium currents in hippocampal pyramidal cells. *Prog Brain Res*, 83:161–187.

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

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

Swadlow H (1988) Efferent neurons and suspected interneurons in binocular visual cortex of the awake rabbit: Receptive fields and binocular properties. *J Neurophysiol*, 59:1162–1187.

Swadlow H (1989) Efferent neurons and suspected interneurons in S-1 vibrissa cortex of the awake rabbit: Receptive fields and axonal properties. *J Neurophysiol*, 62:288–308.

Swadlow H (1990) Efferent neurons and suspected interneurons in S-1 forelimb representation of the awake rabbit: Receptive fields and axonal properties. *J Neurophysiol*, 63:1477–1498.

Swadlow H (1991) Efferent neurons and suspected interneurons in second somatosensory cortex of the awake rabbit: Receptive fields and axonal properties. *J Neurophysiol*, 66:1392–1409.

Swadlow H (1994) Efferent neurons and suspected interneurons in motor cortex of the awake rabbit: axonal properties, sensory receptive fields, and subthreshold synaptic inputs. *J Neurophysiol*, 71:437–453.

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

Tang A, Bartels A, and Sejnowski T (1997) Effects of cholinergic modulation on responses of neocortical neurons to fluctuating input. *Cerebral Cortex*, 7:502–509.

Thomson A and Deuchars J (1994) Temporal and spatial properties of local circuits in neocortex. *Trends Neurosci*, 17:119–126.

Tolhurst D and Dean A (1990) The effects of contrast on the linearity of spatial summation of simple cells in the cat's striate cortex. *Exp Brain Res*, 79:582–588.

Tolhurst D and Thompson I (1981) On the variety of spatial frequency selectivities shown by neurons in area 17 of the cat. *Proc Royal Soc Lond B*, 213:183–199.

Toyama K, Kimura M, and Tanaka K (1981) Organization of cat visual cortex as investigated by cross-correlation technique. *J Neurophysiol*, 46:202–214.

Troyer T and Miller K (1997a) Integrate-and-fire neurons matched to physiological f-I curves yield high input sensitivity and wide dynamic range In Bower J, editor, *Computational Neuroscience: Trends in Research 1997*, pages 197–201 Plenum Press, NY.

Troyer T and Miller K (1997b) Physiological gain leads to high ISI variability in a simple model of a cortical regular spiking cell. *Neural Comput*, 9:971–983.

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

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

Tusa R, Palmer L, and Rosenquist A (1978) The retinotopic organization of area 17 (striate cortex) in the cat. *J Comp Neurol*, 177:213–235.

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

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

Worgotter F and Koch C (1991) A detailed model of the primary visual pathway in the cat: Comparison of afferent excitatory and intracortical inhibitory connection schemes for orientation selectivity. *J Neurosci*, 11:1959–1979.

Yuste R and Tank D (1996) Dendritic integration in mammalian neurons, a century after Cajal. *Neuron*, 16:701–716.

Zucker S (1986) The computational connection in vision: early orientation selection. *Behav Res Meth Instr and Computers*, 18:608–617.

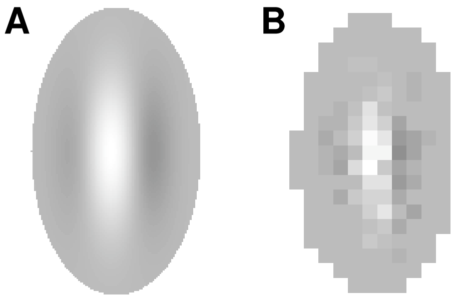

Figure 2. Gabor-shaped cortical RFs. Lighter grays to white: positive values of Gabor function, corresponding to weights of ON-center LGN cells with centers at corresponding spatial positions; darker grays to black: negative values of Gabor function, corresponding to weights of OFF-center cells. **A**. A full Gabor function, used to determine LGN inputs to a cortical cell in the conceptual model. **B**. Typical LGN inputs to a cortical cell in the computational model, after probabilistic sampling from the full Gabor (see Methods). These receptive fields are 'typical'; different cortical cells may have different preferred orientations, spatial phase (relative locations of ON or OFF subregions), spatial location and, in the computational model, different outcomes of the probabilistic sampling. Spatial frequency of sinusoid in Gabor function is 0.8 cycles/degree.

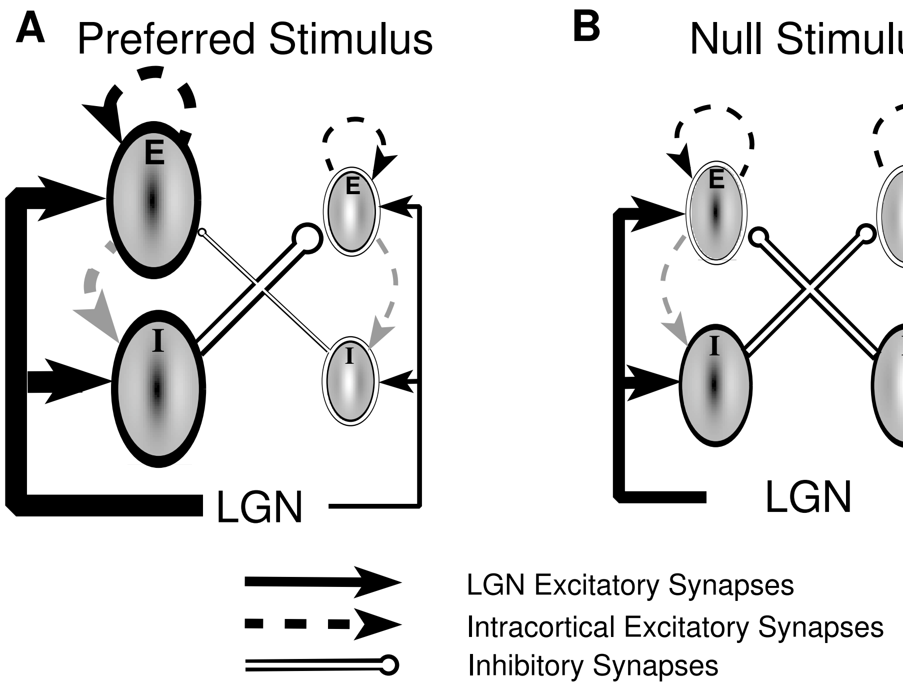

Figure 6. Behavior of model using correlation-based connectivity. Schematic representing behavior of the model in response to preferred (**A**) and null (**B**) stimuli. The excitatory cell described in the Results is in the upper left; its inhibitory "anti-phase partner" is in the lower right. E: excitatory cells; I: inhibitory cells. Solid lines: excitation and depolarization; open lines: inhibition and hyperpolarization. Line thickness and size of RF icon represent magnitude of activity. Dashed lines represent correlation-based excitation, which is included in the complete computational model only (Figs. 8-11). Some simulations were performed without cortical excitatory projections onto inhibitory neurons (gray dashed lines), but this did not substantially affect network behavior (see Fig. 13B).