Computational Modeling of Orientation Tuning Dynamics in Monkey Primary Visual Cortex

M. C. Pugh, D. L. Ringach, R. Shapley and M. J. Shelley

August 20, 1999

to appear in the Journal of Computational Neuroscience 8(2000)143–159

please address correspondence to:
Dr. M. C. Pugh, Department of Mathematics, University of Pennsylvania
209 South 33rd Street, Philadelphia, PA 19104-6395
Phone: (215) 898-8472, FAX: (215) 573-4063, Email: `mpugh@math.upenn.edu`

D. L. Ringach, Departments of Neurobiology and Psychology, University of California at Los Angeles, Los Angeles, CA 90095
R. Shapley, Center for Neural Science, New York University, NY, NY 10003
M. J. Shelley, Courant Inst. for Mathematical Sciences, NYU, NY, NY 10012

Abstract

In the primate visual pathway, orientation tuning of neurons is first observed in the primary visual cortex. The LGN cells that comprise the thalamic input to V1 are not orientation tuned, but some V1 neurons are quite selective. Two main classes of theoretical models have been offered to explain orientation selectivity: feedforward models, in which inputs from spatially aligned LGN cells are summed together by one cortical neuron; and feedback models, in which an initial weak orientation bias due to convergent LGN input is sharpened and amplified by intracortical feedback. Recent data on the dynamics of orientation tuning, obtained by a cross-correlation technique, may help to distinguish between these classes of models. To test this possibility, we simulated the measurement of orientation tuning dynamics on various receptive field models, including a simple Hubel-Wiesel type feedforward model: a linear spatio-temporal filter followed by an integrate-and-fire spike generator. The computational study reveals that simple feedforward models may account for some aspects of the experimental data, but fail to explain many salient features of orientation tuning dynamics in V1 cells. A simple feedback model of interacting cells is also considered. This model is successful in explaining the appearance of Mexican-hat orientation profiles, but other features of the data continue to be unexplained.

Keywords: Cortical dynamics, orientation tuning, monkey, primary visual cortex, layers.

Introduction

How visual cells in the primary visual cortex (area V1) become selective to the orientation of a visual stimulus is an important question in visual neuroscience (see, for example, Das (1996); Reid and Alonso (1996); Sompolinsky and Shapley (1997)). The question is how V1 cells become strongly tuned for orientation since it is known that cells in the thalamus are not. Hubel and Wiesel, who first discovered the property of orientation tuning in V1 cells, proposed a simple circuitry to explain orientation tuning (Hubel and Wiesel, 1962, 1968). They suggested that the inputs from LGN cells with spatially aligned receptive fields are summed together by one cortical neuron. This spatial alignment of receptive fields generates a preference for a particular orientation in the cortical cell. We refer to this, and similar proposals, as feedforward models. An alternate class of models are feedback (or recurrent) models, which suggest that the feedforward tuning is initially weak and only provides a small orientation bias to the cortical cell. This bias is then amplified and sharpened by intracortical feedback (Wörgötter et al., 1991; Douglas et al., 1995; Somers et al., 1995; Ben-Yishai et al., 1995; Maex and Orban, 1996; Carandini and Ringach, 1997). Clearly, these models are not mutually exclusive; there is a continuum of models depending on the relative strength of the feedforward and feedback connections (Ben-Yishai et al., 1995).

In this work we study recent data on the dynamics of orientation tuning obtained by a cross-correlation technique (Ringach et al., 1997b). The method used for measuring tuning dynamics is a modification of the cross-correlation method and shares many of its properties (de Boer and Kuyper, 1968; Marmarelis and Marmarelis, 1978; Jones and Palmer, 1987; Victor, 1992; Reid et al., 1997). We believe such data may help to determine where the cortical circuitry is best represented along the continuum between feedforward models and feedback models.

As shown below, there are clear differences of dynamics between cell classes in V1. The aim of this paper is to consider these qualitative features, rather than fitting models to data. To achieve this goal, we simulated the experimental protocol on various types of receptive field models, including a simple Hubel-Wiesel type feedforward model: a spatio-temporal filter followed by an integrate-and-fire spike generator. The computational study reveals that simple feedforward models may account for some aspects of the experimental data, but fail to explain many salient features, including the Mexican-hat profile of the orientation distributions and the active inhibition seen at off-optimal orientations ("cross-orientation" inhibition). We then studied a feedback network of interacting cells. A similar model is used in the earlier work of Ben-Yishai et al. (1995). We find that the model with lateral coupling better approximates the behavior of some V1 neurons. However, some aspects of the experimental data remain unexplained.

This paper is organized as follows. We first describe the experiment and the most salient features of the data. A theoretical and computational study of the feed-forward model is then presented. We then study of the behavior of a small network of interacting oriented neurons. Finally, we discuss the relevance of our findings to the modeling of orientation tuning and cortical function.

The Experiment

The visual stimuli used are made up of sinusoidal luminance gratings,

$$I(x, y; \theta, \phi) = A\,(1 + \epsilon \sin[\,\omega(x \cos\theta - y \sin\theta) - \phi\,])\,. \tag{1}$$

Here, $(x, y) \in [-1, 1] \times [-1, 1]$ represents coordinates in visual space and θ and ϕ are fixed parameters whose values range from -90° to 90° and from 0° to 360° respectively. The angle θ corresponds to the orientation of the grating in space while ϕ represents a spatial phase relative to the origin. The variable ω represents the spatial frequency of the grating. A is the mean luminance of the grating and $0 \leq \epsilon \leq 1$ is its contrast. Images of uniform luminance, A, with no spatial structure ($\epsilon = 0$), will be referred to as "blanks".

The visual stimulus used in the experiments of Ringach et al. (1997a) consists of a rapid sequence of randomly selected images, each of which is either a "blank" or a grating at one of N possible angles: $\theta_1 = -90^\circ$, $\theta_2 = -90^\circ + \Delta\theta$, $\ldots$,$\theta_N = -90^\circ + (N - 1)\Delta\theta$, and one of M possible phases: $\phi_1 = 0^\circ$, $\ldots$, $\phi_M = 0^\circ + (M - 1)\Delta\phi$; where $\Delta\theta = 180^\circ/N$ and $\Delta\phi = 360^\circ/M$. Blanks are represented by a dummy orientation with index $N + 1$; $\theta_{N+1} = blank$.

A random sequence of images, $\{I_i\}$, is generated from an associated random sequence of orientation angles and phases $\{(\theta_{n(i)}, \phi_{m(i)})\}$ where $n(i)$ is chosen from $\{1, \ldots, N + 1\}$ with probability $1/(N + 1)$ and $m(i)$ is chosen from $\{1, \ldots, M\}$ with probability $1/M$. θ and ϕ are piecewise constant functions in time, changing every T ms (the duration of one "frame" in the movie). In the physiology experiments T was 16.6 ms, $N = 18$, and $M = 8$.

The responses of cells to the visual stimulus are recorded and analyzed using the following cross-correlation technique. First, a correlation time $\tau \geq 0$ is selected. An $(N + 1) \times M$ array of counters, C, is zeroed. We read through the recorded spike train. If a spike occurred at time t we find the orientation, θ_i, and spatial phase, ϕ_j, of the grating shown at time $t - \tau$. We then increment the corresponding counter, $C(i, j; \tau)$, by one. If the image was a blank, then $i = N + 1$ and j is chosen randomly from $1 \ldots M$. We apply this procedure to each spike in the spike train. In this way, all the spikes are represented in the counter array.

Averaging over spatial phases, we define $Pr(\theta_i; \tau) = \sum_j C(i, j; \tau) / \sum_{i,j} C(i, j; \tau)$. $Pr(\theta_i; \tau)$ is the relative frequency at which a grating with orientation θ_i appeared in the stimulus sequence τ ms before a spike. Because the relative frequency approximates the probability distribution, in the following we call $Pr(\theta; \tau)$ the "probability" that a grating at angle θ appeared τ ms before a spike. The dynamics of orientation tuning are studied by measuring how $Pr(\theta_i; \tau)$ evolves with the correlation time τ. Blanks provide a dynamic baseline to which the responses to gratings can be compared. In our interpretation, if $Pr(\theta_i; \tau) > Pr(blank; \tau)$, for a fixed τ, then the orientation θ_i is excitatory relative to the blank. Similarly, if $Pr(\theta_i; \tau) < Pr(blank; \tau)$ then θ_i is inhibitory.

How will $Pr(\theta_i; \tau)$ behave as a function of τ? For $\tau = 0$, for example, we do not expect to see any influence of the visual stimulus on the output train of action potentials since there is a time delay between the retina and the evoked activity of cortical cells. Therefore, $Pr(\theta; 0)$ should be a uniform (flat) distribution. Similarly, for very large correlation times τ we also expect to obtain a uniform distribution for $Pr(\theta; \tau)$; this is because cells have only finite memory. For intermediate values of τ one would expect to see a smooth peaked distribution which represents the preference of the cell for particular orientations. There will be one value of τ for which the cell will be better tuned than at any other time; we will refer to this value as the "optimal τ".

Experimental Results

In this section we describe some of the salient features of the experimental results in Ringach et al. (1997a, 1998) that we would like the mathematical models to explain. No obvious differences were found between the dynamics of simple and complex cells. Both classes of neurons exhibited the range of behaviors described below. The perspective plots in Fig. 1 illustrate typical results obtained from orientation selective cells in monkey V1 at different values of the correlation time τ, with τ given beside each graph. Each plot represents the measurements from a single cell. The graphs within the perspective plots show the distributions $Pr(\theta;\tau)$. The dashed line represents the baseline given by the probability of a blank, $Pr(blank;\tau)$. The range of values for the correlation times τ was chosen to bracket the response period for each cell. The x-axis represents the orientation θ between $-90°$ and $90°$, and the y-axis represents the value of $Pr(\theta;\tau)$. The vertical scale bar in the lower graph of each perspective plot has a height of $1/(N+1)$, which is the value for a uniform distribution on the orientation domain.

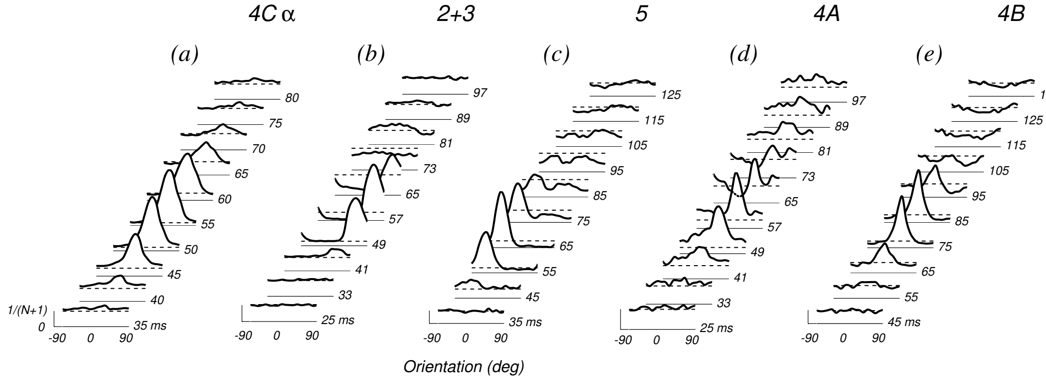

Figure 1: Examples of orientation tuning dynamics from monkey V1 cells. (a) Orientation tuning dynamics from a cell in layer 4Cα. The graphs within the perspective plots show the distributions $Pr(\theta;\tau)$ at different values of τ. The dashed line represents the baseline given by the probability of a blank, $Pr(blank;\tau)$. The x-axis represents the orientation θ between $-90°$ and $90°$, and the y-axis represents the value of $Pr(\theta;\tau)$. The vertical scale bar in the lower graph of each perspective plot has a height of $1/(N+1)$, which is the value for a uniform distribution on the orientation domain; (b) Orientation dynamics from a cell in layer 2+3 ; (c) Orientation dynamics from a cell in layer 5; (d) Orientation tuning from a cell in layer 4A. The profile at 57 ms is "transparent" with the profile at 65 msec indicated by a dotted line; (e) Orientation tuning from a cell in layer 4B.

Fig. 1(a) illustrates the response obtained from a cell located in layer $4C\alpha$, one of the main input layers of V1. For $\tau < 40$ ms and $\tau > 75$ ms the distributions are uniform. The distribution develops into a unimodal distribution and peaks at $\tau = 55$ ms. The distribution's peak occurs at the same angle as the preferred angle of the cell as found using conventional drifting grating stimuli. Notice that, in this case, the probability of the orientation orthogonal to the preferred one and the probability of a "blank" behave in a very similar way. In fact, this behavior does not hold for all layer 4C cells; there are cells for which the blanks and the orthogonal orientation do not have similar dynamics. More intricate dynamical responses are observed outside layer 4C (Ringach et al., 1997b); see Figs. 1(b–e).

Consider the responses depicted in Fig. 1(b). For $\tau < 40$ ms and $\tau > 90$ ms the distributions are uniform. The response develops and peaks at $\tau = 57$ ms. The distribution's peak occurs at the same angle as the preferred angle of the cell as found using conventional drifting grating stimuli. The distribution relaxes to a uniform profile around $\tau = 73$ ms, at which time the blank is more probable than any angle. We observe an "inversion" of the probability profile at $\tau = 81$ ms. At this time the probability of the preferred orientation of the cell is lower than that of any other angle; instead, the most probable angle $\tau = 81$ ms before a spike is the angle *orthogonal* to the preferred angle. Many cells show such "inversions", and this is one of the properties we would like the models to explain.

Another important aspect of the behavior shown in Fig. 1(b) is the dynamics of the probability of a "blank". Initially, for $\tau < 57$ ms, the probability of a blank behaves in a very similar way to the probability of the orientation orthogonal to the preferred angle. However, this is not the case for larger values of τ. For example, for 57 ms $\leq \tau \leq$ 73 ms the probability of an angle that is not close to the preferred angle is *less* than the probability of a "blank". This suggests active inhibition at off-optimal orientations. We note that this behavior also "inverts" at $\tau = 81$ ms. An explanation for the dynamics of this baseline is another goal of the mathematical modeling.

The results shown in Fig. 1(c) show a response pattern similar to that seen in Fig. 1(b). A unimodal distribution develops and peaks at $\tau = 65$ ms followed by a slight "inversion" of the distribution at $\tau > 115$ ms. However, we also see that for 75 ms $\leq \tau \leq$ 105 ms the distribution is clearly bimodal; it has local maxima at the preferred orientation and at the orientation orthogonal to it and has local minima flanking the preferred orientation. We refer to these profiles as "Mexican-hat" shaped — observe the shape at $\tau = 75$ ms. Such distributions are characteristic of many of the well-tuned cells found in V1. The dynamics of the "blank" baseline in Fig. 1(c) is quite similar to that in Fig. 1(b). In particular, for 65 ms $\leq \tau \leq$ 85 ms off-preferred angles appear to be inhibited – their probabilities are below the blank baseline.

In Fig. 1(d) we present another example of a cell that develops a Mexican-hat distribution profile. Early in the dynamics, for $\tau = 41$ ms, the distribution is quite broad. The distribution then sharpens and by $\tau = 57$ ms the cell is much better tuned. A Mexican-hat profile is evident at $\tau = 65$ ms. Finally, in Fig. 1(e) we present another example of a well-tuned cell that develops a Mexican-hat profile at $\tau = 95$ ms and which later, at $\tau = 115$ ms, exhibits an "inversion".

The dynamical features shown in Fig. 1 are typical of V1 cells. The goal of this study is to explore what kind of models may explain the phenomena of "inversions", the development of Mexican-hat shaped profiles, and the dynamics of the "blank" baseline.

A Feed-forward model

We first present a feed-forward model that has been considered extensively in the literature as a good candidate to explain the spatio-temporal properties of simple cells (Movshon et al., 1978; Citron and Emerson, 1983; Jones and Palmer, 1987; McLean and Palmer, 1989; DeAngelis et al., 1993a; McLean et al., 1994; Reid et al., 1997). It consists of a linear spatio-temporal filter followed by an integrate-and-fire spike generator with leakage:

$$\frac{dv}{dt}(t) = -\lambda(v(t) - v_r) + DC + \int_0^t \iint H(x, y, t - s)I(x, y, s)\ dx\ dy\ ds. \tag{2}$$

$H(x, y, t)$ is the impulse response of a linear spatio-temporal filter, corresponding to a linearized model of the input from the LGN, incorporating inputs from both on-center and off-center cells (Palmer and Davis, 1981; Ferster, 1986; Heeger, 1992b,a). $I(x, y, s)$ is the image shown at time s and λ is the rate of leakage ($1/\lambda$ is the membrane time constant). The leakage has been set to zero in most of the computations presented here; although we do discuss the effects of having nonzero leakage in various cases. DC represents DC input and v_r is the resting potential. A spike is generated when the membrane voltage variable $v(t)$ reaches the firing threshold of -50 mV, the voltage is then reset to $v_r = -70$ mV (Connors et al. (1982); McCormick et al. (1985); Somers et al. (1995); Carandini et al. (1996)). In the model we also mimic the effects of an inhibitory reversal potential: the voltage is not allowed to decrease past a lower barrier of -90 mV (see appendix).

We assume that the kernel H is space-time separable, hence

$$\frac{dv}{dt}(t) = -\lambda(v(t) - v_r) + DC + \int_0^t G(t - s)r(s)ds. \tag{3}$$

where $r(s)$ is the spatial dot product of the image at time s and a spatial kernel,

$$r(s) = \iint K(x, y; \theta_K, \phi_K)\ I(x, y, s)\ dx\ dy. \tag{4}$$

The spatial kernel is taken to be a windowed Gabor function (Marcelja, 1980; Daugman, 1985),

$$K(x, y; \theta_K, \phi_K) = K_0\ W(d)\ e^{-\frac{d^2}{L^2}}\ \sin(\ \omega(x \cos\theta_K - y \sin\theta_K) - \phi_K\ ). \tag{5}$$

This two-dimensional spatial kernel is parametrized by an orientation angle θ_K and phase ϕ_K (the preferred orientation and phase of the simulated cell). In the above, ω is the spatial frequency of the carrier, and $d^2 = x^2 + y^2$. The function is windowed to reflect that in the experiment the stimuli are shown in a circular patch: $W(d) = 1$ if $d < 1$ and 0 otherwise. The length-scale L determines the rate of decay away from the center of the receptive field. Fig. 2(a) shows the Gabor function when $\omega = 3\pi$, $\theta_K = 0$, $\phi_K = 0°$ and $L\omega = 4.2$. We use these parameters for the rest of the article. (Ringach et al., 1997c) found Gabor-like receptive fields with multiple subfields in monkey V1.

In our simulations we use a biphasic temporal kernel:

$$G(t) = \begin{cases} a_e(t/\tau)^5 e^{-t/\tau} & \text{for } 0 \le t \le .05 \\ a_e(t/\tau)^5 e^{-t/\tau} - a_i \left(\frac{(t-.05)}{\tau}\right)^3 e^{-(t-.05)/\tau} & \text{for } t \ge .05 \end{cases} \tag{6}$$

where $a_e = 1.67$, $a_i = 16.7$, $\tau = .01$ sec (see Fig. 2(b)). The temporal kernel is chosen so that $\int_0^\infty G(t)\ dt = 1$ and $\max(G)/\min(G) \sim -5$, with its maximum occurring at 50 ms, its minimum around 80 ms, and G decaying to zero around 150 ms (Gielen et al., 1981). Such a choice of G introduces both time lag and memory effects. G models a magnocellular cell; it has both a positive and a negative response and has a positive mean.

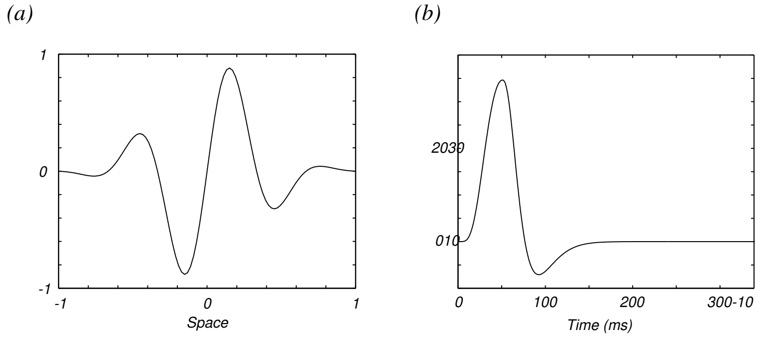

Figure 2: Spatial and temporal filters of the feedforward model. (a) The spatial kernel is the windowed Gabor function used in the simulations (see Eqn. (5)). Plotted here is a cross-section of the kernel, at $y = 0$, $-1 \leq x \leq 1$. (b) The temporal kernel is as given in Eqn (6).

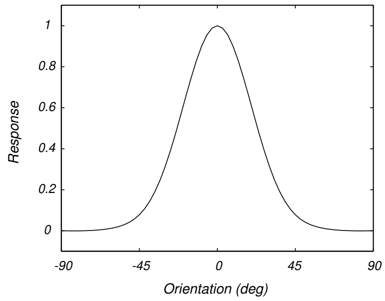

Figure 3: Orientation tuning of the feedforward spatial kernel. The graph shows a plot of $r_{0^\circ}(\theta)$. It shows the orientation tuning of the LGN input in the feedforward model.

Interesting behaviors and analytical results can already be observed in a system whose temporal kernel is a δ-function. In this case, the evolution equation is local in time,

$$\frac{dv}{dt}(t) = -\lambda(v(t) - v_r) + DC + r(t). \tag{7}$$

As shown in the Appendix, in this situation the cross-correlation functions can be studied analytically. We find that leakage can cause sharpening in the cross-correlation functions. Also, in the absence of leakage there is instantaneous decorrelation; $Pr(\theta, \tau)$ becomes flat for all $\tau > 0$.

Computational Results for the Feed-forward Model

We now present results for the feedforward model. To consider a specific example, assume a spatial kernel $K(x, y; \theta_K, \phi_K)$ as shown in Fig. 2(a). This kernel is odd-symmetric about $x = 0$. For the visual stimulus we choose a new grating pattern every $T = 17$ ms. The gratings, described by Eqn. (1), are selected from a family of gratings with 60 equally spaced orientation angles and 6 equally spaced spatial phases. The number of angles and phases is essentially arbitrary; there is no biological limitation on how close two angles or phases can be before their difference is imperceptible. For a fixed phase ϕ, we first compute the spatial inner product of a grating with the spatial kernel:

$$r_\phi(\theta) = \epsilon A \iint K(x, y; 0^\circ, 0^\circ) \sin\left(\, \omega(x \cos\theta - y \sin\theta) - \phi \,\right)\, dx\, dy. \tag{8}$$

Because the kernel K is odd-symmetric, $r_\phi(\theta) = \cos(\phi) r_{0^\circ}(\theta)$. Fig. 3 shows r_{0°, on $[-90^\circ, 90^\circ)$. To abbreviate, in the sequel we refer to the spatial response component $r_\phi(\theta)$ as simply the "response".

The LGN input does not, on average, provide DC input because the six response functions $r_0(\theta)$, $r_{60^\circ}(\theta)$, $r_{120^\circ}(\theta)$, $r_{180^\circ}(\theta)$, $r_{240^\circ}(\theta)$, and $r_{300^\circ}(\theta)$ sum together to zero on $[-90^\circ, 90^\circ)$. It follows that the expected value of the LGN input at any time is zero. DC input is provided by the term DC.

The response scales as the product $\epsilon A K_0$; the product of the contrast, the mean luminance, and the kernel "gain" K_0 (see Eqn. 5). We choose $K_0 = 1.2013$ so that $\sum r_{0^\circ}(\theta_i) = N$ when $A = \epsilon = 1$. This choice is arbitrary and does not affect the results. We then fix the mean luminance A and vary the contrast ϵ to yield firing rates in the experimentally-observed range: between 1 and 20 spikes/sec.

Given a random sequence of images, the voltage equation (3) is numerically integrated in time. We use a modified second-order Runge-Kutta method (M. Shelley, unpublished results) with 1 ms time-steps. The initial voltage is the resting potential $v_r = -70$ mV. When the voltage reaches an upper threshold of -50 mV, it is reset to v_r and the spike is registered by being included in the cross-correlation counters $C(i, j; \tau)$ for τ between 1 and 340 ms. The voltage is not allowed to decrease below -90 mV.

Fig. 4 shows the cross-correlation functions as a function of the correlation time τ. The cross-correlation function $Pr(\theta; \tau)$ rises and peaks at about 55 ms, then falls, returning to baseline at 75 ms, followed by an "inversion" between 75 ms and 115 ms. As shown in Fig. 1(b), "inversions" are sometimes observed in the experimental data. We also considered temporal kernels different from (6). We found that strictly positive kernels led to no inversion. The kernel (6) has a positive mean with significantly more "positive" area than negative. We considered kernels with less of an imbalance between positive and negative area, as well as kernels with zero mean. In such cases, the inversion was sometimes replaced by a second, smaller, peak rising and falling between 75 ms and 115 ms. This is caused by the visual stimulus containing images that are exactly out of phase relative to the preferred phase.

If the system leading to the measurement were completely linear then $Pr(\theta; \tau)$ would be a uniform distribution. (The averaging over phases would completely flatten $Pr(\theta; \tau)$.) The resetting of the voltage after a spike event is the nonlinearity in the model that causes $Pr(\theta; \tau)$ to be nonuniform.

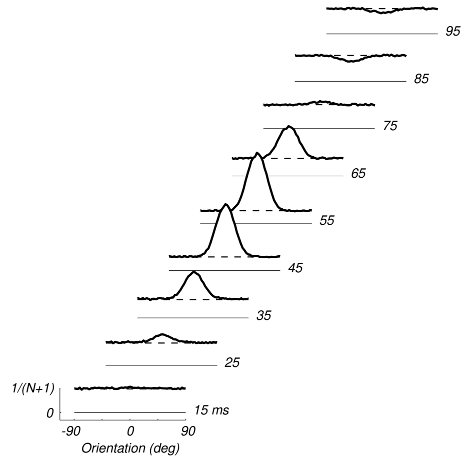

Figure 4: Orientation tuning dynamics of a feedforward model. The plotting conventions are the same as in Fig. 1. The simulation was stopped when 200,000 spikes were collected. There was no leakage and no DC input, $\lambda = 0, DC = 0$, and $A\epsilon = 994.6$. This resulted in a mean interspike time of 108 ms (9.24 spikes/sec) with a standard deviation of 167 ms.

The dashed line corresponds to the baseline — the probability of having shown a “blank” τ ms before a spike. The baseline dynamics coincide with the dynamics of the cross-correlation function at the angle orthogonal to the preferred orientation, $Pr(90^\circ, \tau)$. This is expected because the dot product between the spatial kernel and any grating orthogonal to the preferred orientation equals zero, which is also the value generated by a “blank”.

The voltage equation (3) has three terms; the first term is a leakage term, the second term is DC input, and the third term is LGN input. In the absence of leakage, there are two mechanisms to adjust the mean firing rate — the DC input and the product ϵA. Intuitively, if one holds ϵA fixed and increases the DC input, then one is decreasing the number of spikes that are correlated to the visual stimulus. For this reason, one would expect that the cross-correlation functions would show less tuning the stronger the DC input. We test this expectation by first fixing the contrast $\epsilon = 1/2$ and choosing the mean luminance A so that in the absence of DC input the mean firing rate is around 1 spike per second; the DC input is then increased. The results for four values of the DC component, resulting in mean firing rates of 15, 5.1, 2.4, and 1.1 spikes per second, are presented in Fig. 5.

The cross-correlation functions from the four computations all show the same qualitative dynamics shown in Fig. 4. A peak rises at the preferred orientation, falls, and is followed by an inversion. However, both the maximum height of the peak and the maximum depth of the inversion decrease as the DC input increases. In Fig. 5(a), we plot $Pr(0^\circ, \tau)$ as a function of the correlation time τ, for the four DC inputs. The peak achieves its maximum at 53 ms. The larger the DC input, the smaller the peak. This suggests a loss of tuning. To confirm this, we plot $Pr(\theta, 53)$ as a function of angle θ in Fig. 5(b). Almost all correlation to the preferred angle 0° is lost at the highest level of DC input, as expected.

Another way to vary the firing rate is to hold the DC input fixed and to then vary the quantity ϵA. Maximal tuning is expected for $DC = 0$; this is when the spikes are maximally correlated to the visual stimulus. We found that fixing $DC \neq 0$ and then varying $A\epsilon$ yields the same phenomena. The leakage λ was again set to zero. Four values of ϵA were selected resulting in mean firing rates of 23, 15, 7.1, and 1.1 spikes per second. The ratio of the smallest value of ϵA to the largest value is .14. So if the largest value corresponded to contrast $\epsilon = 1$, then the smallest contrast was $\epsilon = .14$.

The cross-correlation functions have qualitative dynamics similar to those shown in Fig. 4. Fig. 6(a) shows $Pr(0^\circ; \tau)$ as a function of τ. This figure demonstrates that $Pr(0^\circ; 50)$ increases as the firing rate decreases (as ϵA decreases), suggesting a sharpening. To confirm this sharpening, in Fig. 6(b) we plot $Pr(\theta; 50)$ as a function of θ. For almost all firing rates, the peak is followed by an inversion. Fig. 6(c) depicts $Pr(\theta; 84)$ as a function of θ. A value of $\tau = 84$ ms was chosen as being the time at which the inversion was largest at the 15 spikes/second simulation. The inversion is seen in all but the 1.1 spikes/sec simulation.

The above simulations were done with no leakage, $\lambda = 0$. We find that if one fixes $A\epsilon$ and DC and then increases the leakage, and plots the analogues of the figures in Figs. 6, one sees the same behavior with firing rate: the greater the leakage (and hence the lower the firing rate) the greater the sharpening.

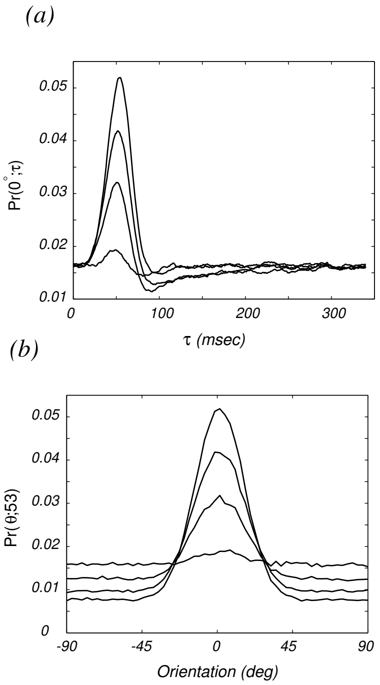

Figure 5: Dependence of dynamics on DC input when $\epsilon A = 270$ is fixed. Four values of DC are selected, $DC = 0, 40, 100, 300$, resulting in mean firing rates of 1.1, 2.4, 5.1, and 15 spikes per second. There is no leakage, $\lambda = 0$. (a) The curves represent the dynamics at the preferred orientation angle, $Pr(0°, \tau)$ at the four different DC input levels. The peak height decreases monotonically as the firing rate increase (as the DC input increases). (b) Orientation probability distributions for the time at which tuning is greatest, 53 ms, as a function of DC input.

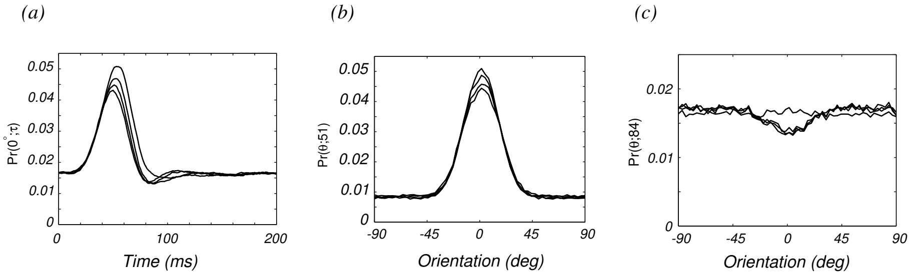

Figure 6: Dependence of dynamics on mean spike rate. There is no leakage and no DC input, $\lambda = 0, DC = 0$. The mean firing rate is increased by increasing ϵA. Four values of $\epsilon A = 1934, 1371, 828.6, 276.2$ are selected, resulting in mean firing rates of 1.1, 7.1, 15, and 23 spikes per second. (a) The curves represent the dynamics at the preferred orientation angle, $Pr(0°, \tau)$, for the four different mean spike rate levels. As the spike rate decreases the peak height increases. (b) Orientation probability distributions for the time at which tuning is greatest, 51 ms, as a function of mean spike rate. (c) Orientation probability distributions for the time at which an inversion occurs, 84 ms, at the different mean rates elicited by the four different levels of ϵA. Except for the lowest mean rate, all the distributions have the same inversion of response at $\theta = 0°$.

The computations presented above used a spatial kernel that is anti-symmetric about the y-axis: $\phi_K = 0$ in formula (5). If the spatial kernel is not anti-symmetric, then the response to a blank and the response to a grating at the orthogonal orientation are not constrained to be equal to each other: $r_\phi(90°) \neq r_\phi(blank)$. One may ask if this would make it possible to observe different behaviors between the dynamics of "blanks" and of the orthogonal orientation in the cross-correlation functions (as seen in the experimental data).

We tested this possibility by performing simulations using spatial kernels with $\phi_K = \pi/4, 2\pi/4, \ldots 7\pi/4$. Here, the response $r_\phi(90°)$ to the orthogonal angle can differ from the response to a blank by up to 1.1% (for $\phi_K = 5\pi/4$). This variation is relatively small — the response to the preferred angle differs from the response to a blank by 48%. Thus, while we found that other spatial kernels may result in the baseline and orthogonal orientations having different dynamics the differences are slight and far from the magnitude of the effect seen in the experimental data. Also, the experimental data normally shows that, initially, all orientations are excitatory relative to the "blank", while late in the response many of the off-optimal orientations are below the "blank" baseline and appear to be actively inhibited (Ringach et al. (1997a)). We did not observe this sort of dynamics in our computations.

We conclude that a simple feedforward model is able to account for some of the features seen in the dynamical experimental data. In particular, the "inversions" seen in the orientation distributions may be a consequence of biphasic temporal impulse responses of the input signals. As magnocellular LGN cells usually have biphasic temporal responses (Gielen et al., 1981), this implies that one should observe "inversions" in layer 4Cα, which receives a strong magnocellular input. Interestingly, the experimental data show that "inversions" are *not* seen in layer 4Cα or 4Cβ (Ringach et al., 1997b). Thus, the feedforward model predicts a type of response in layer 4Cα which is not actually seen.

The simulations of the feed-forward model show that the probability of the angle orthogonal to the preferred orientation and the probability of a "blank" behave in essentially the same way. This behavior is seen in some of the V1 data, such as the cell in layer 4C shown in Fig. 1(a). However, as depicted in Fig. 1(b–e), other cells do not show this type of behavior. Furthermore, the feed-forward model also fails to generate the Mexican-hat profiles seen in many cells outside layer 4C.

A Network Model with Lateral Coupling

We now study a network of interacting neurons similar to the one proposed by Ben-Yishai et al. (1995), Somers et al. (1995), and by Carandini and Ringach (1997). We consider a ring of model cells in which each neuron is assumed to have a different preferred angle, and the population of neurons spans the entire orientation range. The connections are all-to-all. As in Carandini and Ringach (1997), we assume that each neuron has both excitatory and inhibitory effects on its neighbors — the population is not divided into excitatory and inhibitory subpopulations (see Fig. 7). We make the further assumption that the strengths of the connections are a function of the difference between neurons's preferred angles; we did not consider connections that depend on the preferred phases of the neurons or the physical positions of the neurons in V1. Labeling neurons by their preferred angles, the strength of the excitatory (inhibitory) connection between two neurons, a_e, satisfies $a_e(\theta_j, \theta_k) = a_e(\theta_j - \theta_k)$.

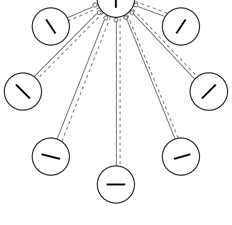

Figure 7: The recurrent feedback network model. One layer of neurons interact with each other via excitatory (solid) and inhibitory (dashed) connections. The preferred orientation for each neuron is indicated schematically with an oriented bar. The connections are all-to-all; only the connections into the 0° neuron are shown here.

The voltage equations for the network are a natural generalization of the integro-differential equation for a single neuron (3). The voltage of the kth neuron satisfies:

$$
\begin{aligned}
\frac{dv_k}{dt}(t) &= -\lambda(v_k(t) - v_r) + DC && (9)\\
&\quad + \int_0^t G(t-s) \iint K(x,y;\theta_k,\phi_k)\, I(x,y,s)\, dx\, dy\, ds && (10)\\
&\quad + C_e \sum_j a_e(\theta_k - \theta_j) \int_0^t G_e(t-s)\, \delta(s - t_s(j))\, ds && (11)\\
&\quad + C_i \sum_j a_i(\theta_k - \theta_j) \int_0^t G_i(t-s)\, \delta(s - t_s(j))\, ds. && (12)
\end{aligned}
$$

C_e and C_i are positive constants that reflect both the balance between excitatory lateral input and inhibitory lateral input as well as the strength of lateral coupling relative to the LGN input. The temporal kernel, G, and spatial kernel, K, are the same as in the feedforward model. All neurons have the same resting potential, $v_r = -70$ mV, are reset to v_r when reaching the upper threshold $v_{th} = -50$ mV, and cannot decrease past the lower barrier of -90 mV. The same stimulus $I(x,y,t)$ is "shown" to all neurons, which are assumed to have spatially overlapping receptive fields. However, since each neuron has a different preferred orientation, they will respond in different ways to the stimulus.

The feedback connections between neurons are assigned temporal dynamics to provide a rough model of synaptic transmission. This is done with temporal kernels $G_e(t)$ and $G_i(t)$ for the lateral connections: $G_e(t) = 20.84\,(t/.0004)^5 e^{-t/.0004}$ and $G_i(t) = 4.17\,(t/.002)^5 e^{-t/.002}$. Both have been chosen so that $\int_0^\infty G_{e,i}(t)\, dt = 1$. We take the inhibitory connections to be slightly slower than the excitatory connections (Somers et al., 1995; Ferster et al., 1996); G_e has its peak at 2 ms while G_i has its peak at 10 ms.

Fig. 8(a) shows these temporal kernels. They are convolved against spike trains, represented as a train of delta functions: $t_s(j)$ is the string of times at which the jth neuron has produced spike events.

Finally, the excitatory effect of the jth neuron's spike train on the kth neuron is weighted by $a_e(\theta_k - \theta_j) > 0$ and the inhibitory effect is weighted by $a_i(\theta_k - \theta_j) < 0$. The functions a_e and a_i are chosen to give a stronger weight to neurons with similar preferred orientations and a weaker weight to neurons with dissimilar preferred orientations. Fig. 8(b) shows the lateral orientation kernels; they have Gaussian shapes. We choose a_i to be broader than a_e (Somers et al., 1995; Ben-Yishai et al., 1995).

In equation (9–12), the following quantities have been normalized:

$$
\int G(s)\, ds = 1 \qquad \sum_{k=1}^{N} \iint K(x,y;0^\circ,0^\circ) \sin\left[\omega(x\cos\theta_k - y\sin\theta_k)\right]\, dx\, dy = N
$$

$$
\int G_e(s)\, ds = 1 \qquad \int G_i(s)\, ds = 1 \qquad \sum_j a_e(\theta_k - \theta_j) = 1 \qquad \sum_j a_i(\theta_k - \theta_j) = -1.
$$

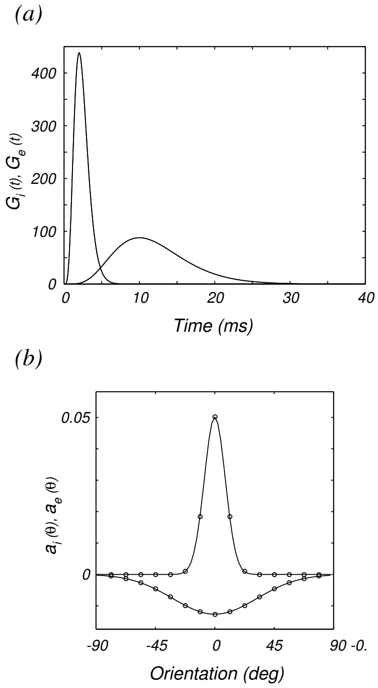

Figure 8: The temporal and spatial weighting of cortical excitation and inhibition in the feedback network model. (*a*) The lateral time kernels for excitation, $G_e(t)$, and inhibition, $G_i(t)$. These are graphs of synaptic current as a function of time after a spike elicited a synaptic potential. The narrow and tall curve is $G_e(t)$ and the shorter and broader curve is $G_i(t)$. (*b*) The orientation dependence of cortical excitation and inhibition are graphed as relative strength vs angle. Excitation is plotted as positive, and is narrower than inhibition. Both are Gaussian functions centered at the preferred angle of the neuron. For 16 neurons in a ring, $a_e(\theta) = .5641\ e^{-(\theta/11.25)^2}$ and $a_i(\theta) = -.1418\ e^{-(\theta/45)^2}$.

Equation (9–12) is of the form

$$\frac{dv_k}{dt} = -\lambda \cdot \text{leakage} + DC + \epsilon A \cdot \text{feed-forward} \tag{13}$$

$$+C_e \cdot \text{lateral excitation} + C_i \cdot \text{lateral inhibition}, \tag{14}$$

Because of the above normalization, the relevant parameters are λ, ϵA, C_e, and C_i. In the computations presented here there is no leakage ($\lambda = 0$). We verified that the observed phenomena persisted in the presence of leakage. If $C_e = C_i$ we call the network "balanced". We define the "lateral coupling ratio" as the ratio between the LGN input and the lateral coupling: $(C_e + C_i)/(2\epsilon A)$. We performed computations for 16 neurons in a ring, with gratings at 80 possible angles and 6 possible phases. The neurons all have the same preferred phase $\phi_k = 0°$.

First we present orientation dynamics obtained in a balanced network with lateral coupling ratio of .2451, which we found produces maximal sharpening. Fig. 9 shows the cross-correlation functions as a function of the correlation time τ. This figure is to be compared to Fig. 4, which is for the strictly feed-forward model. In both figures, the baseline and the orthogonal angle have the same dynamics. Consider the probability distribution of both models at $\tau = 55$ ms. The peak is taller and narrower in Fig. 9 than in Fig. 4. At $\tau = 75$ ms, there is a "Mexican-hat" type profile in Fig. 9, which is absent in Fig. 4. The two figures show different dynamics for $\tau = 85 \ldots 95$ ms: there is an inversion in Figure 4 that is absent in Figure 9.

Fig. 10(a) contrasts the cross-correlation functions of the feed-forward model at $\tau = 54$ ms and the network model at $\tau = 54$ ms. 54 ms is optimal for both the feed-forward computation and the

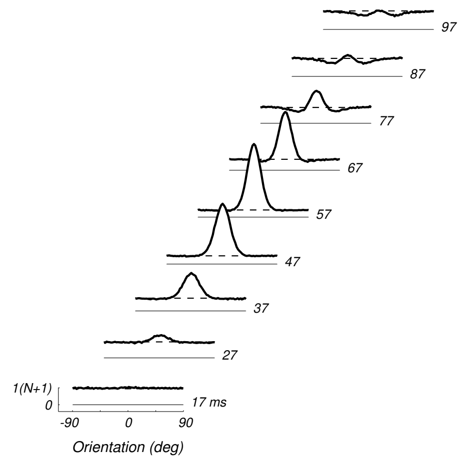

Figure 9: Orientation tuning dynamics of the feedback network model. The plotting conventions are the same as in Figs. 1 and 4. There is no leakage or DC input, $\lambda, DC = 0$. The visual stimulus has $A\epsilon = 416.2$. The network is balanced, $C_e = C_i = 102$, with a lateral coupling ratio of 0.245. The mean interspike time is 366 ms, mean firing rate 2.7 spikes per second. A total of 148,000 spikes were sampled.

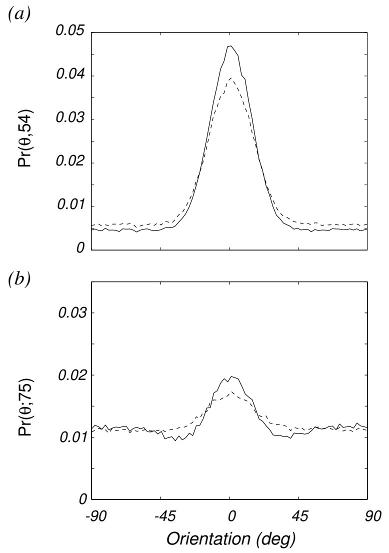

Figure 10: Comparisons of orientation probability distributions for the feedforward and feedback models for the same simulation as in Figure 9. (a) Orientation probability distributions for feedback and feedforward models at their respective peak times. Solid: $Pr(0°;54)$, feedback model. Dashed: $Pr(0°;54)$, strictly feed-forward. (b) Orientation probability distributions $Pr(\theta;75)$ for feedback and feedforward models at 75 ms, time at which a Mexican-hat profile is seen in the feedback model's response. Solid: feedback model. Dashed: strictly feed-forward.

network computation in that $Pr(0°;\tau) < Pr(0°;54)$ for all other correlation times τ. There is a clear sharpening in the profile; $Pr(0°;54) - Pr(blank,54) = .0337$ versus $Pr(0°;54) - Pr(blank,54) = .0420$. The probability of the preferred orientation occurring 54 ms before a spike event has increased by 25%. Fig. 10(b) compares the cross-correlation functions at $\tau = 75$ ms. There is a "Mexican-hat" shaped distribution profile in the simulation of the feedback network. We considered a range of different times-scales in the lateral coupling $G_e(t)$ and $G_i(t)$, including faster inhibition than excitation. We found that this Mexican-hat shaped distribution is robust. The time-scales do affect the degree of tuning; the maximum height of the peak and the time it occurs at are affected.

Next we present the effects of increasing the amount of lateral coupling in a balanced network, $C_e = C_i$. We do this by holding $A\epsilon$ fixed and increasing $C_e = C_i$. Fig. 11 shows mean spike rate

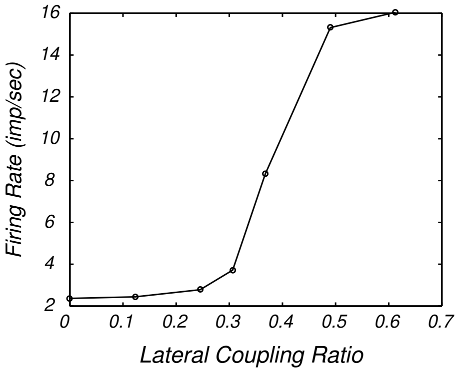

Figure 11: Effects of lateral coupling in the feedback network model on mean firing rate. Seven runs with $\lambda = DC = 0$, $A\epsilon = 416.2$. Balanced network, $C_e = C_i$, lateral coupling ratios $0, .1225, .2451, .3063, .3676, .4901, .6127$. The mean spike rate increases as lateral coupling ratio is increased.

as a function of the lateral coupling ratio. As the lateral coupling ratio increases, the firing rate increases. This is a nonlinear dynamic effect that results from the different time-scales of excitation and inhibition in the temporal kernels G_e and G_i. Specifically, since the network is balanced, every spike event will have both an inhibitory effect and an excitatory effect; the strengths are equal, but the excitation is faster. This leads to a net increase of the mean firing rate. Increased lateral coupling increases the mean firing rate in the ring model and is accompanied by a sharpening of tuning (see below). This is to be contrasted with the feed-forward model where a decrease in the mean firing rate is accompanied by sharpening.

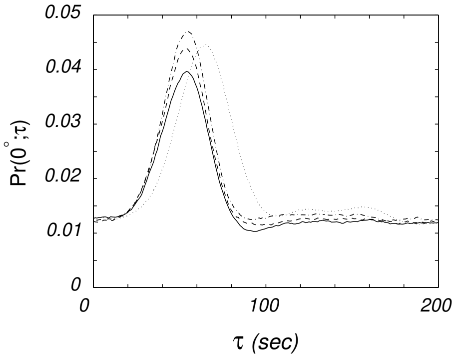

Figure 12: Effects of lateral coupling in a balanced feedback network model on orientation tuning dynamics. $Pr(0°; \tau)$ is plotted as a function of τ. The lateral coupling ratios for these curves are: Solid:= 0,(feedforward); Dashed:= 0.1225; Dot-dashed: = 0.2451; Dotted: = 0.3676; This is for the same simulations as in Figure 11, so the mean firing rates may be gauged from that figure. The peak height of this function is an indicator of the maximal sharpness of orientation tuning.

Fig. 12 shows $Pr(0°, \tau)$ for four of the computations shown in Fig. 11. It is found that lateral coupling provides sharpening. However, we find that when the ratio of lateral coupling to feedforward becomes much larger, the degree of sharpening begins to decrease. From the four runs presented in the figure, the largest degree of sharpening occurs for the lateral coupling ratio of .2451, the value used in Figs. 9 and 10.

The feedback network model considered develops a Mexican-hat profile similar to that observed in some V1 cells. However, the *magnitude* of the effect in the experimental data (see Fig. 1(c–e)) is larger than the one generated by the model (Fig. 9). Furthermore, in the feedback network model, as in the feedforward model, the probability of the angle orthogonal to the preferred one behaves in a similar way as the probability of a “blank”. The feedback network model fails to show "cross-orientation" inhibition effects seen in the data. We considered broader lateral inhibition, including kernels with cross-orientation inhibition. We found that this did not cause the blanks and the orthogonal angle to have different dynamics. The Mexican-hat profile persists, but is weaker.

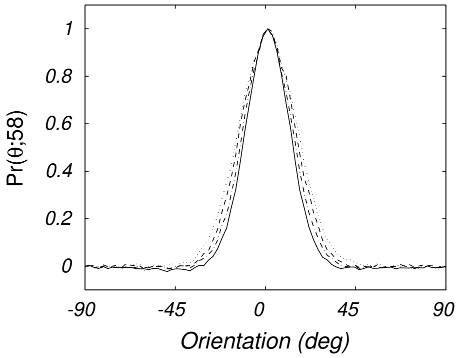

Figure 13: This graph presents normalized orientation probability distributions, at the time of maximum tuning, for four different network computations, with different ratios of excitation to inhibition in the cortical lateral coupling. All simulations have $\lambda = DC = 0$, and $A\epsilon = 416.2$. In order to compare the half-widths, the curves are normalized by first subtracting the probability of the orthogonal orientation and then scaled to have a maximum of one. The values used were zero lateral coupling (dotted); $C_e = C_i = 102$, a balanced network (dashed); $C_i = 3C_e = 306$, net inhibitory (dot-dashed); and $C_i = 12C_e = 612$, even more inhibitory (solid curve).

Up to this point we have calculated the responses of networks which had balanced excitation and inhibition: $C_e = C_i$. We found somewhat better agreement with data in networks with stronger inhibition than excitation, $C_e < C_i$. For these calculations, we first fixed the ratio C_e/C_i. We then varied C_e and C_i, thus varying the lateral coupling ratio. We find sharpening is largest for a particular coupling ratio. Fig. 13 presents the cross-correlation functions, at the time of maximum tuning, for four different network computations where C_e and C_i were chosen in this manner. As in Fig. 12, the balanced network $C_e = C_i \neq 0$ exhibits more sharpening than the feed-forward model $C_e = C_i = 0$. Furthermore, it shows that there is even more sharpening in the networks with $3C_e = C_i$, and $12C_e = C_i$. For the feed-forward model, the maximum value of $Pr(0°, \tau) - Pr(90°, \tau)$ is 0.0278. For the $12C_e = C_i$ computation, this maximum of $Pr(0°, \tau) - Pr(90°, \tau)$ is 0.0424; i.e., lateral coupling which is mainly inhibitory significantly increases the probability of the preferred orientation occurring before a spike.

Discussion

In this study, we evaluated previous models of cortical cells by calculating their orientation tuning dynamics and comparing the result to experimental data. We first considered a well-studied model of cortical simple cells, consisting of a linear spatio-temporal filter and a static integrate-and-fire spike generator. This kind of model has been shown to explain many properties of simple cells when the cortex is under a constant level of luminance and contrast gain control (Movshon et al., 1978; Citron and Emerson, 1983; Jones and Palmer, 1987; McLean and Palmer, 1989; Albrecht and Geisler, 1991; Maex and Orban, 1991; Reid et al., 1991; Tolhurst and Dean, 1991; DeAngelis et al., 1993a,b; McLean et al., 1994; Reid et al., 1997). Therefore, it was of special interest to investigate if such model can account for the phenomena observed in the dynamics of orientation tuning.

We found that this is not the case. The model failed in that (*i*) it did not generate the Mexican-hat profiles observed in some of the responses, (*ii*) it was unable to generate situations where the dynamics of the "blank" and that of the orthogonal orientation differed substantially, and (*iii*) could not be made to generate the active inhibition seen at off-optimal orientations in the late part of the orientation tuning distributions. The model, however, exhibited "inversions". This is somewhat puzzling because "inversions" are not seen in layer 4C (Ringach et al., 1997b). Thus the feedforward model, while capturing an aspect of the responses seen outside layer 4C, actually predicts an effect that is not seen in this layer. It appears, therefore, that the behavior of cells in layer 4C (in particular those in 4Cα which receive biphasic temporal input from the LGN magnocellular layers) are not explained by the feedforward model. One possibility is that 4Cα cells do not fit the description of a space-time separable feedforward kernel, as some of these neurons are known to be directionally selective (Hawken et al., 1988).

In the second part of the paper we considered a class of models that has been proposed in the literature as a putative mechanism for generating sharp orientation selectivity in the cortex (Andrews, 1965, 1967; Ben-Yishai et al., 1995; Somers et al., 1995; Douglas et al., 1995; Carandini and Ringach, 1997). This class of models are described as a network model with "lateral" feedback, in which neurons with similar preferred orientations excite each other while neurons with different orientations preferences inhibit each other. We found that such a model generates orientation distributions which are more similar to those observed in the experimental data. In particular, the model generates unimodal distributions early in the response (for $\tau < 50$ msec) and Mexican-hat shaped distributions later in the response (for $\tau > 60$ msec). However, it is clear that the magnitude of the Mexican-hat profiles is somewhat smaller in the simulations than in the real data. While the feedback model seems to be a step in the right direction, there are additional features of the data that remain unexplained. The network model, like the feedforward model, fails to provide an adequate explanation for the fact that, in most cases, the orthogonal orientation and the "blank" behave differently. In appears that in order to explain the full range of dynamical effects seen in the data we need to refine our ideas and models of neuronal interactions in V1 cortical circuits.

As a next step in this direction we plan to measure experimentally the responses of LGN cells when stimulated with the image sequences used in the experiments of Ringach et al (1997). It might be the case that the failure of the models we considered is a consequence of a poor mathematical representation of the LGN input to the cortex. By measuring the dynamics of LGN cells, and incorporating this data into our models, we hope to advance towards our goal of establishing the relative role of thalamic and cortical signals in the generation of orientation tuning.

Acknowledgments: M. P. was partially supported by an NSF post-doctoral fellowship and NSF-DMS–9709128. D. L. R. and R. S. were supported by NSF-IBN-9720305, NIH-EY01472, and a Sloan Foundation grant to NYU's Program in Theoretical Neuroscience. We thank M. J. Hawken, D. McLaughlin, D. Tranchina, and J. Wielaard for their valuable comments and criticisms.

References

Albrecht DG, Geisler WS (1991) Motion selectivity and the contrast-response function of simple cells in the visual cortex. Vis Neurosci 7:531–546.

Andrews DP (1965) Perception of contours in the central fovea. Nature 205:1218–1220.

Andrews DP (1967) Perception of contour orientation in the central fovea. Part I. Short lines. Vis Res 7:975–997.

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

Carandini M, Mechler F, Leonard CS, Movshon JA (1996) Spike train encoding by regular-spiking cells of the visual cortex. J Neurophysiol 76(5):3425–3441.

Carandini M, Ringach DL (1997) Predictions of a recurrent model of orientation selectivity. Vis Res 37:3061–3071.

Citron MC, Emerson RC (1983) White noise analysis of cortical directional selectivity in cat. Brain Res 279:271–277.

Connors BW, Gutnick MJ, Prince DA (1982) Electrophysiological properties of neocortical neurons in vitro. J Neurophysiol 48:1302–1320.

Das A (1996) Orientation in visual cortex: a simple mechanism emerges. Neuron 16:477–480.

Daugman JG (1985) Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. J Opt Soc Am A 2(7):1160–9.

de Boer E, Kuyper P (1968) Triggered correlation. IEEE Trans on Biomedical Engineering 15:169–179.

DeAngelis GC, Ohzawa I, Freeman RD (1993a) Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. ii. linearity of temporal and spatial summation. J Neurophysiol 69:1118–35.

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

Douglas RJ, Koch C, Mahowald M, Martin KA, Suarez HH (1995) Recurrent excitation in neocortical circuits. Science 269:981–985.

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

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

Gielen CC, van Gisbergen JA, Vendrik AJ (1981) Characterization of spatial and temporal properties of monkey LGN Y-cells. Biol Cyb 40(3):157–170.

Gut A (1988) Stopped random walks : limit theorems and applications. Springer Verlag, New York.

Hawken MJ, Parker AJ, Lund JS (1988) Laminar organization and contrast sensitivity of direction-selective cells in the striate cortex of the old world monkey. J Neurosci 8:3541–3548.

Heeger DJ (1992a) Half-squaring in responses of cat striate cells. VISUAL NEUROSCI 9(5):427–443.

Heeger DJ (1992b) Normalization of cell responses in cat striate cortex. VISUAL NEUROSCI 9(2):181–197.

Hubel DH, Wiesel TN (1962) Receptive fields, binocular interaction and functional architecture of cat's visual cortex. J Physiol Lond 160:106–154.

Hubel DH, Wiesel TN (1968) Receptive fields and functional architecture of monkey striate cortex. J Physiol Lond 195:215–245.

Jones JP, Palmer LA (1987) The two-dimensional spatial structure of simple receptive fields in the cat striate cortex. J Neurophysiol 58:1187–1258.

Maex R, Orban GA (1991) Subtraction inhibition combined with a spiking threshold accounts for cortical direction selectivity. Proc Natl Acad Sci USA 88(9):3549–3553.

Maex R, Orban GA (1996) Model circuit of spiking neurons generating directional selectivity in simple cells. J Neurophysiol 75:1515–1545.

Marcelja S (1980) Mathematical description of the responses of simple cortical cells. J Opt Soc Am 70(11):1297–1300.

Marmarelis PN, Marmarelis VZ (1978) Analysis of Physiological Systems: The White Noise Approach. NY: Plenum Press.

McCormick DA, Connors BW, Lighthall JW (1985) Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. J Neurophysiol 54:782–806.

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

McLean J, Raab S, Palmer LA (1994) Contribution of linear mechanisms to the specification of local motion by simple cells in areas 17 and 18 of the cat. Vis Neurosci 11:271–94.

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

Palmer LA, Davis TL (1981) Receptive-field structure in cat striate cortex. J Neurophysiol 46(2):260–76.

Reid RC, Alonso JM (1996) The processing and encoding of information in the visual cortex. Current Opinion in Neurobiology 6:475–480.

Reid RC, Soodak RE, Shapley RM (1991) Directional selectivity and spatiotemporal structure of receptive fields of simple cells in cat striate cortex. J Neurophysio 66:505–29.

Reid RC, Victor JD, Shapley RM (1997) The use of m-sequences in the analysis of visual neurons: linear receptive field properties. Vis Neurosci 14:1015–1027.

Ringach DL, Hawken MJ, Shapley R (1997a) Dynamics of excitatory and inhibitory mechanisms shaping the orientation tuning of neurons in V1. Soc Neurosci Abstr 23:1544, part 2.

Ringach DL, Hawken MJ, Shapley R (1997b) Dynamics of orientation tuning in macaque primary visual cortex. Nature 387:281–284.

Ringach DL, Hawken MJ, Shapley R (1998) Spatial-phase dependent and independent response components of oriented neurons in macaque V1. In Invest. Ophthal. & Vis. Sci. (Suppl.), volume 39, p. S683.

Ringach DL, Sapiro G, Shapley R (1997c) A subspace reverse-correlation technique for the study of visual neurons. Vis Res 37:2455–2464.

Somers DC, Nelson SB, Sur M (1995) An emergent model of orientation selectivity in cat visual cortical simple cells. J Neurosci 269:5448–5465.

Sompolinsky H, Shapley R (1997) New perspectives on the mechanisms for orientation selectivity. Current Opinion in Neurobiology 7:514–522.

Tolhurst DJ, Dean AF (1991) Evaluation of a linear model of directional selectivity in simple cells of the cat's striate cortex. Vis Neurosci 6(5):421–428.

Victor JD (1992) Nonlinear systems analysis in vision: overview of kernel methods. In Pinter R, Nabet B (eds.), Nonlinear vision: Determination of neural receptive fields, function and networks, volume 1, pp. 1–37, CRC Press, Cleveland, Ohio.

Wörgötter F, Niebur E, Koch C (1991) Quantification and comparison of cell properties in cat's striate cortex determined by different types of stimuli. J Neurophysiol 66:1163–1176.

Appendix

Effects of leakage, phases, temporal kernels

Equation with δ-function forcing and excitatory responses

We consider the feedforward model (3) where the temporal kernel G is a δ-function. The membrane voltage is then the solution of

$$\frac{dv}{dt}(t) = -\lambda(v(t) - v_r) + r(t), \tag{15}$$

with the resting potential, v_r, the upper voltage threshold, and lower barrier as described for the feed-forward model. We assume the DC input is zero. The forcing of the voltage equation, $r(t)$, is piecewise constant in time, allowing us to solve the equation explicitly between spike events. We denote the times at which the forcing changes by the sequence $t_0, t_1, \ldots, t_n, \ldots$. For $t \in [t_n, t_{n+1})$,

$$v(t) = \begin{cases} v(\tilde{t}_n) + r(t_n)(t - \tilde{t}_n) & \text{if } \lambda = 0 \\ v(\tilde{t}_n) + \frac{r(t_n) - \lambda(v(\tilde{t}_n) - v_r)}{\lambda}\left(1 - e^{-\lambda(t - \tilde{t}_n)}\right) & \text{if } \lambda \neq 0 \end{cases}$$

Here, $\tilde{t}_n = \max(t_n, t_{spike})$ where t_{spike} is the time of the last spike and $v(t_{spike}) = v_r$.

As before we use a spatial kernel (5) with preferred angle $\theta_K = 0^\circ$ and preferred phase $\phi_K = 0^\circ$. Consider first the equation with non-negative responses $r(\theta) \geq 0$. This case arises if the phase of the gratings is fixed and such that they yield only nonnegative responses.

As we are considering the case of N possible angles and one fixed phase, there are $N+1$ possible values for $r(t_n)$. If there is no leakage, $\lambda = 0$, or if there is weak leakage, $0 < \lambda \leq \min\{r(t_n)\}/(-v_r)$, one can view the voltage's evolution before a spike as a directed random walk (i.e., steps of positive or zero length). The spike occurs at the first time the random walk crosses the upper threshold, after which the state is reset to the resting potential. This is a first passage-time problem.

We first consider the "instantaneous" correlation function, $Pr(\theta; 0)$. In constructing this function, a histogram is made by considering the angles of the images that drove the voltage past the threshold. As the voltage is immediately reset once it crosses the threshold, the voltage at time $t - \Delta t$ must have been below the threshold. Hence, the "instantaneous" correlation function must exclude all angles that generate negative steps. In other words, images which yield $r < 0$ would be excluded, as would images which yield responses too weak to overcome the leakage. One can prove the following:

Theorem: Let the voltage be governed by equation (15).
No leakage (see, for example, Gut (1988)) If $\lambda = 0$ and $\{r(t_n)\} \geq 0$, then the instantaneous cross-correlation function is determined by the spatial response function $r(\theta)$,

$$Pr(\theta; 0) = \frac{r(\theta)}{\sum r(\theta_i)}$$

Furthermore, the process immediately decorrelates, i.e.,

$$Pr(\theta; k\Delta t) = \frac{1}{N+1} \qquad k > 0.$$

Weak leakage If $0 < \lambda \leq \min\{r(t_n)\}/(-v_r)$, then the instantaneous cross-correlation function is determined by the response function, the upper threshold, the resting potential, and the leakage:

$$Pr(\theta; 0) = \frac{r(\theta) + \lambda v_r}{\sum (r(\theta_i) + \lambda v_r)}$$

Again, the process immediately decorrelates: $Pr(\theta; \tau) = 1/(N+1)$ for $\tau > 0$.

The theorem implies that weak leakage causes a "sharpening" of the instantaneous orientation distribution. Such leakage-induced sharpening is observed generally in all our numerical simulations. However, the above case is the only case for which we can prove sharpening occurs. In addition, we note that the leakage-induced sharpening is obtained by translating $r(\theta)$ by a factor λv_r and normalizing. The profile does not change its convexity by such a transformation. Thus, leakage effects cannot be responsible for the formation of a Mexican hat profile.

The theorem applies to the case where the leakage is weak relative to the forcing. This is to be contrasted with the case where the leakage is strong relative to the forcing. Fig. 14(a) shows a pair of numerically computed instantaneous correlation functions, as well as the normalized response function for this case. The exclusion of angles which yield weak responses is clear. Fig. 14(b) shows computational evidence that strong leakage causes a loss of instantaneous decorrelation.

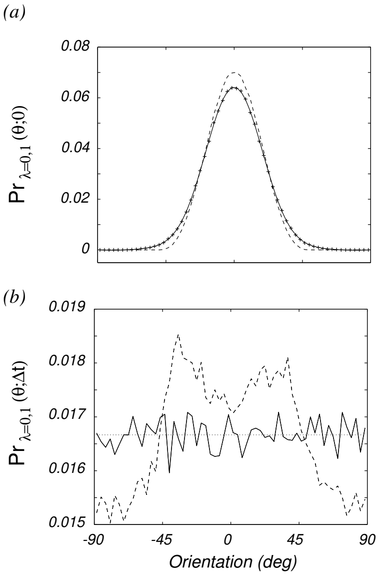

Figure 14: All images shown are in phase, $\phi_i = 0°$. (a) The effect of leakage on orientation tuning bandwidth. Solid: normalized response function, crosses: $Pr_{\lambda=0}(\theta;0)$, the instantaneous correlation function for the no leakage case; dashed: $Pr_{\lambda=1}(\theta;0)$, the instantaneous correlation function for the strong leakage case. (b) The effect of leakage is to introduce correlation. Solid: $Pr_{\lambda=0}(\phi;\Delta t)$ the cross-correlation function for $\tau = \Delta t$ and $\lambda = 0$, dashed: $Pr_{\lambda=1}(\theta;\Delta t)$ the cross-correlation function for strong leakage, $\lambda = 1$, dotted: $1/(N+1)$ Here, $A\epsilon = 260.2$, with $300,000$ spikes sampled, mean firing rate of 3.58 spikes/sec.

Equation with δ-function forcing and both excitatory and inhibitory responses

If we allow again gratings to have random spatial phases then the K can also produce inhibitory responses. In the absence of leakage, the voltage evolves like an undirected random walk with, at most, $MN + 1$ possible step lengths. If the average response over all spatial phases is zero for each angle and if there is no lower barrier for the voltage, then the random walk has no positive drift. The presence of a lower barrier introduces a effective positive drift even when the average response is zero.

The possibility of inhibitory responses causes the loss of instantaneous decorrelation for $\tau > 0$ even when there is no leakage. This is seen in a simple example: assume there are only two possible responses $+1$ if the angle shown is $0°$ and -1 if the angle is not $0°$. For such responses, *both* the instantaneous and the $\tau = \Delta t$ cross-correlation functions are 1 at $0°$ and zero elsewhere.

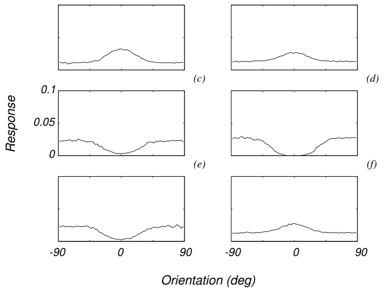

Figure 15: The images shown can be at any of six equally spaced spatial phases. Plotted here are the phase-angle correlation functions, $C(\theta_i, \phi_j; \Delta t)$, as a function of the orientations θ_i for the six phases ϕ_j. Each C has been normalized to sum to one; allowing them to be interpreted as the relative frequency that an image at angle θ and phase ϕ occurs before a spike event. Here, $A\epsilon = 1040.9$, leakage present, $\lambda = .6$, and $160,500$ spikes sampled.

To gain insight into the type of effects one can expect when both excitatory and inhibitory responses are present, we plot the normalized counts in $C(\theta_i, \phi_j)$ for different values of the phase ϕ_j as a function of the orientations θ_i. We call these "phase-angle correlation functions". The phase-angle correlation functions for $\tau = \Delta t$ are non-zero for almost all angles and all phases (see Fig. 15). If (θ_0, ϕ_0) has a non-zero probability of having been shown the time step before a spike, this means that however negative the response $r_{\phi_0}(\theta_0)$ might have been, there was a larger positive response at the next time step, since otherwise there would have been no way to cross the voltage threshold and generate a spike. Hence, in the absence of leakage, the only way for the $\tau = \Delta t$ phase-angle correlation function to be zero for (θ_0, ϕ_0) is if $-r_{\phi_0}(\theta_0) \geq R_{max}$, where R_{max} is the largest possible positive response. In such a case, the angle θ_0 would be excluded from the ϕ_0th phase-angle cross-correlation function for $\tau = \Delta t$. As the responses shown in Figure 3 are symmetric in ϕ, the only exclusion is $(\theta_0, \phi_0) = (0, 180°)$. In the presence of leakage, angles near zero can also be excluded from the 180° phase-angle cross-correlation function. Such behavior is shown in Figure 15(D).

This observation generalizes: if there is a very large negative response to a certain image, such that $-r_{\phi_0}(\theta_0) \geq kR_{max}$, then this angle will be excluded from the ϕ_0th phase-angle cross-correlation function for correlation times $\tau = 0, \ldots, k\Delta t$. Furthermore, one cannot have non-zero probability of showing an image at angle θ for $\tau = k\Delta t$ before a spike and yet have a zero probability for the same angle at $\tau = (k+1)\Delta t$ before a spike. Thus, as the correlation time τ increases the number of angles for which there is zero probability decreases.