Visual Neuroscience (1997), **14**, 293–309. Printed in the USA.

Comparison of contrast-normalization and threshold models of the responses of simple cells in cat striate cortex

D.J. TOLHURST[1] AND D.J. HEEGER[2]

[1]The Physiological Laboratory, Downing Street, Cambridge, CB2 3EG, U.K.
[2]Department of Psychology, Stanford University, Stanford, CA 94305

(RECEIVED April 8, 1996; ACCEPTED August 5, 1996)

Abstract

In almost every study of the linearity of spatiotemporal summation in simple cells of the cat's visual cortex, there have been systematic mismatches between the experimental observations and the predictions of the linear theory. These mismatches have generally been explained by supposing that the initial spatiotemporal summation stage is strictly linear, but that the following output stage of the simple cell is subject to some contrast-dependent nonlinearity. Two main models of the output nonlinearity have been proposed: the *threshold model* (e.g. Tolhurst & Dean, 1987) and the *contrast-normalization model* (e.g. Heeger, 1992*a*,*b*). In this paper, the two models are fitted rigorously to a variety of previously published neurophysiological data, in order to determine whether one model is a better explanation of the data. We reexamine data on the interaction between two bar stimuli presented in different parts of the receptive field; on the relationship between the receptive-field map and the inverse Fourier transform of the spatial-frequency tuning curve; on the dependence of response amplitude and phase on the spatial phase of stationary gratings; on the relationships between the responses to moving and modulated gratings; and on the suppressive action of gratings moving in a neuron's nonpreferred direction. In many situations, the predictions of the two models are similar, but the contrast-normalization model usually fits the data slightly better than the threshold model, and it is easier to apply the equations of the normalization model. More importantly, the normalization model is naturally able to account very well for the details and subtlety of the results in experiments where the total contrast energy of the stimuli changes; some of these phenomena are completely beyond the scope of the threshold model. Rigorous application of the models' equations has revealed some situations where neither model fits quite well enough, and we must suppose, therefore, that there are some subtle nonlinearities still to be characterized.

Keywords: Visual cortex, Simple cells, Thresholds, Linear summation, Contrast normalization, Nonlinear responses, Response suppression, Directional selectivity

Introduction

Hubel and Wiesel (1959) provided the first and lasting definition of simple cells in cat striate cortex. These are neurons whose receptive fields can be mapped into distinct excitatory and inhibitory subregions and, most importantly, the geometry of the most effective visual stimuli can be predicted from the geometry of these subregions. Phrased more formally, this is a proposal that simple cells show linear spatiotemporal summation.

Numerous studies of simple cell behavior have investigated the extent to which summation can indeed be considered to be linear. These studies have included tests of the additivity of the responses to combinations of simple stimuli (Maffei et al., 1979; Schumer & Movshon, 1984; Tolhurst & Dean, 1987; Pollen et al., 1988), comparisons of the responses to moving and counterphase-modulated gratings (Reid et al., 1987, 1991; Albrecht & Geisler, 1991; Tolhurst & Dean, 1991; DeAngelis et al., 1993), and comparisons of receptive-field maps with the inverse Fourier transforms of the spatial-frequency sensitivity curves (R. De Valois et al., 1978; Movshon et al., 1978*a*; Andrews & Pollen, 1979; Maffei et al., 1979; Kulikowski & Bishop, 1981; Glezer et al., 1982; Webster & R. De Valois, 1985; Jones & Palmer, 1987; Shapley et al., 1991; DeAngelis et al., 1993).

In almost every study, there have been systematic mismatches between the experimental observations and the predictions of the linear theory, showing that the behavior of simple cells cannot be thought of as strictly linear. In a small proportion of simple cells, the observed nonlinearities are of such a form that they can only result from fundamental nonlinearities of spatial summation (Movshon et al., 1978*a*; Dean & Tolhurst, 1983; Mullikin et al., 1984; Tolhurst & Dean, 1990). For instance, some simple cells respond with an unmodulated elevation of activity to high spatial-frequency

Reprint requests to: D.J. Heeger, Department of Psychology, Stanford University, Stanford, CA 94305.

gratings. In the same neurons, the sign of the response to a bar stimulus in some parts of the receptive field is *not* reversed when the sign of the stimulus is reversed from bright to dark. These nonlinearities are just as powerful at low contrasts as at high contrasts and they are similar to the nonlinearities reported by Enroth-Cugell and Robson (1966) for Y-cells in the cat's retina. It is widely believed that there must be a push–pull arrangement of excitatory and inhibitory synapses impinging on the simple cell (Palmer & Davis, 1981; Glezer et al., 1982; Tolhurst & Dean, 1987, 1990; Ferster, 1988; Ferster & Jagadeesh, 1992), and it may be that these nonlinearities of spatial summation arise from mismatches in this push–pull mechanism (Tolhurst & Dean, 1987, 1990; Atick & Redlich, 1990). This kind of nonlinear behavior will not be considered any further in this paper.

For most simple cells, on the other hand, the failures of the linear model vary in magnitude depending upon stimulus contrast (e.g. Tolhurst & Dean, 1987, 1990, 1991; Albrecht & Geisler, 1991; Albrecht, 1995). Perhaps, the initial spatiotemporal summation stage *is* strictly linear in these simple cells, but the following output stage is subject to some contrast-dependent nonlinearity. Two main models of the output nonlinearity have been proposed to account for many of the discrepancies between measurement and linear prediction. These are the *threshold model* (e.g. Tolhurst & Dean, 1987) and the *contrast-normalization model* (e.g. Heeger, 1992*a*,*b*).

The threshold model

The threshold model was introduced formally by Movshon et al. (1978*a*) to rationalize some mismatches between their experimental observations and the linear predictions. The model was developed most thoroughly by Schumer and Movshon (1984), Tolhurst and Dean (1987), and Tadmor and Tolhurst (1989), although many other authors have called on it naturally to explain discrepancies in their results (e.g. Andrews & Pollen, 1979; Kulikowski & Bishop, 1981; Glezer et al., 1982; R. De Valois et al., 1985; Robson et al., 1988; DeAngelis et al., 1993).

In this model, it is presumed that the simple cell sums the influences of the light falling in the different parts of its receptive field to give a strictly linear underlying response (continuous fluctuations of membrane potential), which must then exceed some threshold level of depolarization before an overt response of action potentials is seen by the experimenter. For example, a low contrast stimulus would evoke no action potentials. But once the contrast exceeded some threshold value a response would be evident, and its amplitude would increase with further increases in contrast (Ikeda & Wright, 1974; Tolhurst et al., 1981; Dean, 1981*a*).

As well as the threshold, another nonlinearity is usually present in the responses of simple cells: response amplitudes saturate at moderate to high contrasts (e.g. Maffei & Fiorentini, 1973; Albrecht & Hamilton, 1982; Ohzawa et al., 1982). Although the threshold model *per se* does not explain such response saturation, we can add a *post hoc* maximum firing rate to the model, forcing it to exhibit this behavior.

This threshold model seems to be justified by the known biophysics of action potential generation in neurons of the cerebral neocortex. In response to injected current, a threshold depolarization value must indeed be exceeded before action potentials are generated and, due to the refractory period between action potentials, there is indeed a maximal rate of firing in response to current injection (e.g. Koike et al., 1970; Ogawa et al., 1981; McCormick et al., 1985; Mason & Larkman, 1990; Jagadeesh et al., 1992).

However, there are several problems with this model. First, the firing-rate limit imposed by the refractory period is generally much higher than that typically measured in response to visual stimuli. Second, the model fails to explain that saturation begins at a given contrast rather than at a given (maximal) firing rate (Albrecht & Hamilton, 1982; Li & Creutzfeldt, 1984; Skottun et al., 1987; Albrecht & Geisler, 1991; Tolhurst & Dean, 1991; Geisler & Albrecht, 1992; Heeger, 1992*a*; Carandini & Heeger, 1994; Albrecht, 1995). Third, the threshold model does not provide any explanation at all for the phenomenon of *nonspecific suppression*; a simple cell's response to a preferred stimulus can be reduced by the simultaneous presence of almost any other visual stimulus, which may on its own have no overt effect (e.g. Maffei & Fiorentini, 1976; Dean et al., 1980; Hammond & MacKay, 1981; Morrone et al., 1982; K. De Valois & Tootell, 1983; Gulyas et al., 1987; Bonds, 1989; Bauman & Bonds, 1991; DeAngelis et al., 1992).

The contrast-normalization model

Contrast normalization was originally proposed by Robson (1988) and Bonds (1989) specifically to provide explanations for response saturation and nonspecific suppression. The model has been expanded and formalized by Heeger (1991, 1992*a*,*b*, 1993), by Albrecht and Geisler (1991), and by Carandini and Heeger (1994) who have shown that it is capable, in principle, of explaining a wide variety of phenomena including many of those which were originally attributed to a threshold (see also DeAngelis et al., 1993; Jagadeesh et al., 1993; Tolhurst & Heeger, 1997; Carandini et al., 1997; Nestares & Heeger, 1997). The overall motivation of the normalization model and its detailed synaptic mechanisms are surprisingly similar to Marr's (1970) general theory of cerebral neocortex, and to Grossberg's theoretical work on nonlinear neural networks (for review, see Grossberg, 1988).

The normalization model supposes that the simple cell's underlying linear response is subject to two nonlinearities. First, the underlying linear response is normalized, divided by a quantity proportional to the pooled activity of a large number of other neurons. The response of each neuron is no longer dependent solely on the contrast of stimulus components that it prefers; rather, the response is normalized or rescaled with respect to the total contrast or energy within the stimulus. This model explains response saturation because the divisive suppression increases with stimulus contrast. It also explains nonspecific suppression because a given model neuron is suppressed by many other neurons, including those with very different stimulus preferences. The second nonlinearity is that the neuron's response is halfwave rectified and then squared, to give "half-squaring." This expansive behavior, at low stimulus contrast, is very similar to a threshold nonlinearity.

Sometimes, it is the expansiveness of response at low contrasts which is the crucial feature for modelling the behavior of real neurons; normalization *per se* does not contribute when total stimulus contrast remains constant throughout an experiment. However, in most of the cases, it is the combination of expansiveness and normalization that is required to model a neuron's behavior. Although the model includes expansiveness as well as normalization, it is convenient to refer to the model as the *normalization model*.

The plan of this paper

In this paper, we compare these two models of the output nonlinearity of simple cells head-to-head. We fit the two models to a

variety of neurophysiological data, most of which have been published previously by one of us (D.J.T.) and which were supposed to show the success of the threshold model. We re-evaluate data on the additivity of responses to paired bar stimuli (Tolhurst & Dean, 1987); on the relationship between the line-weighting function and the inverse Fourier transform of the spatial-frequency tuning function (Movshon et al., 1978*a*; Tadmor & Tolhurst, 1989); on the amplitude of response to sinusoidal gratings of different spatial phases (Robson et al., 1988); and on the relationships between the responses to modulated and moving gratings (Tolhurst & Dean, 1991). We show that the contrast-normalization model usually fits the data at least as well, and often better, than the threshold model. We also show that the normalization model is naturally able to account for some phenomena that were previously unexplained or unnoticed. Lastly, we present some previously unpublished results, related to the study of Dean et al. (1980) on the divisive inhibition produced by stimuli moving in the neuron's nonpreferred direction. The normalization model explains a variety of phenomena in these suppression experiments that find no immediate explanation in the threshold model.

Methods

Neurophysiological procedures

The experimental methods have mostly been described in the papers where the data were first presented; general details are found in Movshon et al. (1978*a*,*b*), Tolhurst and Thompson (1981), and Dean and Tolhurst (1983). Simple cells were recorded extracellularly from the *area centralis* representation of area 17 of adult cats using tungsten-in-glass microelectrodes. The neurons were classified after Hubel and Wiesel (1959) using the quantitative criteria discussed by Dean and Tolhurst (1983). The cats were anesthetized by i.v. infusion of barbiturates, supplemented by ventilation with nitrous oxide. The animals were also paralyzed by i.v. infusion of gallamine triethiodide to prevent eye movements. The state of anesthesia was assessed by monitoring heart rate and the EEG.

Visual stimuli were presented on a bright monochrome raster display, and a computer compiled peristimulus time histograms (PSTH) of the action potentials generated in response. A particular stimulus was presented for only a few seconds or temporal cycles at a time, interleaved at random with short presentations of the other stimuli in the experiment. These short epochs were repeated several times, so that the final PSTH might represent the summed response to 50–200 temporal cycles. Sinusoidal gratings were generally of the optimal spatial frequency and orientation; the temporal frequency of modulation or movement was usually 2 Hz. Bar stimuli were of optimal orientation, and their contrast was modulated sinusoidally in time at 1 Hz.

The metric of response was usually the amplitude of the Fourier component in the PSTH whose frequency was the same as the modulation frequency of the stimulus. Response is expressed as impulses per second (ips). In some Figs. (3, 5 and 6), responses are plotted as if they were negative. This simply indicates a shift of 180 deg in the temporal phase of the response; the response amplitudes are, of course, positive.

The directional suppression experiment

Figs. 8–10 show the results of experiments that have not been described before in detail (after Dean et al., 1980). Thirty-five experiments were performed on 17 simple cells which had profound directional selectivity, so that they gave almost no response to sinusoidal gratings moving in the nonpreferred direction. The neurons were activated by gratings of optimal spatial frequency and orientation moving in the preferred direction at 2 Hz. This was done at a variety of contrasts. Suppression of responses was caused by the simultaneous presentation of a grating of the same spatial frequency and orientation moving in the nonpreferred direction at 2.5 Hz. The contrast of the suppressing grating could also be varied. In a typical experiment, control conditions with just the activating grating present would be interleaved randomly with test conditions where the suppressing grating was also present.

In this experiment, a stimulus presentation lasted 4.5 s and the results from the initial 0.5 s were discarded. A PSTH was compiled over the next 4 s, when there were eight complete cycles of the activating grating and ten complete cycles of the suppressing grating if it was present. Thus, the temporal phase difference of the two gratings went through two complete cycles. Each time a given pair of gratings was presented together, the starting phase difference was randomized. Thus, over the many repeats of a full experiment, we would expect all possible phase differences to be represented. The results were analyzed by calculating the temporal phase and amplitude of the 2-Hz Fourier component in the PSTH; this represents "the response to the activating grating."

Definitions

Consider that the simple cell gives an underlying linear response to visual stimuli—R. This underlying response is then subject to one or more output nonlinearities, resulting in an overt or measured response—$\bar{R}$. These two quantities are related by

$$\bar{R} = f(R) \tag{1}$$

$$R = f^{-1}(\bar{R}) \tag{2}$$

where $f(\,)$ is a nonlinear output function and $f^{-1}(\,)$ is the inverse of that function. The nonlinearity need not be completely reversible; it may not always be possible to estimate R from $\bar{R}$.

The threshold model

The threshold model is described most easily for an experiment in which overt response amplitude ($\bar{R}$) is measured as a function of stimulus contrast (c). If R is the linear underlying response, then

$$R = K \cdot c \tag{3}$$

where K, the neuron's responsivity, depends on the spatial and temporal parameters (e.g. spatial frequency, temporal frequency, orientation, etc.) of the stimulus. Now, an overt response $\bar{R}$ will be generated only when a threshold response value T is exceeded:

$$\bar{R} = K \cdot c - T \quad \text{for } K \cdot c > T \tag{4a}$$

$$\bar{R} = 0 \quad \text{for } K \cdot c \leq T \tag{4b}$$

At high contrasts, the neuron's response saturates and so we must add another *post hoc* limit to eqn. (4a):

$$\bar{R} = \bar{R}_{\max} \quad \text{for } (K \cdot c - T) \geq \bar{R}_{\max} \tag{4c}$$

where $\bar{R}_{\max}$ is the response amplitude at saturation.

The threshold model was fit to the experimentally measured response amplitude as a function of contrast (Fig. 2 open circles,

Fig. 4, and Fig. 7), by searching for the best-fitting values of the three parameters (K, T, and $\bar{R}_{max}$) in eqn. (4). Details of the fitting algorithms are given below. The value of T is expressed as a percentage of the largest measured response to any of the stimuli in the experiment.

The two-bar interaction data (filled symbols in Fig. 2) were fit relatively easily for the threshold model, since the linear theory predicts a parallel upward shift of the control relationship between response and contrast for single bars. Compared to the control data [eqn. (4)], one extra variable parameter was needed; i.e. the amount of the upward shift, equal to the underlying linear response to the second (the dark) bar, R_d.

$$\bar{R} = K \cdot c_b - T + R_d \quad \text{for } \bar{R} > 0 \tag{5}$$

where c_b is the contrast of the first (the bright) bar. In the example shown in Fig. 2, the temporal phases of the responses were very nearly the same for the two bars that were used (cf. Tolhurst & Dean, 1987). Hence, we can use scalar rather than vector arithmetic in eqn. (5).

For the other experiments, the contrast term (c) in eqns. (3) and (4) had to be replaced with the appropriate prediction of the linear model specific to the experiment. For instance, the purely linear model predicts that the amplitude of the responses to modulated gratings as a function of stimulus spatial phase (ϕ), depicted in Fig. 5 would obey

$$R = K \cdot |\cos(\phi - \Phi)| \tag{6}$$

where R is the linear response amplitude, and Φ is the stimulus spatial phase that evokes the maximal response. Equations analogous to eqns. (4a) and (4b) were used to account for the threshold:

$$\bar{R} = K \cdot |\cos(\phi - \Phi)| - T \quad \text{for } K \cdot |\cos(\phi - \Phi)| > T \tag{6a}$$

$$\bar{R} = 0 \quad \text{for } K \cdot |\cos(\phi - \Phi)| \leq T \tag{6b}$$

In this experiment, the form of the results was not influenced by response saturation; hence, the saturation parameter ($\bar{R}_{max}$) was ignored. Thus, the fitting algorithm sought the best values for K, T, and Φ in the above equation. In addition, the fitting routine kept track of the temporal phase of the responses and used them as additional constraints for the fits. In other words, the routine recognized that two stimulus spatial phases on either side of the "null position" might evoke virtually the same response amplitudes but with response phases shifted by 180 deg.

To fit the polar plot of Fig. 6 (response amplitude *vs.* response phase for temporally modulated grating stimuli), we needed to find the best-fitting values for both the major and the minor axes of the putative ellipse. Hence, we used a total of four parameters: K_1 (response amplitude at the best stimulus spatial phase), K_2 (response amplitude at the worst stimulus spatial phase), T (threshold), and Φ (best stimulus spatial phase).

The contrast-normalization model

According to the contrast-normalization model with strict half-squaring, the overt response $\bar{R}$ of a simple cell to any stimulus is given by

$$\bar{R} = \frac{k \cdot H_0^2}{S^2 + \sum_i H_i^2} \tag{7}$$

where k and S are constants for the particular neuron, H_0 is the halfwave-rectified version of the underlying linear response (R_0) of the neuron from which we are recording, and H_i represents the halfwave-rectified responses of all the neurons in the normalization pool. This equation can be written in different ways to match different experimental situations. For experiments (Fig. 2 open circles, Fig. 4, and Fig. 7) in which response amplitude ($\bar{R}$) was measured at different contrasts (c), we can use

$$\bar{R} = \bar{R}_{max} \frac{c^2}{\sigma^2 + c^2} \tag{8}$$

where $\bar{R}_{max}$ determines the maximal attainable response amplitude, and σ is the semisaturation contrast (the contrast at which response is half the maximum). This equation expresses a sigmoidal relationship between response and contrast; it is expansive at low contrasts (where $c \ll \sigma$) and saturating at high contrasts (where $c \gg \sigma$). As long as σ is positive, the response $\bar{R}$ will always be a value between 0 and $\bar{R}_{max}$, saturating for high contrasts.

Albrecht and Hamilton (1982) used a more general form of eqn. (8) as an *empirical* description of the relationship between response and contrast:

$$\bar{R} = \bar{R}_{max} \frac{c^n}{\sigma^n + c^n} \tag{9}$$

where n is a variable exponent. In a mechanistic description of the synaptic processes underlying contrast normalization, Carandini et al. (1997) have developed the related formulation:

$$\bar{R} = \bar{R}_{max} \frac{c^n}{\left(\sqrt{\sigma^2 + c^2}\right)^n} \tag{10}$$

The dividing signal depends upon the square of the contrast since it arises from many neurons, and the value of n averaged across the whole population of simple cells is close to 2 (Albrecht & Hamilton, 1982). The normalized underlying response of the neuron is then subject to an expansive output nonlinearity with an overall exponent n.

We used all three of these formulations to fit the response *versus* contrast data described in this paper. In most cases, the variable exponent n turned out to be close to 2 so that, in fact, the three formulations tended to be almost identical. The fitting procedure searched for the best-fitting values of the parameters: $\bar{R}_{max}$, σ, and n.

For the response *versus* spatial phase (ϕ) data of Fig. 5, we used

$$\bar{R} = K \cdot |\cos^n(\phi - \Phi)| \tag{11}$$

and the fitting procedure sought the best values of K, n, and Φ. Since the contrast (c) was fixed and presumably the semisaturation parameter (σ) was also constant throughout the experiment, the full normalization model was not needed for fitting this data set. Rather, the exponent (e.g. squaring) was the only part of the output nonlinearity that we needed to consider. We either allowed n to vary, or we fixed it at a value of 2. The fitting procedure accounted for 180-deg response phase shifts, as described above [eqn. (6)].

The fits to the polar plots of Fig. 6 were only slightly more complex, requiring a total of four parameters: K_1 (response to best spatial phase), K_2 (response to worst spatial phase), n (exponent), and Φ (best spatial phase).

For the directional suppression experiments (Figs. 8–10), it so happened for these neurons that the grating moving in the nonpreferred direction elicited no overt responses. But, even so, superimposing the grating moving in the nonpreferred direction would be expected to change the normalizing signal in the denominator of eqn. (10), which was elaborated to:

$$\bar{R} = \bar{R}_{\max} \frac{c_a^n}{\left(\sqrt{\sigma^2 + c_a^2 + c_s^2}\right)^n} \tag{12}$$

where c_a is the contrast of the activating grating and c_s is the contrast of the suppressing grating. One or the other of these contrasts was fixed in any given experiment, leaving three parameters ($\bar{R}_{\max}$, σ, and n) that were sought by the fitting procedure.

In the experiment investigating the interaction between two bar stimuli (Fig. 2, filled symbols), both bars *do* activate the neuron. Thus, we elaborated both the numerator and the denominator of eqn. (10):

$$\bar{R} = \bar{R}_{\max} \frac{(c_b + k_d \cdot c_d)^n}{\left(\sqrt{\sigma^2 + c_b^2 + c_d^2}\right)^n} \tag{13}$$

where c_b is the contrast of the first bar, c_d is the contrast of the second, and k_d is the ratio of the neuron's responsiveness to the two bar stimuli. This expression has four variable parameters ($\bar{R}_{\max}$, k_d, σ, and n) which were optimized by the fitting procedure.

Fitting procedures

Most of the experimental data were fit to the threshold model and normalization model using the simplex routine AMOEBA taken from Press et al. (1986). The simplex routine was used to minimize the following weighted least-squares error or merit function:

$$\chi^2 = \sum_i \frac{(m_i(\mathbf{p}) - r_i)^2}{w_i} \tag{14}$$

where r_i is the measured response to the ith stimulus condition and $m_i(\mathbf{p})$ is the model's prediction of the response with model parameters, $\mathbf{p}$. The w_i in the denominator is the weight to be given to the ith value; it should be equal to the variance of the response for that stimulus condition. Response variance was not usually measured in the original experiments, but we estimated it by knowing that the response variance of cortical neurons is directly proportional to response amplitude, independently of what precise stimulus configuration evoked that response (Tolhurst et al., 1981; Dean, 1981*b*; Tolhurst et al., 1983; Skottun et al., 1987; Vogels et al., 1989; Geisler & Albrecht, 1995). We had to allow a minimal nonzero variance for very small or zero response values:

$$w_i = r_i + \epsilon \tag{15}$$

where the constant ϵ was set arbitrarily to be 2.5% of the largest response (r) recorded in that experiment. It should be noted that, since these were weighted fits, they were dominated by the data with smallest variance (i.e. the smallest response amplitudes). In some cases (e.g. Fig. 5A), the best weighted fit might not appear, at first sight, to be correct: the fit to the high amplitude points has, in fact, been sacrificed in order to provide a good fit at low amplitudes.

When fitting the normalization model, the merit function [eqn. (14)] usually had a shallow minimum as a function of the expansive exponent n, which could be changed considerably on either side of the best-fit value with only small changes in χ^2.

We quote values for χ^2 to allow *comparison* of the best fits of different models for one data set at a time. The actual values of χ^2 should be interpreted cautiously because estimates, not measurements, of the response variances are used. Furthermore, the actual values are influenced by such arbitrary decisions as whether response should be expressed as impulses per second or as impulses per stimulus trial.

Fitting the line-weighting data

The fits to the line-weighting data (Fig. 3) were the most indirect. Tadmor and Tolhurst (1989) discuss the procedure in detail for the threshold model. The neuron's measured spatial-frequency response function was inverse Fourier transformed. The even-symmetric transform with amplitude A_E was added to the odd-symmetric transform with amplitude A_O. The resulting estimate of the underlying linear line-weighting function was shifted in position by an amount x_0, and was subjected to a threshold T [cf. eqn. (4b)]. The fitting procedure sought the best-fitting values of these four parameters (A_E, A_O, x_0, and T).

To fit the line-weighting data with the half-squaring contrast-normalization model, we needed only three parameters (A_E, A_O, and x_0). We assumed that the neuron's measured spatial-frequency tuning curve had itself been subject to half-squaring. Hence, we computed the square-root of the spatial-frequency tuning curve, before taking its inverse Fourier transform. The resulting even- and odd-symmetric inverse Fourier transforms were added with weights A_E and A_O. Finally, this function was squared and shifted in position by an amount x_0, before it was compared with the measured line-weighting data. We also examined a more general expansive nonlinearity, by taking the nth root of the spatial-frequency tuning curve before calculating the Fourier transforms, and then raising the model line-weighting function to the nth power before fitting it to the experimental data.

The fitting algorithm used for the line-weighting data was different from that used for the other data. We used the MRQMIN routine of Press et al. (1986) to optimize the parameters of the model. Furthermore, we minimized the *unweighted* sum of squared deviations between the data and the model predictions, for reasons that are discussed by Tadmor and Tolhurst (1989).

Results

Additivity of responses to paired bar stimuli

Fig. 1 illustrates an experiment first described by Tolhurst and Dean (1987) to test the additivity of the responses of simple cells to two simple stimuli. The PSTHs show the averaged responses of a simple cell to sinusoidal temporal modulation of bar stimuli. In Fig. 1A, a single bar was placed in the neuron's dominant excitatory region, and the temporal phase of modulation was such that the bar began by getting brighter than the uniform background (see schematic). The response waveform looks like a halfwave-rectified sine wave, which seems compatible with quasilinear behavior in a neuron with zero spontaneous activity (Movshon et al., 1978*a*). During the course of the experiment, the neuron gave an overall response ($\bar{R}_b$) of 239 action potentials to this stimulus. In Fig. 1B, we see the response ($\bar{R}_d$) when a bar was presented in the neuron's dominant inhibitory region; the temporal modulation applied to the bar was 180 deg out of phase with respect to Fig. 1A so that now

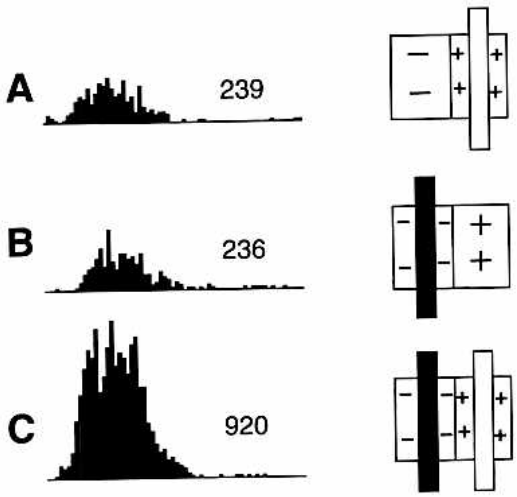

Fig. 1. The PSTHs show the failure of additivity of the responses to two excitatory stimuli. The binwidth is 16 ms, and each PSTH represents the sum of the responses to 200 stimulus cycles. The contrast of each stimulus was sinusoidally modulated at 1 Hz. The numbers beside each histogram show the total number of action potentials evoked. A: A bright bar with contrast 0.13 was presented in the dominant excitatory region of the receptive field (see schematic). B: A dark bar of contrast 0.075 was presented in the dominant inhibitory region. C: Both bars are presented simultaneously. Their contrasts were modulated in antiphase so that the temporal phases of the responses were almost the same. See Tolhurst and Dean (1987, their Fig. 2B) for more details.

the stimulus began by getting darker than the background. The response to this dark bar in the inhibitory region (236 action potentials) had nearly the same temporal phase as that to the bright bar in the excitatory region.

Fig. 1C shows the response ($\bar{R}_{b+d}$) when the two bars were presented simultaneously. If summation in the simple cell were strictly linear, we would expect

$$\bar{R}_{b+d} = \bar{R}_b + \bar{R}_d \tag{16}$$

This is a test of the Law of Superposition, and clearly it has not been obeyed: the response to the combined stimulus (920 action potentials) was much larger than predicted (475 action potentials). If the failure of the linear prediction were due only to an output nonlinearity, then in terms of eqn. (2):

$$f^{-1}(\bar{R}_{b+d}) = f^{-1}(\bar{R}_b) + f^{-1}(\bar{R}_d) \tag{16a}$$

Tolhurst and Dean (1987) were able to rationalize the discrepancy using a threshold model. Suppose that the *underlying* responses (R_b and R_d) of the neuron *are* summed linearly, but that some threshold response value (T) must be exceeded before the overt responses ($\bar{R}_{b+d}$, $\bar{R}_b$, and $\bar{R}_d$) are measurable. Then we would expect

$$(\bar{R}_{b+d} + T) = (\bar{R}_b + T) + (\bar{R}_d + T) \tag{17}$$

where $\bar{R}_b = R_b - T$ is the overt response to the bright bar, $\bar{R}_d = R_d - T$ is the overt response to the dark bar, $R_b = K \cdot c$ is the underlying linear response to the bright bar, and c is the contrast of the bright bar [cf. Methods, eqn. (3)]. The behavior of the neuron in Fig. 1 would be compatible with linear summation, if it had a response threshold of about 450 action potentials.

On the other hand, Heeger (1992*b*) has shown that this kind of mismatch between simple linear prediction and experimental measurement can also be explained on the half-squaring hypothesis. Here, it is presumed that the overt response ($\bar{R}$) is the square of the underlying linear response (R). Then, we would expect

$$\sqrt{\bar{R}_{b+d}} = \sqrt{\bar{R}_b} + \sqrt{\bar{R}_d} \tag{18}$$

Application of this relation to the data of Fig. 1 gives an almost perfect match, as it does for the other neuron illustrated by Tolhurst and Dean (1987, their Fig. 2A). In fact, the full contrast-normalization model would predict a more complicated relationship, since the normalizing signal changes when a second bar stimulus is added [see eqn. (13)]. In this example, however, the effect of normalization will be small because the contrast energy of the bar stimuli was low compared to the best-fitting semisaturation parameter ($\sigma = 0.32$, see Fig. 2B).

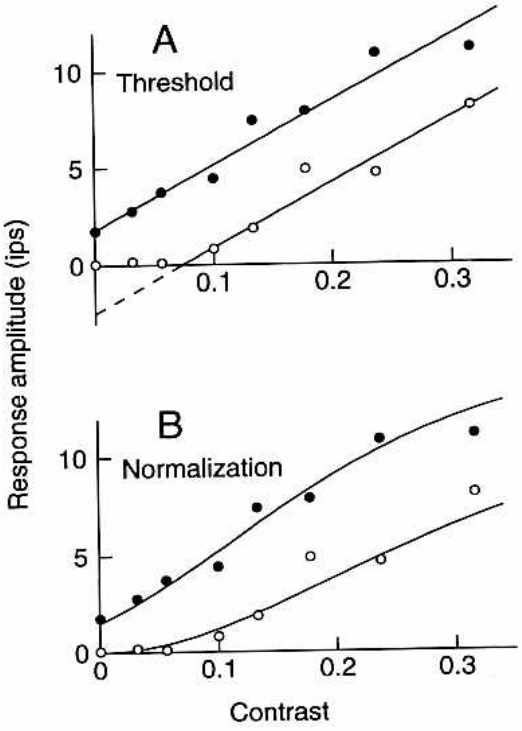

Fig. 2. A detailed examination of the additivity of excitatory responses for the same experiment as Fig. 1. The open circles in (A) and (B) show the amplitude of the response to the bright bar in the neuron's dominant excitatory region as a function of stimulus contrast (c_b). The filled symbols show the effect of adding a dark bar of fixed contrast 0.075 (c_d) in the dominant inhibitory region. The lines through these data in (A) are parallel and are the best joint fit of the threshold model to the two sets of data [eqn. (4) is fit to the open circles, and eqn. (5) is fit to the filled circles; three parameters: $K = 33.03$, $R_d = 4.25$, $T = 29.9\%$ of response to the highest contrast bright bar; $\chi^2 = 1.21$]. The dashed line shows the supposed subthreshold responses to single bars. The curves in (B) are the best joint fit of the normalization model to both data sets [eqn. (10) fit to the open circles, and eqn. (13) fit to the filled circles; four parameters: $\bar{R}_{max} = 14.09$, $\sigma = 0.32$, $k_d = 1.47$, $n = 2.05$; $\chi^2 = 1.64$]. The threshold model needed one fewer fitting parameters than the normalization model since there is little sign of response saturation at high contrasts; $\bar{R}_{max}$ can be ignored for the threshold fit. Data replotted from Tolhurst and Dean (1987, their Fig. 3B). They fitted regression lines to the "linear" portions of the two data sets: open circles—slope = 32.5, S.E. = 5.38; filled circles—slope = 38.2, S.E. = 2.90.

Tolhurst and Dean (1987) performed a more elaborate experiment with the intention of establishing that the threshold constant T remained the same for a variety of different stimulus conditions. Fig. 2 illustrates such an experiment, from which the examples of Fig. 1 were taken. The data are the same in the two parts of the figure; the models fitted are different. The open circles show how the response amplitude depended on the contrast of a bar presented in the neuron's dominant excitatory region. The filled symbols in Fig. 2 show how the response *versus* contrast relationship for the bar in the excitatory region was changed by the simultaneous presence of a second bar in the neighboring inhibitory region of the receptive field. The second bar was of fixed contrast (0.075) and was modulated in antiphase (see Fig. 1B) so that the responses to the two bar stimuli added in the same temporal phase.

The threshold model (Fig. 2A) predicts a parallel upwards shift by a fixed value, equal to the underlying linear response to the dark bar when presented alone ($R_d = \bar{R}_d + T$). The lines drawn through these data in Fig. 2A are the best-fitting threshold model applied jointly to the open [eqn. (4)] and filled [eqn. (5)] symbols, with shared parameter T (there are three fitting parameters altogether). It fits reasonably well.

The prediction of the contrast-normalization model (Fig. 2B) is more complicated, and needs four fitting parameters [see Methods, eqn. (13)]. The curves drawn through the data are solutions of the model [eqns. (10) and (13)] with the best overall fit to the two data sets of Fig. 2B. The parameters n and σ are shared between the two data sets. It is important to note that this model predicts that the relationship for paired bars will *not* be simply parallel to the control for single bars (cf. the threshold model). Rather, the two curves should diverge slowly. Such divergence is evident in the data of Fig. 2, and was indeed reported by Tolhurst and Dean (1987) for two out of their three experiments (their Figs. 3 and 6).

Overall the two models describe the data almost equally well, with the threshold model being slightly the better ($\chi^2 = 1.21$ for threshold model, three parameters; $\chi^2 = 1.64$ for normalization model with $n = 2.05$, four parameters). The threshold model is superior largely because it is the better fit to the very low-amplitude responses to low-contrast bright bars when presented alone (open circles). For instance, a bar of contrast 0.056 evoked a response of only 85 action potentials in 200 repetitions of the stimulus (an average of one action potential for every three stimulus presentations). Had the neuron produced an average of one action potential for each stimulus repetition, then the normalization model would have been the better fit to the data of Fig. 2. Moreover, in another very similar experiment on this neuron (unpublished), the normalization model provided the slightly better fit.

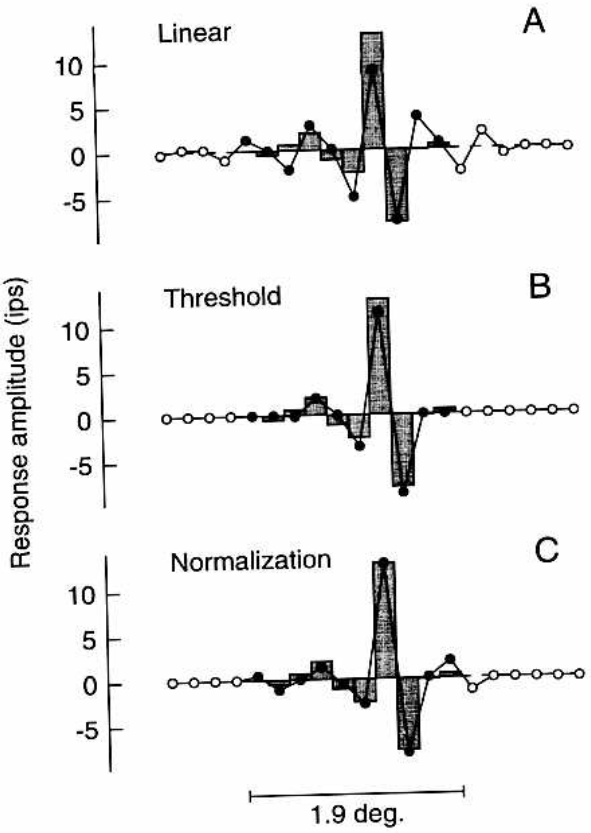

Fig. 3. The histogram shows the measured receptive-field profile of a simple cell, as measured with bars of the same width as the histogram blocks. The receptive field was mapped with bars 0.19 deg wide and a contrast of 0.56 over a total extent of 1.9 deg. Preliminary experiments had shown that there was no significant activity outside of this extent. The circles show the best-fitting predictions of the profile under a number of models. The filled circles show the prediction within the spatial extent that the field was actually mapped; the open circles show predictions outside the mapped region. A: A simple linear prediction from the inverse Fourier transform of the spatial frequency tuning curve (0% threshold model). The residual sum of squares is 10.48, with three fitting parameters (A_E, A_O, x_0); for these data, we did not use a weighted fit (see Methods). B: A threshold, T, equal to 35% of the maximal response recorded in the experiment is incorporated in the model (residual sum of squares = 1.34, with four parameters). C: The model assumes that the measured responses to both sinusoidal gratings and to bars have been subject to half-squaring (residual sum of squares = 0.794, with three parameters, one fewer than the threshold model since n was fixed at 2.0). A similar fit with $n = 2.3$ gave a residual sum of squares of 0.735. Further details in Tadmor and Tolhurst (1989, their Fig. 4).

Responses to single bars: The line-weighting function

There have been many studies in which the receptive-field map has been compared with the inverse Fourier transform of the spatial-frequency tuning curve (see Introduction for references). Fig. 3A shows the measured receptive-field profile of a simple cell (histogram bars) along with the best-fitting linear prediction from the spatial-frequency tuning (reproduced from Tadmor & Tolhurst, 1989, their Fig. 4). Details of the least-squares fitting procedure are given in the Methods. As in most such studies, the inverse Fourier transform predicts that there should have been more excitatory and inhibitory subregions within the receptive field than were actually measured experimentally with bars or spots. Several authors have proposed that this discrepancy between measured and predicted receptive-field profile can be attributed to a threshold in the neuron's responses (Andrews & Pollen, 1979; Kulikowski & Bishop, 1981; Glezer et al., 1982; R. De Valois et al., 1985; DeAngelis et al., 1993).

Tadmor and Tolhurst (1989) addressed this question quantitatively, and Fig. 3B shows the best-fitting prediction from the inverse Fourier transform, once a threshold has been allowed. Incorporation of the threshold causes a great improvement in the fit (residual sum of squares reduced by a factor of about 8). However, very high threshold values were often needed for these fits, and inspection of Fig. 3B suggests that the model is now underestimating the number and strengths of the subsidiary regions in the receptive field.

Heeger (1992*b*) and DeAngelis et al. (1993) have shown that the discrepancies between measurement and linear prediction (as

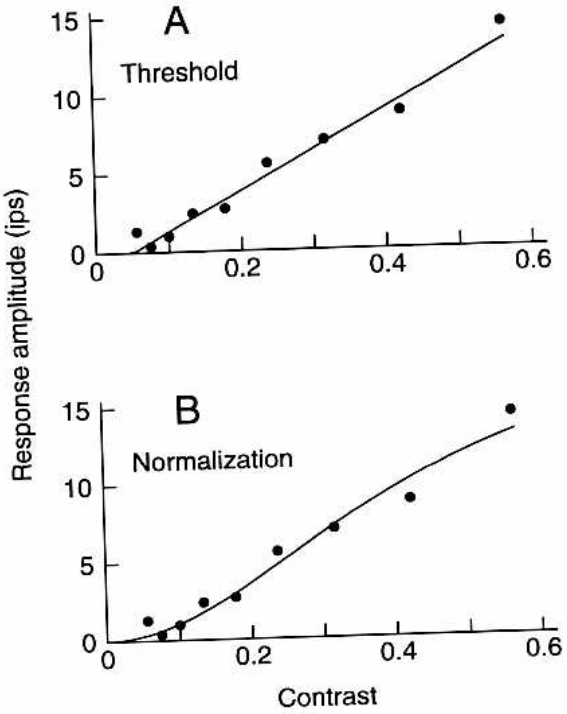

Fig. 4. For the same neuron as in Fig. 3, the amplitude of response is plotted against the contrast of a bar of width 0.19 deg, presented in the neuron's most sensitive excitatory region. Further details in Tadmor and Tolhurst (1989, their Fig. 7). The line in (A) is the best fit of the threshold model [eqn. (2)]; χ^2 was 1.72 (two parameters since $\bar{R}_{max}$ was ignored) and the threshold parameter T was only about 9.5% (cf. the value of 35% needed to fit the receptive-field profile in Fig. 3B). The curve in (B) is the best fit of the normalization model [eqn. (10); three parameters: $\bar{R}_{max}$ = 25.17, n = 1.79, σ = 0.56; χ^2 = 1.45].

in Fig. 3A) are also compatible with the contrast-normalization model. Fig. 3C shows the same receptive-field profile, along with the best-fitting prediction of the normalization model using an expansive exponent of 2. The residual sum-of-squares is less than that for the threshold model (see figure legend). Inspection of Fig. 3C shows that the model does describe the weak flanking regions of the receptive field rather well. We also examined the effect of changing the expansive exponent in the normalization model (see Methods); the data were best fit with n of 2.3.

Tadmor and Tolhurst (1989, their Fig. 7) also analyzed response *versus* contrast measurements from the same neuron. We replot these data (Fig. 4) and have fit them with the predictions of the threshold model [eqn. (4)] and the normalization model [eqn. (8)]. The χ^2 of the best fits is almost the same in the two cases. However, the magnitude of the threshold parameter from this fit (Fig. 4A) was only about 9.5%, compared to the value of 35% needed to fit the line-weighting data in Fig. 3B. There was much greater consistency with the normalization model: the best-fit exponent for the line-weighting data was 2.3, while the best-fit exponent for the response *versus* contrast measurements of Fig. 4B was 1.79. These values are very close, considering that the minima in the fitting procedure were quite shallow (see Methods).

Response amplitude at different spatial phases

According to linear theory, response amplitudes evoked by stationary sinusoidal gratings would vary sinusoidally with the grating's spatial phase (e.g. Enroth-Cugell & Robson, 1966; Movshon et al., 1978*a*). Fig. 5 shows the results of such an experiment on a simple cell, replotted from Robson et al. (1988). As with other predictions of the simple linear model, this one is not fulfilled in detail. Robson et al. (1988) supposed that deviation from a perfect sinusoid could be explained if the neuron had a threshold. The curves drawn through the data of Fig. 5A show the best-fitting prediction of the threshold model [see eqn. (6), Methods for details]. This fit may not, at first sight, seem convincing but it must be remembered that this is a *weighted fit* (see Methods) which will tend to fit the low-amplitude data at the expense of the high-amplitude data. Fig. 5B shows the same data but fitted with the contrast-normalization model [eqn. (11), n = 2.35].

The normalization model (χ^2 = 0.81, three parameters) fits better than the threshold model (χ^2 = 2.24, also three parameters). However, inspection of Fig. 5 suggests that the difference in performance of the two models may hinge almost entirely on their

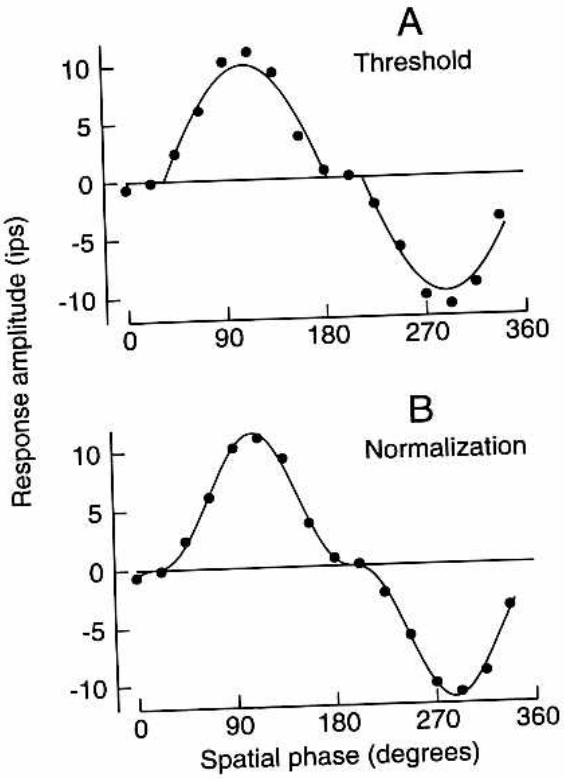

Fig. 5. The amplitude of response of a simple cell is plotted against the spatial phase of a stationary sinusoidal grating of optimal orientation and spatial frequency. The contrast of the grating was modulated sinusoidally in time at 2 Hz. Negative responses represent a phase shift of 180 deg in the temporal phase of response. Gratings were presented only over a spatial phase range of 180 deg, and the data have been negated and repeated to complete the full cycle. The Michelson contrast was 0.31. A: The data are compared with the best-fitting sinusoid after applying a threshold [eqn. (6a); T = 38% of the maximal response evoked in the experiment; χ^2 is 2.24, three parameters]. The relationship between response and grating contrast was not determined for this neuron, so we do not have an independent estimate of the threshold value. B: The curve is the best-fitting solution of the contrast-normalization model [eqn. (11); K = 11.46, n = 2.35; χ^2 is 0.81, three parameters including Φ]. The data were also fit with an exact half-squaring normalization model (not shown; K = 10.69, χ^2 = 1.53, two parameters), and with a strictly linear model, i.e. a pure sinusoid (not shown; χ^2 = 20.35, two parameters). Further details in Robson et al. (1988, their Fig. 2A).

respective abilities to fit just one of the data points. The response at a phase of 0 deg/180 deg had a very low amplitude (only 25 action potentials were recorded in 100 repetitions of the stimulus). Had the response actually been eight action potentials instead of 25, then the threshold model would have been slightly the better fit.

Polar plots of response to modulated gratings, and directional selectivity

In most simple cells, the temporal phase of the response also changes with the spatial phase of a modulated sinusoidal grating (Movshon et al., 1978*a*). According to the linear model, a polar plot of response amplitude as a function of *response* phase should produce an ellipse centered on the origin. The degree of directional selectivity shown by the neuron will affect the aspect ratio of the ellipse. For a neuron with no directional selectivity, the ellipse should degenerate into a straight line; for a neuron with total directional selectivity, the ellipse should become a perfect circle. For neurons with varying degrees of directional selectivity, the ellipses should have an aspect ratio in between these two extremes.

Fig. 6 shows an example of such data replotted from Tolhurst and Dean (1991, their Fig. 1A). The responses to moving gratings of the same spatial and temporal frequencies were measured in the same experiment. An ellipse is, in fact, a poor fit to the data (Fig. 6A) because the responses near the minor axis are smaller than expected. This "wasp-waist" was described by Movshon et al. (1978*a*) who proposed that it could be explained by the existence of a threshold. Fig. 6B shows the best-fitting version of the threshold model and there is, indeed, a considerable improvement in the fit (χ^2 has fallen by a factor 9, for the addition of one extra fitting parameter, T). However, the threshold parameter T had to be very high (74%) to achieve this fit.

The expansive nonlinearity in the normalization model would also predict a wasp-waist (Heeger, 1992*b*, 1993) and Fig. 6C shows the best fit of the model, with $n = 10.1$. Although this may seem to be a very high value, the minimum in the fitting routine was so shallow that the fit was only a little better than with n of 4.0.

We can try to obtain independent support for the estimates of the threshold or of the exponent, from the form of the relationship between response amplitude and contrast for moving gratings. Fig. 7 shows these results (previously unpublished), fit with the threshold model (A) and the normalization model (B). The two fits are almost equally good. However, there is a major inconsistency with the threshold model: the best-fit threshold parameter in Fig. 6B was 74%, while in the experiment of Fig. 7A it was only 1.9%. The normalization model fares little better: the best fit to the data of Fig. 7B required an exponent n of 1.30, which is very different from the value of 10.1 needed to fit the wasp-waist of Fig. 6C. This discrepancy is illustrated by the solid curve in Fig. 6D which shows the best-fitting solution of the expansive nonlinearity in the normalization model with $n = 1.3$; the fit hardly shows a wasp-waist at all.

We made an overall evaluation of the performance of each model by fitting the ellipse (Fig. 6) and contrast experiment (Fig. 7) jointly. For the threshold model, the procedure found the single value of the threshold parameter T that best fit *both* sets of data. The normalization model sought the best-fitting value of the exponent n to fit both sets of data. The data points from the two experiments were weighted according to their actual response values (see Methods). This meant that the best fits tended to be those that fit the low-amplitude responses to low stimulus contrasts in Fig. 7 well, at the expense of ignoring the wasp-waist of Fig. 6. Overall, the threshold model needed a threshold T of 12% of the response to a grating in the best spatial phase ($\chi^2 = 17.98$, six parameters altogether). The best-fitting exponent in the normalization model was 2.29 ($\chi^2 = 11.08$, six parameters altogether). The normalization model is the better fit overall, but both models show inconsistencies; we shall return to this point in the Discussion. Fig. 6D (dashed curve) shows the best-fit solution of the normalization model with this overall n of 2.29.

Predictions of directional index and response amplitude

We can use the same data (Figs. 6 and 7) to evaluate the two models in another way. For a linear neuron, the responses to moving gratings should be predictable from the responses to modulated gratings. In addition, the directional index, a measure of the neuron's directional selectivity, should be the same when computed from the responses to moving gratings and to modulated gratings. However, these linear predictions are only partially successful (Reid et al., 1987, 1991; Albrecht & Geisler, 1991; Tolhurst & Dean, 1991; DeAngelis et al., 1993).

Tables 1 and 2 examine these predictions and the effects of various models of the output nonlinearity, for the same neuron used for Figs. 6 and 7. Table 1A lists the amplitudes of the responses to moving and modulated gratings all of the same spatial and temporal frequency, and of the same contrast. Table 1B lists the directional indices calculated from the response amplitudes. First, the directional index was calculated, assuming the linear model, from the responses to moving gratings:

$$DI = \frac{R_p - R_n}{R_p + R_n} \tag{19}$$

where R_p is the preferred direction response and R_n is the nonpreferred direction response. The directional index was also calculated from the responses to modulated gratings:

$$DI = \frac{R_2}{R_1} \tag{20}$$

where R_1 is the response to the best spatial phase and R_2 is the response to the worst spatial phase.

The first row of Table 1B shows that the simple linear prediction from modulated gratings [eqn. (20)] is a significant underestimate compared to the calculation from moving gratings [eqn. (19)], as has been reported before. The remaining rows of the table repeat these calculations, supposing that the underlying responses do obey the relations of eqns. (19) and (20) but that the overt measured responses have been subject to some output nonlinearity. Using the terminology of eqn. (2), the predictions become

$$DI = \frac{f^{-1}(\bar{R}_p) - f^{-1}(\bar{R}_n)}{f^{-1}(\bar{R}_p) + f^{-1}(\bar{R}_n)} \tag{19a}$$

$$DI = \frac{f^{-1}(\bar{R}_2)}{f^{-1}(\bar{R}_1)} \tag{20a}$$

Accounting for either a high threshold (from the best fit to Fig. 6B) or a low threshold (from overall fit to Figs. 6 and 7) does not improve the agreement between the two estimates of directional index. An expansive nonlinearity with exponent 10.1 (from best fit to Fig. 6D) is also unimpressive. However, if we presume that the responses have been subject to expansion with exponent 2.29 (overall fit to Figs. 6 and 7), then the two estimates of direc-

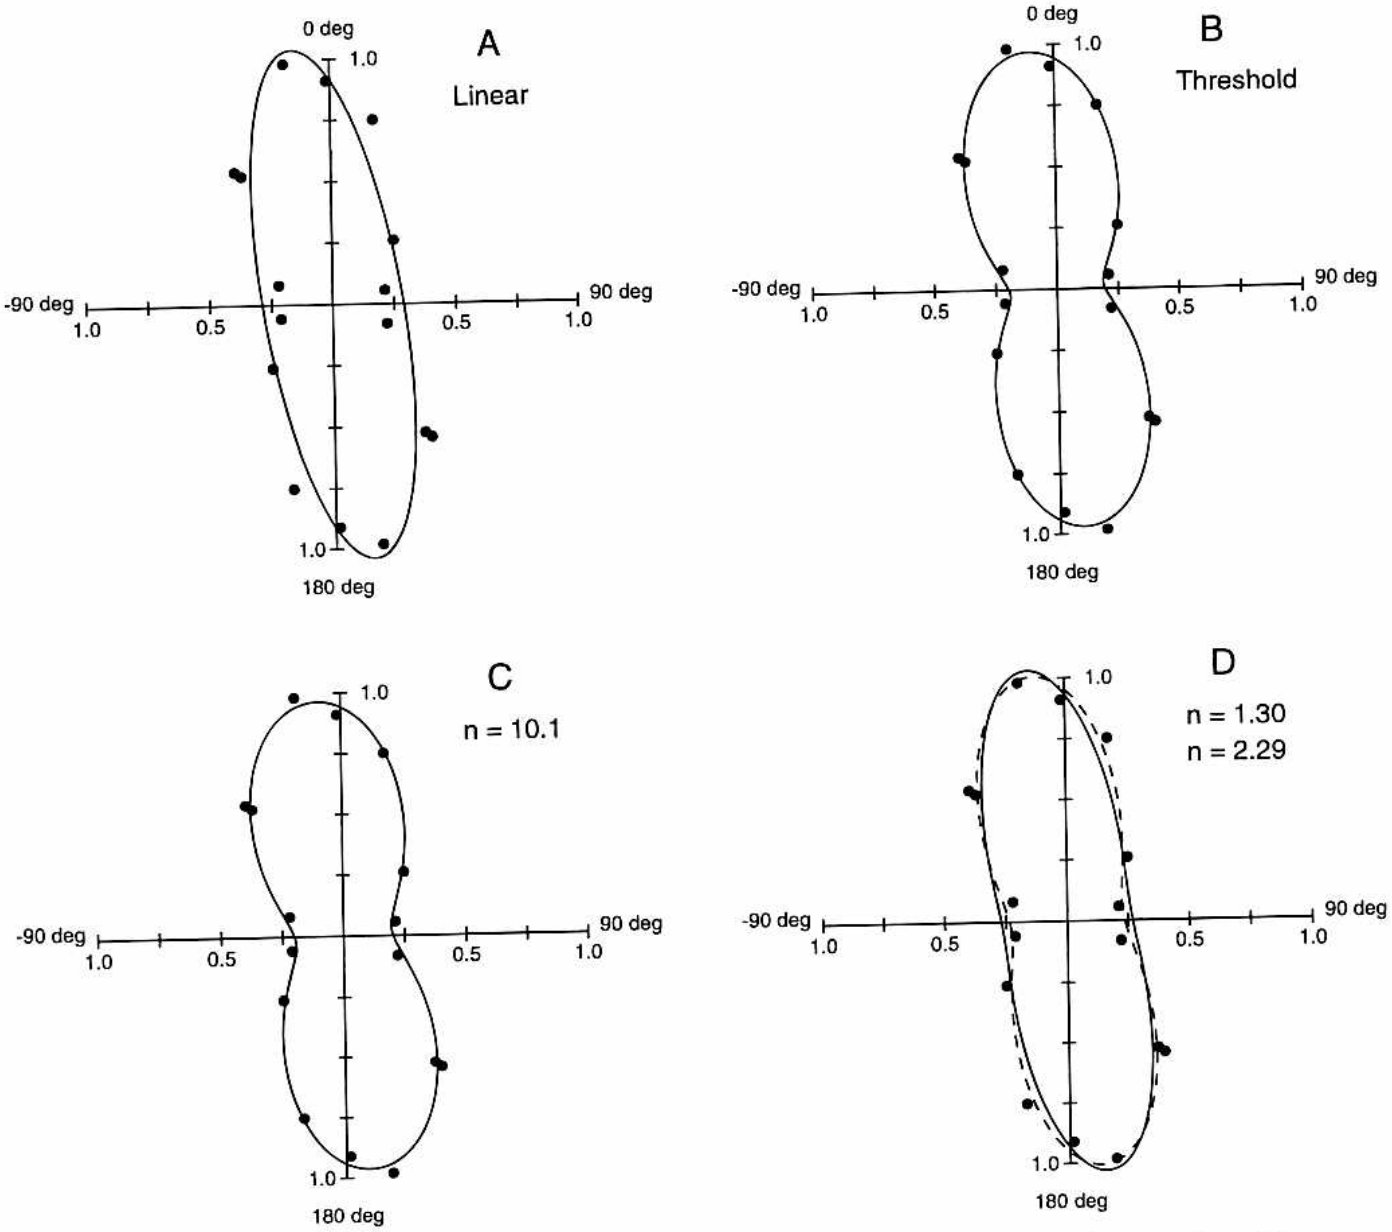

Fig. 6. A polar plot of response amplitude and response phase for an experiment in which the spatial phase of a temporally modulated sinusoidal grating was varied systematically. This is a different neuron from Fig. 5. The grating's contrast was modulated sinusoidally at 2 Hz, from a maximal Michelson contrast of 0.31. In the same experiment, the responses were measured for gratings of the same spatial frequency and contrast moving in the preferred and nonpreferred directions through the receptive field. They, too, had a temporal frequency of 2 Hz. Details in Tolhurst and Dean (1991, their Fig. 1A). The four panels show the best-weighted fits to four different models. A: A simple linear model predicts an ellipse ($\chi^2 = 11.51$, three parameters: K_1, K_2, Φ). B: A threshold model with threshold T equal to 74% of the overt response at the best spatial phase ($\chi^2 = 1.29$, four parameters). C: Contrast normalization with exponent n as a parameter ($n = 10.1$; $\chi^2 = 1.12$, four parameters). D: Contrast normalization with n fixed at 1.3 (solid curve; $\chi^2 = 7.36$, three parameters) or at 2.29 (dashed curve; $\chi^2 = 2.81$, three parameters). The axes show response in arbitrary units; the firing rate in the best spatial phase was 53 ips.

tional index are closer. The discrepancy almost disappears when the exponent is 1.30 (the best fit to Fig. 7B).

Linear theory (Reid et al., 1987, 1991) also shows that the amplitudes of response to moving gratings should be predictable from those to modulated gratings:

$$R_p = R_1 + R_2 \tag{21}$$

$$R_n = R_1 - R_2 \tag{22}$$

The linear model generally fails; in particular, the prediction usually overestimates the amplitude of the response in the nonpreferred direction (Reid et al., 1987, 1991; Tolhurst & Dean, 1991).

Table 2 shows the success of different models in predicting the responses to moving gratings from those to modulated gratings. The first row of the table shows the simple linear model making the familiar failures of prediction. The remaining rows show the predictions, supposing various output nonlinearities. For instance,

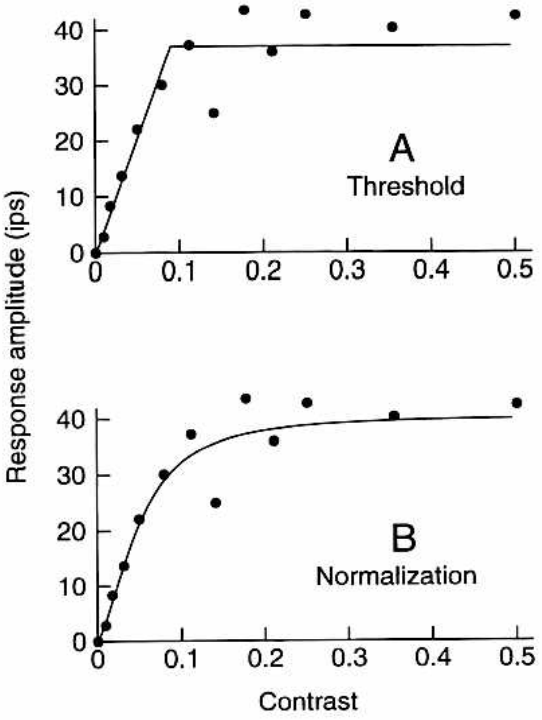

Fig. 7. For the neuron of Fig. 6, the amplitude of response for moving gratings is plotted against stimulus contrast. A: The lines show the best-fitting version of the threshold model [eqn. (4); $\chi^2 = 9.26$, three parameters]; the threshold parameter T is only 1.9%, very different from that estimated from the polar plot (Fig. 6B). B: The curve shows the best fit of the contrast-normalization model [eqn. (10); $\bar{R}_{max} = 40.45$; $\sigma = 0.065$; $n = 1.30$; $\chi^2 = 6.79$, three parameters].

the general solution for responses in the preferred direction is, in terms of eqns. (1) and (2),

$$\bar{R}_p = f(f^{-1}(\bar{R}_1) + f^{-1}(\bar{R}_2)) \tag{21a}$$

Attempting to account for a threshold (whether high or low) makes the discrepancies between prediction and measurement even worse. However, the contrast-normalization model expects these discrepancies. This is partly because of the expansiveness (e.g. half-squaring) in the neuron's output, but there is also an important new consideration: there is a difference in the normalization signals that arise from moving and modulated gratings (Heeger, 1993; Tolhurst & Heeger, 1997). The time-averaged contrast energy of the modulated gratings is less than that of the moving gratings, and so modulated gratings will invoke less normalization for a given contrast. Hence, the proper predictions of response in the preferred and nonpreferred directions become

$$\bar{R}_p = \left(\sqrt[n]{\bar{R}_1} + \sqrt[n]{\bar{R}_1}\right)^n \cdot s \tag{21b}$$

$$\bar{R}_n = \left(\sqrt[n]{\bar{R}_1} - \sqrt[n]{\bar{R}_1}\right)^n \cdot s \tag{22b}$$

where n is the expansive exponent and s is a contrast-dependent scaling factor that accounts for the lower normalizing effect of modulated gratings (Tolhurst & Heeger, 1997). From Tolhurst and Heeger [1997, their eqn. (21)],

$$s = \left(\frac{\sqrt{\sigma^2 + (c^2/2)}}{\sqrt{\sigma^2 + c^2}}\right)^n \tag{23}$$

The models with contrast normalization generally fare better than the threshold model, although the exponent of 10.1 is much too extreme in its "corrections." The amplitude of response in the nonpreferred direction is best predicted by an exponent between 1.30 (from the best fit to Fig. 7B) and 2.29 (from the overall fit to Figs. 6 and 7). The response in the preferred direction is fit most convincingly by the normalization model with $n = 1.30$.

Table 1A. *Definition of terms (Heeger, 1993; Tolhurst & Heeger, 1997) and summary of response amplitudes, expressed relative to the response to modulated gratings of the best spatial phase for the experiment illustrated in Fig. 6*

Grating stimulus	Abbreviation	Measured response amplitude
Modulated in best spatial phase	$\bar{R}_1$	1.000
Modulated in worst spatial phase	$\bar{R}_2$	0.212
Moving in preferred direction	$\bar{R}_p$	0.884
Moving in nonpreferred direction	$\bar{R}_n$	0.298

Table 1B. *Calculations of directional index of the supposedly linear underlying responses of the neuron*[a]

Model	Directional index calculated from: Moving	Modulated
Simple linear model	0.496	0.212
Threshold model (T = 74%, from Fig. 6B)	0.220	0.547
Threshold model (T = 12%, overall fit)	0.412	0.296
Exponent, n = 10.1 (Fig. 6D)	0.054	0.858
Exponent, n = 2.29 (overall fit)	0.234	0.508
Half-squaring	0.265	0.461
Exponent, n = 1.30 (Fig. 7B)	0.395	0.303

[a]For each model, the index is calculated from the supposedly linear underlying responses to moving gratings [eqns. (19) or (19a)] and from those to modulated gratings [eqns. (20) or (20a)], assuming different nonlinear output models. A model would be successful if the two calculations of directional index gave almost the same value (e.g. see last row).

Table 2. *Predictions of the response amplitudes for moving gratings from the responses to modulated gratings, for the same neuron as Figs. 6 and 7, and Table 1*[a]

Model	Response amplitude calculated from modulated responses: Preferred direction (cf. 0.884)	Nonpreferred direction (cf. 0.298)
Simple linear model	1.212	0.788
Threshold model (T = 74%, from Fig. 6B)	1.950	0.048
Threshold model (T = 12%, overall fit)	1.330	0.668
Exponent, n = 10.1 (Fig. 6D)	15.71	0.000
Exponent, n = 2.29 (overall fit)	1.176	0.091
Half-squaring	1.084	0.148
Exponent, n = 1.30 (Fig. 7B)	0.924	0.409

[a]The linear model shows the solutions to eqns. (21) and (22). The contrast normalization models show solutions to eqns. (21b) and (22b) with values of s from eqn. (23). The responses actually measured were 0.884 in the preferred direction and 0.298 in the nonpreferred direction. A successful model would need to match both these values (e.g. see last row).

Overall, the contrast-normalization model with an exponent of 1.3–2.0 has provided the best description of the various relationships in Tables 1 and 2. The threshold model has not performed well.

Response suppression by gratings moving in the nonpreferred direction

Dean et al. (1980) briefly reported a series of experiments in which they studied the suppressive effect of "masking" gratings moving in the simple cell's nonpreferred direction. At the time, the experiments were thought to be examining a directionally selective form of inhibition, responsible for the neuron's own directional selectivity. We now believe that this effect is just one example of a more general, nonspecific suppression (Robson, 1988; Bonds, 1989; Heeger, 1992*a*; and many others cited in the Introduction).

Fig. 8 shows the example originally described by Dean et al. (1980, their Fig. 1). The open symbols show the responses of a highly directionally selective simple cell as a function of the contrast of a single grating moving in the preferred direction of motion. The filled symbols show that these responses were reduced by the simultaneous presence of a second (suppressing) grating, which had a relatively high, fixed contrast (0.4) and which moved in the neuron's nonpreferred direction. The suppressing grating evoked almost no response when presented alone (filled circle on ordinate). The data do seem to support the proposition of Dean et al. (1980) that the suppressing stimulus has decreased the neuron's gain or responsivity by an arithmetic operation akin to division since the suppressing grating has caused a change in slope of the response *versus* contrast relation.

Dean et al. (1980) had no theory to explain why a suppressing contrast of 0.4 should cause the particular magnitude of effect that was observed, and they concentrated on the responses to low-contrast stimuli because saturation phenomena would have been difficult to explain. The real success of the normalization model is shown in this experiment since it naturally encompasses the data. The two curves drawn through the data in Fig. 8 are the *single* best-fit solution to the normalization model [eqn. (12)], with c_s of zero for the open circles and of 0.4 for the filled circles. Only three parameters are needed and these are shared between the two curves ($\bar{R}_{max}$, n, and σ). Furthermore, these are the same three parameters as are used to fit a single "control" response *versus* contrast relation [see eqn. (10)].

The forms of the suppression curves can be more complicated than is shown in Fig. 8, especially when higher activating contrasts are used. Figs. 9 and 10 show the results of several experiments performed on two other simple cells. For each neuron (Fig. 9A and Fig. 10A), the relation between response and activating contrast was studied in the presence of suppressing gratings of several different contrasts (c_s). These results are similar to that in Fig. 8, with the obvious extra finding that the amount of response suppression increased as the suppressing contrast was increased (Morrone et al., 1982; Bonds, 1989; Carandini et al., 1997). This is examined explicitly in Figs. 9B and 10B, which show how the response to one or two activating contrasts (c_a) was affected as a function of the contrast of the suppressing grating.

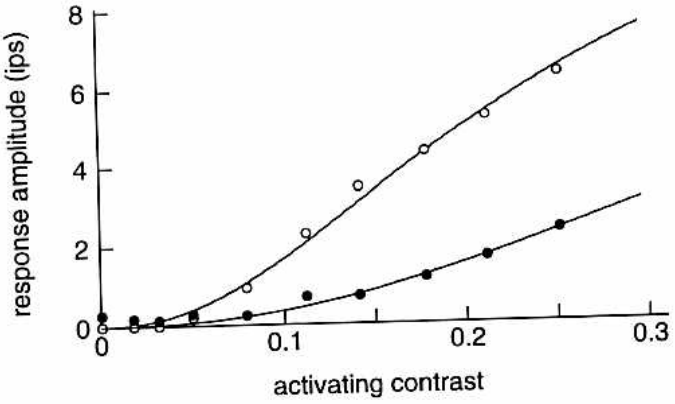

Fig. 8. The relation between response amplitude and contrast for a simple cell that was extremely direction selective. The open circles show the response to sinusoidal gratings moving in the preferred direction at 2 Hz. The filled circles show how the relation was changed by the simultaneous presence of a grating moving in the nonpreferred direction at 2.5 Hz with contrast of 0.4. Replotted from Dean et al. (1980). The curves drawn through the data show the best-fitting solution to the normalization model [eqn. (12): $c_s = 0$ for open circles, $c_s = 0.4$ for filled circles]; the three best-fit parameters, shared between the two graphs were: $\bar{R}_{max} = 12.30$, $n = 2.31$, $\sigma = 0.22$.

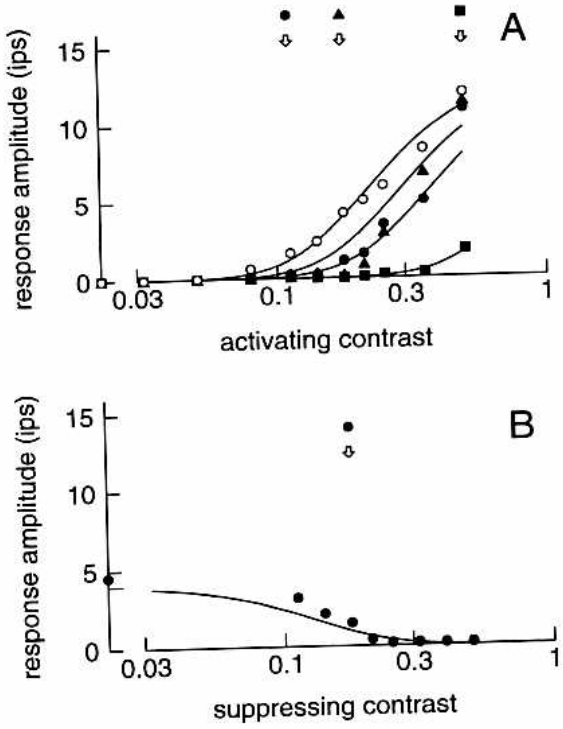

Fig. 9. A: The open circles show the relation between response amplitude and activating contrast for gratings moving in the preferred direction at 2 Hz. The filled symbols show how that relationship was changed by the presence of suppressing gratings moving in the nonpreferred direction at 2.5 Hz. The three suppressing contrasts were studied in three separate experiments, each of which included a full "control" curve; the open circles are the averages of the measurements in the three experiments. The suppressing contrasts were 0.11 (filled circles), 0.18 (filled triangles), and 0.5 (filled squares). The neuron had no spontaneous activity and the suppressing gratings evoked no response when presented alone (overlaid points plotted on the ordinate). B: For the same simple cell, the response to an activating grating of 2 Hz and contrast 0.18 is reduced by the simultaneous presence of suppressing gratings (2.5 Hz) at a variety of contrasts. The response to the activating grating alone is shown on the ordinate. The five curves drawn in this figure are the *single* best-fit solution of the normalization model [eqn. (12)] with only three parameters: $\bar{R}_{max} = 13.11$, $n = 6.44$, $\sigma = 0.12$. The parameters are shared by all five curves which differ in the experimental variables c_a and c_s.

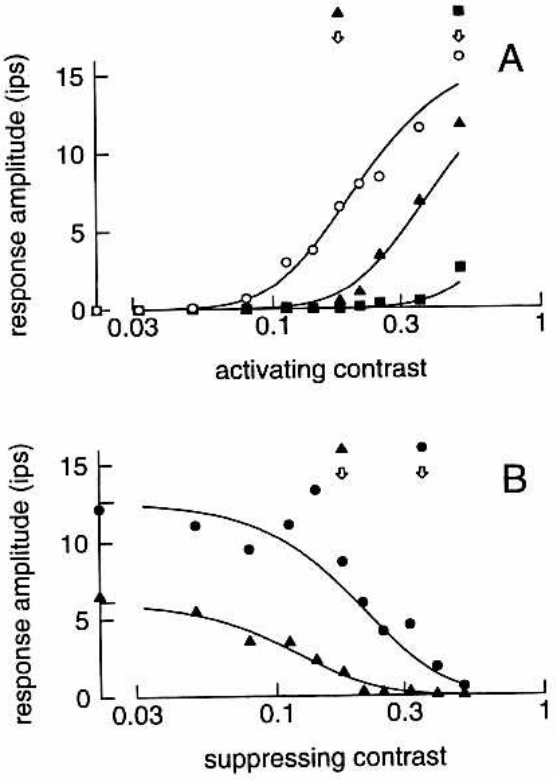

Fig. 10. Like Fig. 9 but for another simple cell recorded in the same animal. A: The effects of only two suppressing contrasts were studied, so that the open circles show the average of two measurements of response. B: The action of suppressing gratings was studied for two activating contrasts in the same experiment. The five curves in the figure are the *single* best-fit solution of the normalization model [eqn. (12)] with only three parameters: $\bar{R}_{max} = 16.87$, $n = 5.81$, $\sigma = 0.12$.

The curves drawn through the data in Figs. 9 and 10 are the best fits of the contrast-normalization model [eqn. (12)]. For each simple cell, the five data sets have been fit with a *single set of three parameters* ($\bar{R}_{max}$, n, and σ); the five curves differ only because they use the specific values of the experimental variables c_a and c_s appropriate to that data set. Obviously, the fits in Figs. 9 and 10 are not perfect and the model may need some refinement. In particular, the fits to the filled circles and filled triangles in Fig. 9A are poor, where a suppressing contrast of 0.11 often had a greater effect than a suppressing contrast of 0.18 (the results of Fig. 9B show a more consistent trend). We will consider in the Discussion whether the experimental design may have contributed to this apparent problem with the normalization model.

However, the model fits well overall, and it is particularly successful in describing the sigmoidal form of the graphs in Figs. 9B and 10B. The threshold model may be able to account for these data, but it is not exactly clear how.

Discussion

In this paper, we have reexamined a variety of experimental observations on cat simple cells, which all show some departure from the predictions of a linear theory. Originally, these departures had been explained on the hypothesis that the neuron had a threshold; that is, the neuron would not produce a measurable response until its underlying linear response exceeded a threshold value (e.g. Movshon et al., 1978*a*; Tolhurst & Dean, 1987, 1991; Robson et al., 1988). It has since been proposed that these same phenomena may also be consistent with the more general proposal that simple cell responses are subject to half-squaring and to contrast normalization (Heeger, 1991, 1992*a*,*b*). It is proposed that the squared response of each neuron is scaled by a normalizing signal based on the total contrast energy of the stimulus.

Half-squaring is a specific implementation of a more general model, where the output might be subject to expansion with an exponent (n) other than 2 (Albrecht & Hamilton, 1982; Albrecht & Geisler, 1991). The exponent of 2 has theoretical attractions since it relates specifically to stimulus energy, but it is true that the expansive normalization model has fit our data better when the exponent was free to find some other value. Even so, the resulting best-fit values of the exponent have generally been close to 2.0, especially as the quality of the fits (value of χ^2) was generally tolerant of changes in n.

We have tried to distinguish between the normalization and threshold models by a head-to-head comparison of their fits to the various experimental data.

Half-squaring and other expansive nonlinearities

In experiments where the stimuli all have the same contrast, we often do not need to consider the full effects of normalization (e.g. Figs. 3 and 5). Then, we need only compare the expansive nonlinearity *per se* with the threshold model. For many of the data sets, the expansive nonlinearity and threshold gave fits that were almost equally good. Overall, the expansive nonlinearity gave the better values of χ^2, even when we fixed the exponent (n) in the model to the value of 2 needed for exact half-squaring. This may suggest that the expansive nonlinearity is the better description of the actual mechanisms of action potential generation.

The threshold model did produce acceptable descriptions of the results in many individual experiments. However, the value of the threshold parameter T needed to achieve those fits often seemed to be unrealistically large, as was noted by Robson et al. (1988) and Tadmor and Tolhurst (1989). In many cases, we would have to accept that up to 75% of the neuron's response to high-contrast stimuli remained below threshold. Furthermore, the fits for two different experiments on the same neuron often produced very different estimates of the threshold parameter (e.g. Figs. 6B *vs.* 7A).

This is somewhat paradoxical, since there is no biophysical evidence that neurons show a squaring or expansive relation between membrane potential and firing rate. There is, however, much evidence that neurons do require a threshold amount of depolarization before action potentials are generated (see Introduction for citations). The biophysical threshold may represent a neuronal approximation to the desirable operation of squaring.

Sometimes, one model seemed to fit better than the other only because it was a better description of a single low-amplitude response at a single stimulus condition. For instance, in Fig. 5, the expansive nonlinearity gives the better fit to the data (χ^2 of 0.81 compared to 2.24 for the threshold model). But this difference in χ^2 may hinge on the response to a single stimulus condition (at a spatial phase of 0 deg), where the neuron produced only 25 action potentials during a total period of 50 s. If, in fact, the neuron had produced eight action potentials for that stimulus condition, the threshold model would have become the better fit (χ^2 of 1.34 compared to 1.49 for the expansive nonlinearity). Similarly, the data of Fig. 2 are better fit by the threshold model (χ^2 of 1.21 *vs.* 1.64 for the expansive nonlinearity); but if the neuron had given a

response to a contrast of 0.056 of only one action potential per stimulus trial instead of one action potential every three trials, the expansive nonlinearity would have become the better fit (χ^2 of 1.36 *vs.* 1.37 for the threshold model).

One crucial feature in the fits of the different models is how suddenly or gradually the neuron begins to give reliable responses to stimuli that evoke little activity (e.g. at low contrasts). A more gradual transition between activity and inactivity is indicative that an expansive nonlinearity will fit better than the abrupt-threshold model. Unfortunately, it is very difficult to measure the crucial small responses reliably enough to allow a definitive proof that one model really is a better fit than the other. The dependence of the fits on low-amplitude responses is especially problematic since we have weighted the data in the fits according to our estimates of the variances of the responses. We suspect that these estimates of variance are not fully accurate, especially for very low-response amplitudes when the statistics of response are not normally distributed (Dean, 1981*b*; Tolhurst, 1989). In addition, the responsiveness of cortical neurons often changes repeatedly during the time scale of an experiment (e.g. Tolhurst et al., 1981, 1983), and this might easily obscure the abruptness of a true threshold.

The receptive-field data (Fig. 3) were better fitted by the expansive output model than by the threshold model originally employed by Tadmor and Tolhurst (1989). But this may not be due to the superiority of the expansive nonlinearity *per se*. Tadmor and Tolhurst (1989) could not compensate for the threshold output nonlinearity on the neuron's spatial-frequency tuning data; they would have had to "invent" the thresholded tails of the tuning curves. They recognized that this interfered with the quality of their fits. With the expansive nonlinearity, on the other hand, it was quite easy to account for the output nonlinearity in the bar *and* in the grating data (simply by computing the square roots of responses). The ease of working with the expansive output model is one good reason for preferring it, all other things being equal.

We have compared the space- and frequency-domain measurements, but only up to an arbitrary amplitude scaling-factor since the temporal properties of the bar and grating stimuli were very different (see Movshon et al., 1978*a* for discussion). In particular, we used two unconstrained scale-factor parameters (A_E and A_O) to fit the data. In the design of future experiments, by choosing the bar and grating stimuli to have the same temporal properties, we might fit the data with only one scale factor (the ratio of even to odd component). Better yet, by measuring response phase as well as amplitude (Hamilton et al., 1989), we might be able to dispense with both of these scale factors, which would allow for stronger tests of the models.

Contrast normalization and response suppression

The basic mathematical formulation of the normalization model [e.g. eqn. (10)] is more elegant and is easier to apply than that for the threshold model [eqn. (4)]. The normalization model is, then, preferable if only as an empirical description of data. This is especially the case for the phenomena of response saturation and nonspecific suppression. These phenomena are outside the scope of the threshold model, but the normalization model provides a coherent theory to account for them.

We have examined a number of experimental situations where we have had to consider the effects of normalization as well as the expansive nonlinearity: i.e. whenever the stimuli under comparison have differed in their total space-averaged or time-averaged contrast energy. It is in these experiments that the normalization model has been particularly successful, and that the threshold model has been most disadvantaged. For instance, the normalization model has given an explanation for the slight divergence in the functions relating response to contrast in the experiment with pairs of bar stimuli (Fig. 2). The threshold model does not predict such behavior.

The normalization model has also been successful in explaining the relationships between the responses to moving and modulated gratings (Tables 1 and 2). Simple linear theories (Reid et al., 1987, 1991; Tolhurst & Dean, 1991) consistently find discrepancies in estimates of the directional index and in the predicted and measured amplitudes of responses to moving gratings. We have confirmed (Table 1) that the discrepancies in the estimates of directional index can be accounted for by expansive nonlinearities alone (Albrecht & Geisler, 1991; DeAngelis et al., 1993). However, the normalization model must be implemented fully to account for the discrepancies in the predictions of the amplitudes of the responses to moving gratings from the responses to modulated gratings, since the time-averaged energy of moving and modulated gratings is different (Tolhurst & Heeger, 1997). Indeed, the normalization model did perform reasonably well, while the threshold model performed very poorly (Table 2).

We have described in detail some experiments designed to examine the suppressive action of gratings moving in the nonpreferred direction for simple cells with profound directional selectivity. The results were interpreted by Dean et al. (1980) as showing that the suppressing grating caused a reduction in the neuron's gain or responsivity, by an arithmetic operation similar to division. When the suppressing contrast (c_s) is high compared to the activating contrast (c_a) or semisaturation contrast (σ), the prediction of the full normalization model [eqn. (12)] can be approximated to

$$\bar{R} \approx \bar{R}_{\max}\left(\frac{c_a}{c_s}\right)^n \tag{24}$$

which does indeed show a divisive action of the suppressing grating.

However, the normalization model makes much more specific predictions than were envisaged by Dean et al. (1980). Figs. 8–10 show experiments with a variety of combinations of activating and suppressing contrasts. The normalization model, expressed simply as eqn. (12), is very successful at describing the pattern of results. In particular, the model properly predicts the magnitude of the suppressing effect and the way in which it depends upon the contrast of the suppressing grating (Figs. 9B and 10B). The most dramatic justification of the model is that it provides a reasonable description of up to four different experiments on a single neuron with many different combinations of activating and suppressing contrast, using only three variable parameters ($\bar{R}_{\max}$, n, and σ). Furthermore, these are the same parameters as would be used to fit the control relation between response and contrast for gratings moving in the preferred direction.

The fits of the model curves in Figs. 9 and 10 are not perfect, of course. Either the model or the experiments need refinement. The effects of different suppressing contrasts in Figs. 9A were examined in different experiments performed at different times. It is possible that the neuron's responsiveness may have changed between experiments, and this may explain the poor fits to the results for suppressing contrasts of 0.11 and 0.18. It would be better to perform a single experiment with all the stimulus conditions interleaved, rather than to run several different experiments sequentially.

We could certainly devise an equation that would be a better fit to the data than is eqn. (12). This might need extra variable pa-

rameters to differentially weight the suppressing actions of given contrasts, or we might give different exponents to the numerator and denominator in the equation. For instance, we have found good fits to some of the data in this paper using models based on the following alternative to eqn. (10):

$$\bar{R} = \bar{R}_{\max} \frac{c^n}{\sigma^2 + c^2} \tag{25}$$

The difference between eqns. (10) and (25) is subtle: does normalization occur before [eqn. (10)] or after [eqn. (25)] the expansive nonlinearity. The latter formulation allows better fits in some neurons where the responses continue to creep upwards at high contrasts without saturating totally. However, we are not trying to provide an empirical description of the data; we are evaluating a specific model of neuronal interactions, and we should change the equations only if there is a physiological or theoretical basis to justify it.

Other forms of nonlinearity

Thus, we have found that the expansive, contrast-normalization model (Heeger, 1991, 1992*a*,*b*, 1993; Carandini & Heeger, 1994; Carandini et al., 1997; Tolhurst & Heeger, 1997) can account very successfully for a wide range of phenomena. However, we should not be led into believing that the normalization model is capable of explaining everything. Some simple cells do show undoubted signs of nonlinear spatial summation, perhaps because of failures in push–pull synaptic arrangements (Movshon et al., 1978*a*; Dean & Tolhurst, 1983; Mullikin et al., 1984; Tolhurst & Dean, 1990; Atick & Redlich, 1990).

For instance, we have found that the half-squaring, contrast-normalization model gives an excellent description of one of the receptive-field profiles shown by Tadmor and Tolhurst (1989, their Fig. 4, our Fig. 3). However, the model does not explain the second profile shown in that paper (their Fig. 6); that neuron was found by Tolhurst and Dean (1987) to be subject to serious nonlinearities of spatial summation. We have also shown an example (Fig. 6, Tables 1 and 2) where the normalization model has been quite successful in describing the detailed relationships between the responses to moving and modulated gratings for one of the simple cells described by Tolhurst and Dean (1991, their Fig. 1A). But again, Tolhurst and Dean (1991, their Fig. 1B) and McLean et al. (1994) found that some simple cells show pronounced directional selectivity even though their responses suggested that their receptive fields were spatiotemporally separable. This would be difficult to explain by a simple output nonlinearity such as half-squaring or contrast normalization.

Indeed, if we look critically at the way in which the various models dealt with the polar plot of response phase and amplitude (Fig. 6), we see various inconsistencies. The threshold model needed a very high threshold to fit the "wasp-waist" of the distorted ellipse (Fig. 6B), and this high value was not confirmed by the same neuron's response *versus* contrast relationship (Fig. 7A). The normalization model needed a very high exponent (4–10, in Fig. 6D) to fit the wasp-waist, but a low value (1.30) to fit the response *versus* contrast curve (Fig. 7B). This low exponent also accounted for relationships between directional index and the amplitudes of response to moving and modulated gratings (Tables 1 and 2). Perhaps, there are other nonlinear phenomena that make the polar plot nonelliptical (Emerson & Huang, 1997). It will take careful experimentation at a variety of stimulus contrasts to dissect such additional nonlinearities from those that we are currently attributing to contrast normalization.

Conclusions

Heeger (1991, 1992*a*,*b*) suggested that the contrast-normalization model (Robson, 1988; Bonds, 1989; Albrecht & Geisler, 1991) might be capable of explaining many observations on cat simple cells, which were previously attributed to a threshold model. We have shown that the normalization model does, indeed, provide a good description of actual experimental data and that it is capable of providing a satisfying explanation for a very wide variety of phenomena. However, by applying the model rigorously to real data, we have found some instances where the normalization model, although impressive, does not perform quite well enough. This suggests that there are other, more subtle nonlinearities in simple cell behavior which need to be recognized.

Acknowledgments

Most of the experiments reported here were carried out in collaborations between D.J. Tolhurst and A.F. Dean, and they were supported by an MRC grant (D.J.T.) and a Wellcome Trust grant to F.W. Campbell. We thank A.F. Dean and R.F. Hess for allowing us to present previously unpublished data collected in collaboration with D.J. Tolhurst. This paper was mostly written while D.J. Tolhurst was on a visit to Stanford, supported by an NIMH Grant (MH50228), and an Alfred P. Sloan Research Fellowship to D.J. Heeger.

References

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

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

Albrecht, D.G. & Hamilton, D.B. (1982). Striate cortex of monkey and cat: Contrast response function. *Journal of Neurophysiology* **48**, 217–237.

Andrews, B.W. & Pollen, D.A. (1979). Relationship between spatial frequency selectivity and receptive field profile of simple cells. *Journal of Physiology* **287**, 163–176.

Atick, J.J. & Redlich, A.N. (1990). Mathematical model of the simple cells in the visual cortex. *Biological Cybernetics* **63**, 99–109.

Bauman, L.A. & Bonds, A.B. (1991). Inhibitory refinement of spatial frequency selectivity in single cells of the cat striate cortex. *Vision Research* **31**, 933–944.

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

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

Carandini, M., Heeger, D.J. & Movshon, J.A. (1997). Linearity and gain control in V1 simple cells. In *Cerebral Cortex, Vol.XII, Cortical Models*, ed. Jones, E.G. & Ulinski, P.S., New York: Plenum Press.

Dean, A.F. (1981*a*). The relationship between response amplitude and contrast for cat striate cortical neurones. *Journal of Physiology* **318**, 413–427.

Dean, A.F. (1981*b*). The variability of discharge of simple cells in the cat striate cortex. *Experimental Brain Research* **44**, 437–440.

Dean, A.F. & Tolhurst, D.J. (1983). On the distinctness of simple and complex cells in the visual cortex of the cat. *Journal of Physiology* **344**, 305–325.

Dean, A.F., Hess, R.F. & Tolhurst, D.J. (1980). Divisive inhibition involved in directional selectivity. *Journal of Physiology* **308**, 84–85P.

DeAngelis, G.C., Ohzawa, I. & Freeman, R.D. (1993). Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. II, Linearity of temporal and spatial summation. *Journal of Neurophysiology* **69**, 1118–1135.

DeAngelis, G.C., Robson, J.G., Ohzawa, I. & Freeman, R.D. (1992). Organization of suppression in receptive fields of neurons in cat visual cortex. *Journal of Neurophysiology* **68**, 144–163.

De Valois, K. & Tootell, R. (1983). Spatial-frequency specific inhibition in cat striate cortex cells. *Journal of Physiology* **336**, 359–376.

De Valois, R.L., Albrecht, D.G. & Thorell, L.G. (1978). Cortical cells: Bar and edge detectors or spatial-frequency analyzers. In *Frontiers in Visual Science*, ed. Cool, S.J. & Smith, E.L., pp. 544–556. New York: Springer Verlag.

De Valois, R.L., Thorell, L.G. & Albrecht, D.G. (1985). Periodicity of striate cortex cell receptive fields. *Journal of the Optical Society of America A*, **2**, 1115–1123.

Emerson, R.C. & Huang, M.C. (1997). Quadrature subunits in directionally selective simple cells: Counterphase and drifting grating responses. *Visual Neuroscience* **14**, 373–385.

Enroth-Cugell, C. & Robson, J.G. (1966). The contrast sensitivity of retinal ganglion cells of the cat. *Journal of Physiology* **187**, 517–552.

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

Ferster, D. & Jagadeesh, B. (1992). EPSP-IPSP interactions in cat visual cortex studied with *in vivo* whole-cell patch recording. *Journal of Neuroscience* **12**, 1262–1274.

Geisler, W.S. & Albrecht, D.G. (1992). Cortical neurons: Isolation of contrast gain control. *Vision Research* **32**, 1409–1410.

Geisler, W.S. & Albrecht, D.G. (1995). Bayesian analysis of identification performance in monkey visual cortex: Nonlinear mechanisms and stimulus certainty. *Vision Research* **35**, 2723–2730.

Glezer, V.D., Tsherbach, T.A., Gauselman, V.E. & Bondarko, V.M. (1982). Spatiotemporal organization of receptive fields of the cat striate cortex. *Biological Cybernetics* **43**, 35–49.

Grossberg, S. (1988). Nonlinear neural networks: Principles, mechanisms and architectures. *Neural Networks* **1**, 17–61.

Gulyas, B., Orban, G., Duysens, J. & Maes, H. (1987). The suppressive influence of moving textured backgrounds on responses of cat striate cortex neurons to moving bars. *Journal of Neurophysiology* **57**, 1767–1791.

Hamilton, D.B., Albrecht, D.G. & Geisler, W.S. (1989). Visual cortical receptive fields in monkey and cat: Spatial and temporal phase transfer function. *Vision Research* **29**, 1285–1308.

Hammond, P. & MacKay, D.M. (1981). Modulatory influences of moving textured backgrounds on responsiveness of simple cells in feline striate cortex. *Journal of Physiology* **319**, 431–442.

Heeger, D.J. (1991). Nonlinear model of neural responses in cat visual cortex. In *Computational Models of Visual Processing*, ed. Landy, M. & Movshon, J.A., pp. 119–133. Cambridge, Massachusetts: MIT Press.

Heeger, D.J. (1992*a*). Normalization of cell responses in cat striate cortex. *Visual Neuroscience* **9**, 181–197.

Heeger, D.J. (1992*b*). Half-squaring in responses of cat simple cells. *Visual Neuroscience* **9**, 427–443.

Heeger, D.J. (1993). Modeling simple-cell direction selectivity with normalized, half-squared, linear operators. *Journal of Neurophysiology* **70**, 1885–1898.

Hubel, D.H. & Wiesel, T.N. (1959). Receptive fields of single neurones in the cat's striate cortex. *Journal of Physiology* **148**, 574–591.

Ikeda, H. & Wright, M.J. (1974). Sensitivity of neurones in visual cortex (area 17) under different levels of anaesthesia. *Experimental Brain Research* **20**, 471–484.

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

Jagadeesh, B., Wheat, H.S. & Ferster, D. (1993). Linearity of summation of synaptic potentials underlying direction selectivity in simple cells of the cat visual cortex. *Science* **262**, 1901–1904.

Jones, J.P. & Palmer, L.A. (1987). An evaluation of the two-dimensional Gabor filter model of simple receptive fields in the cat striate cortex. *Journal of Neurophysiology* **58**, 1233–1258.

Koike, H., Mano, N., Okada, Y. & Oshima, T. (1970). Repetitive impulses generated in fast and slow pyramidal tract cells by intracellularly applied current steps. *Experimental Brain Research* **11**, 263–281.

Kulikowski, J.J. & Bishop, P.O. (1981). Linear analysis of the responses of simple cells in cat striate cortex. *Experimental Brain Research* **44**, 386–400.

Li, C.-Y. & Creutzfeldt, O. (1984). The representation of contrast and other stimulus parameters by single neurons in area 17 of the cat. *Pflügers Archives* **401**, 304–314.

Maffei, L. & Fiorentini, A. (1973). The visual cortex as a spatial frequency analyzer. *Vision Research* **13**, 1255–1267.

Maffei, L. & Fiorentini, A. (1976). The unresponsive regions of visual cortical receptive fields. *Vision Research* **16**, 1131–1139.

Maffei, L., Morrone, M.C., Pirchio, M. & Sandini, G. (1979). Responses of visual cortical cells to periodic and non-periodic stimuli. *Journal of Physiology* **296**, 27–47.

Marr, D. (1970). A theory for cerebral neocortex. *Proceedings of the Royal Society B* **176**, 161–234.

Mason, A. & Larkman, A.U. (1990). Correlations between morphology and electrophysiology of pyramidal neurons in slices of rat visual cortex. II. Electrophysiology. *Journal of Neuroscience* **10**, 1415–1428.

McCormick, D.A., Connors, B.W., Lighthall, J.W. & Prince, D.A. (1985). Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *Journal of Neurophysiology* **54**, 782–806.

McLean, J., Raab, S. & Palmer, L.A. (1994). Contribution of linear mechanisms to the specification of local motion by simple cells in areas 17 & 18 of the cat. *Visual Neuroscience* **11**, 271–294.

Morrone, M.C., Burr, D.C. & Maffei, L. (1982). Functional implications of cross-orientation inhibition of cortical visual cells. I. Neurophysiological evidence. *Proceedings of the Royal Society B* **216**, 335–354.

Movshon, J.A., Thompson, I.D. & Tolhurst, D.J. (1978*a*). Spatial summation in the receptive fields of simple cells in the cat's striate cortex. *Journal of Physiology* **283**, 53–77.

Movshon, J.A., Thompson, I.D. & Tolhurst, D.J. (1978*b*). Spatial and temporal contrast sensitivity of neurones in areas 17 and 18 of the cat's visual cortex. *Journal of Physiology* **283**, 101–120.

Mullikin, W.H., Jones, J.P. & Palmer, L.A. (1984). Receptive-field properties and laminar distribution of X-like and Y-like simple cells in cat area 17. *Journal of Neurophysiology* **52**, 350–371.

Nestares, O. & Heeger, D.J. (1997). Modeling the apparent frequency-specific suppression in simple cell responses. *Vision Research* (in press).

Ogawa, T., Ito, S. & Kato, H. (1981). Membrane characteristics of visual cortical neurons in *in vitro* slices. *Brain Research* **226**, 315–329.

Ohzawa, I., Sclar, G. & Freeman, R.D. (1982). Contrast gain control in the cat visual cortex. *Nature* **298**, 266–268.

Palmer, L.A. & Davis, T.L. (1981). Receptive field structure in cat striate cortex. *Journal of Neurophysiology* **46**, 260–276.

Pollen, D.A., Gaska, J.P. & Jacobson, L.D. (1988). Responses of simple and complex cells to compound sine-wave gratings. *Vision Research* **28**, 25–39.

Press, W.H., Flannery, B.P., Teukolsky, S.A. & Vetterling, W.T. (1986). *Numerical Recipes: The Art of Scientific Computing*. Cambridge, UK: Cambridge University Press.

Reid, R.C., Soodak, R.E. & Shapley, R.M. (1987). Linear mechanisms of directional selectivity in simple cells of cat striate cortex. *Proceedings of the National Academy of Sciences of the U.S.A.* **84**, 8740–8744.

Reid, R.C., Soodak, R.E. & Shapley, R.M. (1991). Directional selectivity and spatiotemporal structure of receptive fields of simple cells in cat striate cortex. *Journal of Neurophysiology* **66**, 505–529.

Robson, J.G. (1988). Linear and non-linear operations in the visual system. *Investigative Ophthalmology and Visual Science* (Suppl.) **29**, 117.

Robson, J.G., Tolhurst, D.J., Freeman, R.D. & Ohzawa, I. (1988). Simple cells in the visual cortex of the cat can be narrowly tuned for spatial frequency. *Visual Neuroscience* **1**, 415–419.

Schumer, R.A. & Movshon, J.A. (1984). Length summation in simple cells of cat striate cortex. *Vision Research* **24**, 565–571.

Shapley, R.M., Reid, R.C. & Soodak, R. (1991). Spatiotemporal receptive fields and direction selectivity. In *Computational Models of Visual Processing*, ed. Landy, M. & Movshon, J.A., pp. 109–118. Cambridge, MA: MIT Press.

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

Tadmor, Y. & Tolhurst, D.J. (1989). The effect of threshold on the relationship between the receptive-field profile and the spatial-frequency tuning curve in simple cells of the cat's striate cortex. *Visual Neuroscience* **3**, 445–454.

Tolhurst, D.J. (1989). The amount of information transmitted about contrast by neurones in the cat's visual cortex. *Visual Neuroscience* **2**, 409–413.

Tolhurst, D.J. & Dean, A.F. (1987). Spatial summation by simple cells in the striate cortex of the cat. *Experimental Brain Research* **66**, 607–620.

Tolhurst, D.J. & Dean, A.F. (1990). The effects of contrast on the lin-

earity of spatial summation of simple cells in the cat's striate cortex. *Experimental Brain Research* **79**, 582–588.

TOLHURST, D.J. & DEAN, A.F. (1991). Evaluation of a linear model of directional selectivity in simple cells of the cat's striate cortex. *Visual Neuroscience* **6**, 421–428.

TOLHURST, D.J. & HEEGER, D.J. (1997). Contrast normalization and a linear model for the directional selectivity of simple cells in cat striate cortex. *Visual Neuroscience* **14**, 19–25.

TOLHURST, D.J. & THOMPSON, I.D. (1981). On the variety of spatial-frequency selectivities shown by neurons in area 17 of the cat. *Proceedings of the Royal Society B* **213**, 183–199.

TOLHURST, D.J., MOVSHON, J.A. & DEAN, A.F. (1983). The statistical reliability of single neurons in cat and monkey visual cortex. *Vision Research* **23**, 775–785.

TOLHURST, D.J., MOVSHON, J.A. & THOMPSON, I.D. (1981). The dependence of response amplitude and variance of cat visual cortical neurones on stimulus contrast. *Experimental Brain Research* **41**, 414–419.

VOGELS, R., SPILEERS, W. & ORBAN, G.A. (1989). The response variability of striate cortical neurons in behaving monkey. *Experimental Brain Research* **77**, 432–436.

WEBSTER, M.A. & DE VALOIS, R.L. (1985). Relationship between spatial-frequency & orientation tuning of striate-cortex cells. *Journal of the Optical Society of America A* **2**, 1124–1132.