MODELING THE APPARENT FREQUENCY-SPECIFIC SUPPRESSION IN SIMPLE CELL RESPONSES.

Oscar Nestares

Instituto de Óptica "Daza de Valdés" (C.S.I.C.), Serrano 121, 28006-Madrid, SPAIN

David J. Heeger

Department of Psychology, Stanford University, Stanford, CA 94305, U.S.A.

Running head: Modeling Frequency-Specific Suppression

Keywords: striate cortex, simple cell, spatial frequency, suppression, contrast normalization.

Abstract

Simple cells in cat striate cortex are selective for spatial frequency. It is widely believed that this selectivity arises simply because of the way in which the neurons sum inputs from the lateral geniculate nucleus. Alternate models, however, advocate the need for frequency-specific inhibitory mechanisms to refine the spatial frequency selectivity. Indeed, simple cell responses are often suppressed by superimposing stimuli with spatial frequencies that flank the neuron's preferred spatial frequency.

In this article, we compare two models of simple cell responses head-to-head. One of these models, the *flanking-suppression model*, includes an inhibitory mechanism that is specific to frequencies that flank the neuron's preferred spatial frequency. The other model, the *nonspecific-suppression model*, includes a suppressive mechanism that is very broadly tuned for spatial frequency. Both models also include a rectification nonlinearity and both may include an additional accelerating (e.g., squaring) output nonlinearity. We demonstrate that both models can be consistent with the *apparent* flanking suppression. However, based on other experimental results, we argue that the nonspecific-suppression model is more plausible. We conclude that the suppression is probably broadly tuned for spatial frequency and that the *apparent* flanking suppression is actually due to distortions introduced by an accelerating output nonlinearity.

1 Introduction

Simple cells in cat striate cortex (area 17) are selective for spatial frequency; each neuron responds most vigorously to a preferred spatial frequency, and the spatial frequency tuning curves typically have relatively narrow bandwidths (e.g., Campbell, Cooper and Enroth-Cugell, 1969; Maffei and Fiorentini, 1973; Robson, Tolhurst, Freeman and Ohzawa, 1988b). However, different models have been proposed for the mechanism(s) underlying spatial frequency selectivity. Linear models (e.g., Movshon, Thompson and Tolhurst, 1978a) purport to explain spatial frequency selectivity simply in terms of the widths and the number of the ON and OFF subregions. For example, a neuron with many, thin subregions would be narrowly tuned for high spatial frequencies. Alternate models advocate the need for frequency-specific inhibitory mechanisms to refine the spatial frequency selectivity (e.g., Bauman and Bonds, 1991).

Indeed, there is some evidence that appears to support frequency-specific inhibition (Movshon, Thompson and Tolhurst, 1978b; DeValois and Tootell, 1983; DeValois, Thorell and Albrecht, 1985; Bauman and Bonds, 1991). Bauman and Bonds (1991), for example, recorded from simple cells while presenting stimuli made of superimposed pairs of moving grating patterns. The first grating (called the base grating) was chosen to have the neuron's preferred spatial frequency and orientation. The second (mask) grating had the same orientation, but its spatial frequency was varied. Bauman and Bonds found that the response to the base grating was often suppressed by superimposing the mask grating. The suppression depended on the spatial frequency of the mask grating, and it was greatest for mask frequencies that flanked the neuron's preferred spatial frequency. These results, of course, violate a strictly linear model. A plausible interpretation is that the suppression results from a frequency-specific inhibitory mechanism.

On the other hand, there is evidence that the suppression is broadly tuned for spatial frequency. Bonds (1989), for example, performed an experiment much like that described above in which suppression was quantified by superimposing a pair of moving gratings. In this case, however, the mask grating was rotated to a very different orientation, at the limit of the neuron's orientation tuning curve, so that the mask grating never evoked a response on its own. Suppression from the mask grating in this case was found to be broadly tuned for spatial frequency. The results of these two experiments can be reconciled by assuming that the suppression might depend on flanking spatial frequencies near the neuron's preferred orientation, but on a broad range of spatial frequencies at different orientations.

In this article, we offer a different explanation: that the suppression is broadly tuned for spatial frequency at all orientations and that the *apparent* flanking suppression is actually due to distortions introduced by an accelerating output nonlinearity.

We compare two models of simple cell responses head-to-head. One of these models includes an inhibitory mechanism that is specific to frequencies that flank the neuron's preferred spatial frequency. We will refer to this model as the *flanking-suppression model*. The other model also includes a suppressive mechanism, but it is not specific to flanking spatial frequencies. We will refer to this model as the *nonspecific-suppression model*. Both models also include a rectification nonlinearity and both may include an additional accelerating (e.g., squaring) output nonlinearity. We demonstrate that both models can be consistent with the *apparent* flanking suppression in the data reported by Bauman and Bonds (1991). Through a careful analysis of the various nonlinearities in the two models, we explain why each of the models succeeds in explaining these results.

However, based on other experimental data, we argue that the nonspecific-suppression model is more plausible. The most critical failure of the flanking-suppression model is that it predicts, contrary to experimental results (Albrecht and Hamilton, 1982; Skottun, Bradley, Sclar, Ohzawa and Freeman, 1987), that spatial frequency tuning bandwidth should vary systematically with stimulus contrast.

2 Models

The two models that we compare are similar to one another. In fact, they are both special cases of a more general model as discussed below. Here, we do not attempt to make the models biologically realistic; they are presented as mathematical abstractions, whose goal is to describe the information transformations rather than the details of the neuronal mechanisms that perform those transformations. The models can, however, be implemented with biologically reasonable mechanisms (Carandini and Heeger, 1994; Carandini, Heeger, and Movshon, 1996).

In both models, an underlying linear response serves as the basis for spatial frequency (and orientation) selectivity. Then, in both models, the linear responses are rectified and normalized to produce an overt (firing rate) response. The exact form of normalization is what differs between the two models.

The response of a linear visual neuron is a weighted sum, over local space and recently past time, of the distribution of light intensity values in the stimulus. According to the linear model, spatial frequency, temporal frequency, and orientation selectivity arise from the shape and location of excitatory (positively weighted) and inhibitory (negatively weighted) subregions of the receptive field.

Simple cells have often been characterized as rectified linear operators (e.g., Movshon *et al.*, 1978a). The rectification guarantees that model responses are always positive, reflecting the fact that extracellular neural responses (firing rates) are by definition positive. Variants of this characterization have used different types of rectification. For example, over-rectification is halfwave-rectification but with a threshold: the neuron has to reach a certain level of excitation before it will fire action potentials. Half-squaring (halfwave-rectification followed by squaring) is quite similar to over-rectification, but with a "softer" threshold (e.g., Heeger, 1992b). In this article, we consider rectification nonlinearities with a variable exponent in which exponents of 1 and 2 yield perfect halfwave-rectification and half-squaring, respectively.

There are a number of problems with the (rectified) linear model of simple cells. One major fault with this model is the fact that simple cell responses saturate (level off) at high contrasts. The response of a halfwave-rectified linear neuron would increase in proportion to stimulus contrast over the entire range of contrasts. The response of a half-squared linear neuron would increase as the square of contrast. A second major fault with the linear model is revealed by testing superposition. A typical simple cell responds vigorously to its preferred orientation but not at all to the perpendicular orientation. For a rectified linear neuron, regardless of the exponent, the response to the superimposed pair of grating stimuli (preferred plus perpendicular) would equal the response to the preferred stimulus presented alone. However, the response of a simple cell to a superimposed pair of orthogonal gratings is about half that predicted (e.g., Bonds, 1989), a phenomenon known as cross-orientation inhibition.

Response normalization was originally proposed by Robson (1988a) to provide explanations for these failures of the linear model. This idea has been expanded and formalized by one of us (Heeger 1991, 1992a,b, 1993; Carandini and Heeger, 1994; Carandini *et al.*, 1996) and by Albrecht and Geisler (1991). These papers and others (e.g., Bonds, 1989; DeAngelis, Robson, Ohzawa and Freeman, 1992; Tolhurst and Heeger, 1996a,b) have shown that response normalization is capable, in principle, of explaining a wide variety of empirical phenomena. 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 much of Grossberg's theoretical work on nonlinear neural networks (for review see Grossberg, 1988).

Response normalization means that each neuron's underlying response is divided by a quantity proportional to the pooled activity of a large number of other neurons from the nearby cortical "neighborhood". The normalization pool includes neurons tuned to all different orientations and a range of spatial frequencies. Activity in this large pool of neurons partially suppresses the response of each individual neuron. The effect of this divisive suppression is that the response of each neuron is normalized (rescaled) with respect to stimulus contrast. The normalization model exhibits response saturation because the divisive suppression increases with stimulus contrast. The normalization model exhibits cross-orientation inhibition because the normalization pool includes neurons with a wide variety of tuning properties, many of which respond to orthogonal gratings.

In the nonspecific-suppression model, the normalization pool includes neurons with a broad range of spatial frequency preferences. In the flanking-suppression model, the normalization pool primarily includes those neurons tuned for particular spatial frequencies flanking the preferred spatial frequency. In both models, the normalization limits the dynamic range of the responses. In the flanking-suppression model, the normalization plays the additional role of sharpening the spatial frequency tuning curve.

2.1 *The Nonspecific-Suppression Model*

According to the nonspecific-suppression model with strict half-squaring, the overt response of a simple cell to any stimulus is given by:

$$R(t) = K \frac{\lfloor L(t) \rfloor^2}{s^2 + \sum_i E_i(t)} \tag{1}$$

where K and s are constants, $\lfloor\cdot\rfloor$ means halfwave-rectification, $L(t)$ is the underlying (orientation and spatial frequency tuned) linear response, and i indexes over the spatial frequencies and orientations included in the normalization pool. The energy, $E_i(t)$, in Eq. 1 is the sum of four half-squared, linear responses with phases in steps of 90°, but with otherwise identical tuning properties. The summation, $\sum_i E_i(t)$, in the denominator includes the term $\lfloor L(t) \rfloor^2$ that appears in the numerator (i.e., each neuron suppresses itself).

For a moving grating stimulus, simple cell responses are often summarized by the amplitude of the first harmonic (equal to the stimulus temporal frequency) of the response time-course. From Eq. 1, response amplitude can be expressed as:

$$|R| = R_{\max} \frac{c^2}{\sigma^2 + c^2} |H(f)|^2, \tag{2}$$

where c is the contrast of the grating, f is the spatiotemporal frequency and orientation of the grating, $H(f)$ is the amplitude and phase of the underlying linear response, $R_{\max}$ is the maximum attainable response, and σ is a new constant that depends on K and s. As long as σ is nonzero, the normalized response will always be a value between 0 and $R_{\max}$, saturating for high contrasts. The σ parameter is often referred to as a semi-saturation constant because it equals the contrast of a moving grating stimulus that evokes half the maximum attainable (fully saturated) response. For the simulation results reported below, we chose a value of $\sigma = 0.1$, near to the average value for cat simple cells of 0.15 (Albrecht and Hamilton, 1982). However, this is not a critical choice because, as explained below, stimulus contrasts were chosen in proportion to the semi-saturation constant.

The squaring in the above equations is mathematically convenient, but in fact, the exponent fit to response-versus-contrast data varies from one neuron to the next (Albrecht and Hamilton, 1982; Sclar, Maunsell, and Lennie, 1990; Albrecht and Geisler, 1991; Tolhurst and Heeger, 1996b), ranging typically between 1 and 4. Based on a mechanistic description of the synaptic processes underlying response normalization, Carandini *et al.* (1996) developed a related formulation with a variable exponent n:

$$R(t) = K \frac{\lfloor L(t) \rfloor^n}{\left[\sqrt{s^2 + \sum_i E_i(t)}\right]^n}. \tag{3}$$

Note that when $n = 2$ this simplifies to Eq. 1. The response amplitude for moving gratings is now given by:

$$|R| = R_{\max} \frac{c^n}{\left(\sqrt{\sigma^2 + c^2}\right)^n} \, |H(f)|^n. \tag{4}$$

When the exponent n is increased (e.g., 3 instead of 2), the slope of the response versus contrast curve is steeper and the effective spatial frequency bandwidth is narrower.

2.2 The Flanking-Suppression Model

This model can be generalized to allow for spatial frequency dependent suppression by assigning different weights to the different spatial frequency bands in the normalization pool:

$$R(t) = K \frac{\lfloor L(t) \rfloor^n}{\left[\sqrt{s^2 + \sum_i w_i \sum_k E_{ik}(t)}\right]^n}, \tag{5}$$

where we have separated the summation over spatial frequency (indexed by i) from the summation over orientation (indexed by k), so that the weights w_i depend only on spatial frequency.

The response amplitude for moving gratings is now expressed as:

$$|R| = R_{\max} \frac{c^n}{\left[\sqrt{\sigma^2 + c^2\, w^2(f)}\right]^n} |H(f)|^n. \tag{6}$$

where $w(f)$ accounts for the frequency dependent weighting on the suppression.

An unfortunate complication in this formulation is that when $w(f) \neq 1$ or when $n \neq 2$, σ no longer corresponds to the semi-saturation contrast. In the simulations, we always adjusted the value of σ so that a moving grating with the preferred spatial frequency and a contrast of 0.1 would evoke half the maximum attainable response.

2.3 Detailed Methods

The underlying linear receptive fields were typically chosen to have spatial frequency bandwidths of 2 octaves (full-width at half height), but for Fig. 1

we chose the spatial frequency bandwidths by hand to match the physiological data. The orientation bandwidths of the underlying linear receptive fields were 60 degrees (full-width at half height), and the temporal frequency bandwidths were large so that the amplitudes of the underlying linear responses would be identical for 2 Hz and 3 Hz gratings.

The overt responses of the simulated neurons exhibit somewhat narrower bandwidths due to the various nonlinearities. Specifically, for our particular choice of the underlying linear receptive fields, using $n = 2$ in Eq. 4 reduces the bandwidth by a factor of 3/4 and using $n = 3$ reduces the bandwidth by a factor of 5/8. Flanking suppression also reduces the bandwidth of the overt responses, particularly for high contrasts (see Fig. 3).

These choices for the bandwidths are not critical for our conclusions. Even so, these values are generally consistent with the range of spatial frequency, temporal frequency, and orientation bandwidths of cat simple cells (Campbell *et al.*, 1968, 1969; Maffei and Fiorentini, 1973; Ikeda and Wright, 1975a and 1975b, Tolhurst and Movshon, 1975; Movshon *et al.*, 1978b; Holub and Morton-Gibson, 1981; Tolhurst and Thompson, 1981; Berardi, Bisti, Cattaneo, Fiorentini and Maffei, 1982; Webster and DeValois, 1985; Jones, Stepnoski and Palmer, 1987; Robson *et al.*, 1988b; Baker, 1990; Saul and Humphrey, 1992).

We simulated the responses of 18 neurons tuned for six orientations (30 degree spacing between preferred orientations) at each of three spatial frequencies (in which the middle of the three frequency bands was tuned for 0.5 c/d and the spacing between bands depended on the bandwidths). We plot simulated responses of one of the neurons from the middle of the three spatial frequency bands. In the nonspecific-suppression model, the simulated neurons were suppressed equally by the entire population (weights set to 1 for all three spatial frequency bands). In this way, the suppression was constant for mask gratings within a broad (depending on the chosen bandwidths) range of spatial frequencies. When the mask spatial frequency was beyond this range, the suppression gradually (again, depending on the chosen bandwidths) decreased.

For the flanking-suppression model, we set the weight corresponding to the middle of the three frequency bands to 0.1 while the other two weights, corresponding to the lowest and highest spatial frequency bands, were set equal to 1. In this way, the simulated neurons were suppressed mainly by the two flanking (higher and lower) spatial frequency bands.

The full set of linear receptive fields were designed so that the summation in the denominator of Eq. 1 would equal the Fourier energy of the stimulus within an annulus of spatial frequencies. In particular, the radial parts of the frequency responses were truncated, raised cosine functions on a logarithmic frequency scale, and the angular parts of the frequency responses were cosines raised to an integer power, i.e.,

$$H(f,\theta) = \begin{cases} \cos^p(\theta)\,\sqrt{\frac{1}{2}[1+\cos[(\pi/b)\,\log_2(f/f_0)]}\ \text{for}\ |\log_2(f/f_0)| < b \\ \\ 0 \qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad \text{otherwise} \end{cases} \tag{7}$$

Here f is radial frequency, θ is angular frequency, p is an integer constant that determines the orientation bandwidth, $p+1$ is the number of orientation bands, b determines the spatial frequency bandwidth, and f_0 is the preferred spatial frequency.

3 Results

3.1 Spatial Frequency Dependent Suppression

Both models were used to simulate the cross-frequency suppression experiment of Bauman and Bonds (1991). Following Bauman and Bonds' experimental design, we simulated responses to pairs of moving sine gratings, a 2 Hz base grating of optimal spatial frequency superimposed on a 3 Hz mask grating of variable spatial frequency, both optimally oriented. Bauman and Bonds chose the contrasts of their stimuli with respect to the contrasts that caused each neuron's response to saturate. Since the model neurons' responses approach saturation asymptotically, we defined the saturation contrast to be that which evoked 97.5% of the maximum attainable response, and we picked the base and mask contrasts with respect to that value. Following Bauman and Bonds, we picked the base contrast to be 37.5% of the contrast that caused our model neurons' responses to saturate, and we used three different contrasts for the mask grating. The middle of the three mask contrasts was equal to the base contrast, and the other two were 10% higher and lower so that the sum of the base and mask contrasts was between 65%–85% of the saturation contrast. We have obtained similar simulation results with all three mask contrasts, so we plot the results only for the middle of the three mask contrasts. Following Bauman and Bonds, the amplitude of the 2 Hz component of the response time-course was used to summarize the responses. We also varied the initial, relative, spatial phases of the two gratings. As shown below, the initial, spatial phase relationship can have a significant effect on the simulated responses.

<u>Figure 1 about here</u>

Figure 1A replots (Bauman and Bonds, 1991) the spatial frequency tuning curve of of a cat simple cell. Figure 1B replots responses of the same neuron to pairs of gratings, as a function of the mask grating's spatial frequency. There is a large effect of superimposing the mask grating when its spatial frequency is slightly above or below the neuron's preferred spatial frequency.

Figures 1C and 1D plot the simulation results for the nonspecific-suppression model with an exponent $n = 2$ and an overt spatial frequency bandwidth of 0.9 octaves (full-width at half-height, chosen by hand so that the simulated responses would appear similar to the physiological data). For a purely linear neuron, there would be no suppressive effect at all; superimposing a 3 Hz mask grating on a 2 Hz base grating would have no effect on the 2 Hz component of the response. The shape of the curve in Fig. 1D arises from a combination of the nonlinear operations, i.e., rectification, squaring, and normalization.

First, we will consider the effect of half-squaring, ignoring the normalization for a moment. Half-squaring gives rise to a 2 Hz distortion product, cross-talk between the two (2 Hz and 3 Hz) components of the response time-course. Even though the mask gratings have the wrong temporal frequency (3 Hz instead of 2 Hz), superimposing a mask grating of optimal spatial frequency enhances the 2 Hz component of the response. This enhancement of the response depends very little on the initial, relative, spatial phases of the two component sinusoids, as can be seen by the complete overlap of the different curves in Fig. 1D.

Now, we will consider the combined effects of half-squaring and normalization. In the nonspecific-suppression model, the normalization pool is very broadly tuned for spatial frequency. The "W" shape in Fig. 1D therefore arises from a broad "U" shaped spatial frequency suppression provided by the normalization, combined with enhancement near the optimal spatial frequency due to the distortion product provided by half-squaring.

Figure 2 about here

Figure 2 shows simulations for both the nonspecific- and flanking-suppression models with exponents ranging from 1 to 3. Figure 2A, for example, shows the simulated responses of the nonspecific suppression model with halfwave-rectification ($n = 1$). Halfwave-rectification also gives rise to a 2 Hz distortion product, but it is phase dependent as can be seen by the differences between the curves in Fig. 2A. For some initial, relative, spatial phases of the base and mask gratings, the distortion product is positive, i.e., the 2 Hz component of the response is enhanced by superimposing the 3 Hz mask. But for other initial phases, the distortion product is negative, i.e., the 2 Hz component of the response is suppressed by superimposing the 3 Hz mask. Even for initial phases that provide the greatest enhancement (squares in Fig 2A), however, the increase in the response is not large enough to account for the "W" shape in the experimental measurements.

Figure 2B shows the simulated responses of the flanking-suppression model with $n = 1$. The "W" shape, largely consistent with the physiological data, is mainly due to the flanking suppression. There are three weights corresponding the three spatial frequency bands that contribute to the normalization pool. For the simulation results in Figs. 2B, 2D, and 2F, the weights on the flanking frequency bands were set to one, and the weight on the center frequency band was set to 0.1 (although center band weights in the range 0.05 - 0.2 all produced similar results).

The simulated responses in several panels of Fig. 2 are qualitatively consistent with the physiological data. In particular, the nonspecific-suppression model produced reasonable results with $n = 2$ (Fig. 2C) and with $n = 3$ (Fig. 2E).

With $n = 2$ and with the base and mask gratings having equal contrasts, the nonspecific-suppression model predicts 0.5 as a lower limit for the relative suppression, producing a "W" shape is perhaps not quite deep enough (compare with Fig. 1A). A higher exponent ($n = 2.5$ or 3), however, produces a sufficient amount of suppression. The flanking-suppression model produced reasonable results with $n = 1$ (Fig. 2B). For larger values of n (Figs. 2D and 2F), flanking suppression produces too much suppression (the "W" is too deep).

Our simulations also demonstrate that the rectification and accelerating nonlinearities (e.g., halfwave-rectification or half-squaring) can give rise to substantial distortions which can result in suppression or enhancement of the response amplitudes, depending on the relative frequencies and the initial, relative, spatial phases of the stimulus components. Bonds (1989) found no phase dependence when using 2 Hz and 3 Hz gratings. But using different combinations of temporal frequencies, DeValois and Tootell (1983) and Pollen, Gaska and Jacobson (1988) did find phase-dependent suppression and enhancement of simple cell responses. These results were interpreted as "phase-dependent inhibition" by DeValois and Tootell. Pollen *et al.* (1988) later argued that rectification alone might be responsible for these effects. For the particular combination of temporal frequencies (2 Hz and 3 Hz) used for the simulations in Fig. 2, the distortion produced by an exponent of $n = 2$ was mostly phase-independent (Figs. 2C and 2D), whereas exponents other than two yielded phase-dependent distortions (other panels in Fig. 2). With different combinations of temporal frequencies the accelerating nonlinearities produced different patterns of phase-dependence. For example, in other simulations (not shown in the figures) with 1 Hz and 3 Hz gratings, an exponent of $n = 1$ produced phase-independent distortions while exponents of $n = 2$ and 3 produced phase-dependent distortions.

In summary, it appears that both models can account for the "W" shaped curves, but for different reasons. In the nonspecific-suppression model, both the rectification and the normalization are critical for the "W" shape. In the flanking-suppression model, the "W" arises mainly because of the flanking suppression. However, other experimental results favor the nonspecific-suppression model, as we discuss next.

3.2 Spatial Frequency Bandwidth Versus Contrast

A critical failure of the flanking-suppression model is the fact spatial frequency bandwidths are invariant to changes in contrast (Albrecht and Hamilton, 1982; Skottun *et al.*, 1987). The nonspecific-suppression model is perfectly consistent with this result, and it was one of the main motivations for proposing the normalization model (Heeger, 1992a). The reason for this behavior of the nonspecific-suppression model can be understood by considering the equations for the response amplitude, Eqs. 2 and 4. These equations express the response as the product of two factors, one that depends only on stimulus contrast and the other that depends on stimulus spatiotemporal frequency and orientation. Changing the contrast affects only the first factor, thereby scaling up/down the responses by the same amount for all frequencies and orientations.

Figure 3 about here

Figure 3 shows the bandwidths of simulated spatial frequency tuning curves as a function of contrast. For all panels, the bandwidths of the underlying linear receptive fields were 2 octaves (full-width at half-height). The overt responses of the simulated neurons exhibited somewhat narrower bandwidths due to the various nonlinearities, as discussed in *Detailed Methods*. For the nonspecific-suppression model (Figs. 3A, 3C, and 3E) the bandwidths are independent of stimulus contrast, as expected.

For the flanking-suppression model (Figs. 3B, 3D, and 3F), however, spatial frequency bandwidth *does* depend on stimulus contrast. The reason for this behavior can be understood by considering Eq. 6. When the weight $w(f)$ in the denominator of Eq. 6 is equal to 1 (i.e., for flanking frequencies), the contrast dependence of the response is given by:

$$\left(\frac{c^2}{\sigma^2 + c^2}\right)^{\frac{n}{2}} .$$

However, when the weight is much less than one (i.e., near the optimal spatial frequency), the contrast dependence of the response is given by:

$$\left(\frac{c^2}{\sigma^2 + w^2 c^2}\right)^{\frac{n}{2}} ,$$

Thus, changing the contrast produces a larger perturbation in the response for frequencies that are closer to the optimal spatial frequency.

4 Discussion

Two different models, nonspecific-suppression and flanking-suppression can account (qualitatively) for the *apparent* frequency-specific suppression in cat simple cell responses. Other experimental results favor the nonspecific-suppression model. The most critical failure of the flanking-suppression model is that it predicts, contrary to experimental results (see *Results* or *Introduction* for citations), that spatial frequency bandwidths should vary systematically with stimulus contrast.

It is commonly believed that information about a visual stimulus, other than its contrast, is represented in terms of the relative responses of collections of neurons. For example, the spatial frequency of a stimulus might be represented with the relative responses of several simple cells tuned for different preferred spatial frequencies. For this view to be correct, spatial frequency bandwidths must be invariant with respect to stimulus contrast. Nonspecific-suppression makes it possible for response ratios to be independent of stimulus contrast, even in the face of response saturation. But flanking suppression fails to provide such invariance.

We conclude that the suppression is probably broadly tuned for spatial frequency and that the *apparent* flanking suppression is actually due to distortions introduced by an accelerating output nonlinearity with an exponent of $n = 2$ or more. It is possible that a modest amount of flanking suppression (i.e., with a central weight in our simulations slightly smaller than the flanking weights) could contribute. But it is also possible that the suppression could be strongest at the *preferred* spatial frequency (i.e., with a central weight in our simulations slightly larger than the flanking weights), as long as the exponent in the output nonlinearity was larger (e.g., $n = 3$ or more).

Quantitative fits of the data are needed to discriminate between these subtle differences. Unfortunately, the currently published data sets are probably not sufficient. To constrain the fits, one would need to do a series of measurements all with the same neuron: (1) responses to moving gratings as a function of contrast and spatial frequency, and (2) responses to pairs of gratings, varying the relative spatial frequencies, the relative temporal frequencies, and the initial, relative, spatial phases of the two component gratings. Altogether these measurements would over-constrain the model parameters: semi-saturation constant, exponent, spatial frequency tuning of the underlying linear receptive fields, and spatial frequency tuning of the normalization pool.

For example, the slope of the response versus contrast curve, the depth of the "W" shaped spatial frequency suppression curve, and the phase-dependence of the enhancement/suppression all depend on the exponent parameter.

Acknowledgement

This research was supported by a fellowship from CSIC to ON, by a grant from CICYT (TIC94-0849), and by a grant from the National Institute of Health (MH50228) and an Alfred P. Sloan Research Fellowship to DJH.

References

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

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

Baker, C. L. (1990). Spatial- and temporal-frequency selectivity as a basis for velocity preference in cat striate cortex neurons. *Visual Neuroscience*, *4*, 101–113.

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

Berardi, N., Bisti, S., Cattaneo, A., Fiorentini, A., and Maffei, L. (1982). Correlation between the preferred orientation and spatial frequency of neurones in visual areas 17 and 18 of the cat. *Journal of Physiology (London)*, *323*, 603–618.

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.

Campbell, F. W., Cleland, B. G., Cooper, G. F., and Enroth-Cugell, C. (1968). The angular selectivity of visual cortical cells to moving gratings. *Journal of Physiology (London)*, *198*, 237–250.

Campbell, F. W., Cooper, G. F., and Enroth-Cugell, C. (1969). The spatial selectivity of visual cells of the cat. *Journal of Physiology (London)*, *203*, 223–235.

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

Carandini, M., Heeger, D. J., and Movshon, J. A. (1996). Linearity and gain control in v1 simple cells. In Jones, E. G. and Ulinski, P. S., editors, *Cerebral Cortex, Vol. XII. Cortical Models*, in press. Plenum Press, New York.

DeAngelis, G. C., Robson, J. G., Ohzawa, I., and Freeman, R. D. (1992). The organization of suppression in receptive fields of neurons in the cat's visual cortex. *Journal of Neurophysiology*, *68*, 144–163.

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

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

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

Heeger, D. J. (1991). Nonlinear model of neural responses in cat visual cortex. In Landy, M. and Movshon, J. A., editors, *Computational Models of Visual Processing*, pages 119–133. MIT Press, Cambridge, MA.

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

Heeger, D. J. (1992b). 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.

Holub, R. A. and Morton-Gibson, M. (1981). Response of visual cortical neurons of the cat to moving sinusoidal gratings: Response-contrast functions and spatiotemporal interactions. *Journal of Neurophysiology*, *46*, 1244–1259.

Ikeda, H. and Wright, M. J. (1975a). Spatial and temporal properties of 'sustained' and 'transient' neurones in area 17 of the cat's visual cortex. *Experimental Brain Research*, *22*, 363–383.

Ikeda, H. and Wright, M. J. (1975b). Retinotopic distribution, visual latency and orientation tuning of 'sustained' and 'transient' cortical neurones in area 17 of the cat. *Experimental Brain Research*, *22*, 385–398.

Jones, J. P., Stepnoski, A., and Palmer, L. A. (1987). The two-dimensional spectral structure of simple receptive fields in cat striate cortex. *Journal of Neurophysiology*, *58*, 1212–1232.

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

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

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

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

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

Robson, J. G. (1988a). Linear and nonlinear operations in the visual system. *Investigative Opthalmology and Visual Science Supplement*, *29*, 117.

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

Saul, A. B. and Humphrey, A. L. (1992). Temporal frequency tuning of direction selectivity in cat visual cortex. *Visual Neuroscience*, *8*, 365–372.

Sclar, G., Maunsell, J. H. R., and Lennie, P. (1990). Coding of image contrast in central visual pathways of the macaque monkey. *Vision Research*, *30*, 1–10.

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

Tolhurst, D. J. and Heeger, D. J. (1996a). Contrast normalization and a linear model for the directional selectivity of simple cells in cat striate cortex. *Visual Neuroscience*, in press.

Tolhurst, D. J. and Heeger, D. J. (1996b). Comparison of contrast normalization and hard threshold models of the responses of simple cells in cat striate cortex. *Visual Neuroscience*, in press.

Tolhurst, D. J. and Movshon, J. A. (1975). Spatial and temporal contrast sensitivity of striate cortical neurons. *Nature*, *257*, 674–675.

Tolhurst, D. J. and 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 of London, B*, *213*, 183–199.

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

Figure Captions

Fig. 1: Cross-frequency suppression. A: Spatial frequency tuning curve of a cat simple cell (data replotted from Bauman and Bonds, 1991). B: Responses of the same neuron to pairs of moving sine gratings, a 2 Hz base grating of optimal spatial frequency superimposed on a 3 Hz mask grating of variable spatial frequency. Base and mask gratings had the same contrast. Dashed line indicates response to base grating alone. C and D: Simulated responses for the nonspecific-suppression model with n=2 (Eqs. 1 and 2). Spatial frequency bandwidth of the underlying linear receptive fields were chosen by hand to be 1.2 octaves so that the bandwidth of the overt responses was 0.9 octaves (full-width at half-height). The different curves (virtually superimposed) in D correspond to different initial, relative, spatial phases between the base and mask gratings: 0° (circles), 45° (squares), 90° (triangles), 135° (diamonds).

Fig. 2: Cross-frequency suppression simulations for a variety of model parameters. A, C, and E: Nonspecific suppression with exponents of $n = 1$, 2, and 3, respectively (Eqs. 3 and 4). B, D, and F: Flanking suppression with exponents of $n = 1$, 2, and 3, respectively (Eqs. 5 and 6). For all panels, the spatial frequency bandwidths of the underlying linear receptive fields were 2 octaves (full-width at half-height). The bandwidths of the overt responses are plotted in Fig. 3. The different curves in each panel correspond to different initial, relative, spatial phases between the base and mask gratings: 0° (circles), 45° (squares), 90° (triangles), 135° (diamonds).

Fig. 3: Simulated spatial frequency bandwidths (full-width at half-height) for a variety of model parameters (same format as Fig. 2). For all panels, the spatial frequency bandwidths of the underlying linear receptive fields were 2 octaves (full-width at half-height). Overt responses of the simulated neurons

exhibit somewhat narrower bandwidths due to the various nonlinearities, as discussed in *Detailed Methods*. A, C, and E: For the nonspecific-suppression model the bandwidths are independent of stimulus contrast. B, D, and F: For the flanking-suppression model, the bandwidths depend on stimulus contrast inconsistent with physiological data.

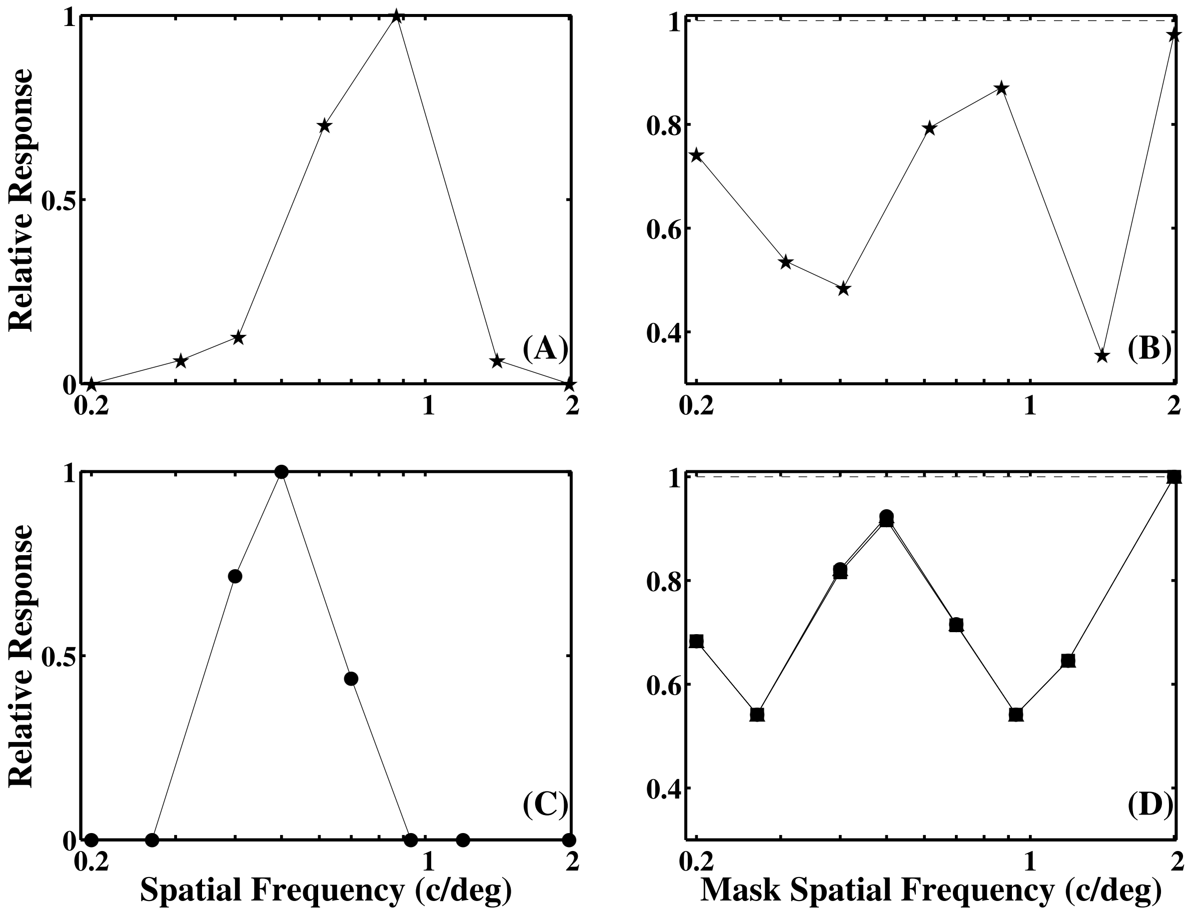

Relative Response
Relative Response
(A)
(B)
(C)
(D)
Spatial Frequency (c/deg)
Mask Spatial Frequency (c/deg)

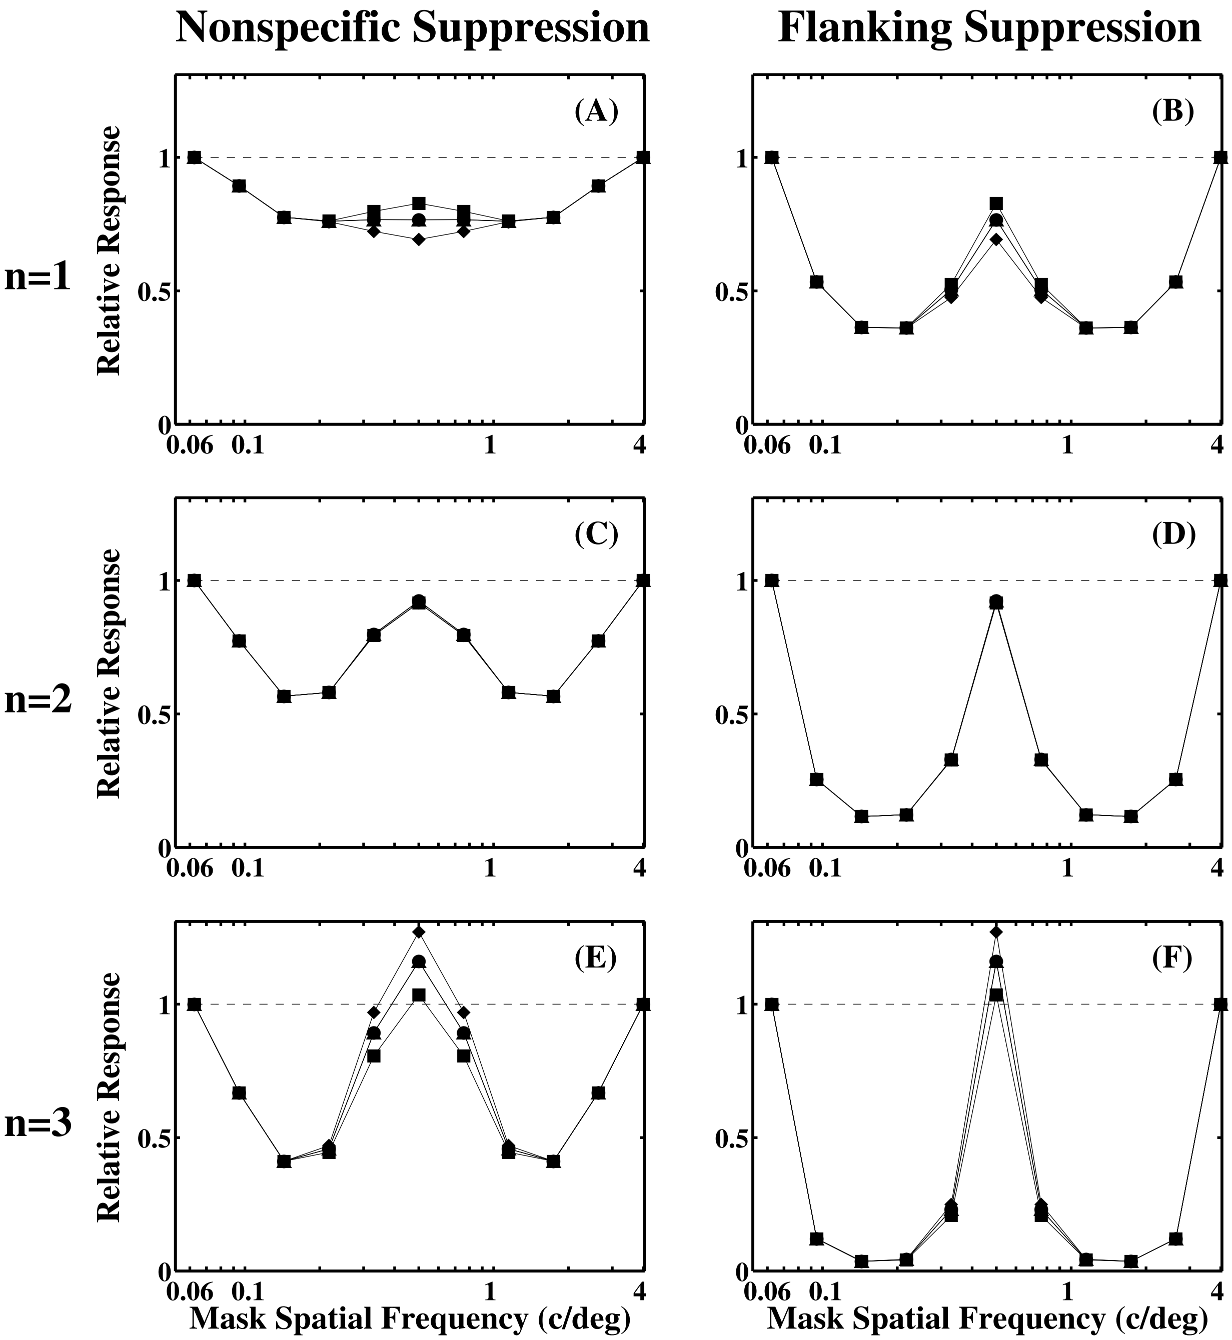
Nonspecific Suppression
Flanking Suppression
n=1
n=2
n=3
(A)
(B)
(C)
(D)
(E)
(F)
Relative Response
0
0.5
1
0.06
0.1
1
4
Mask Spatial Frequency (c/deg)

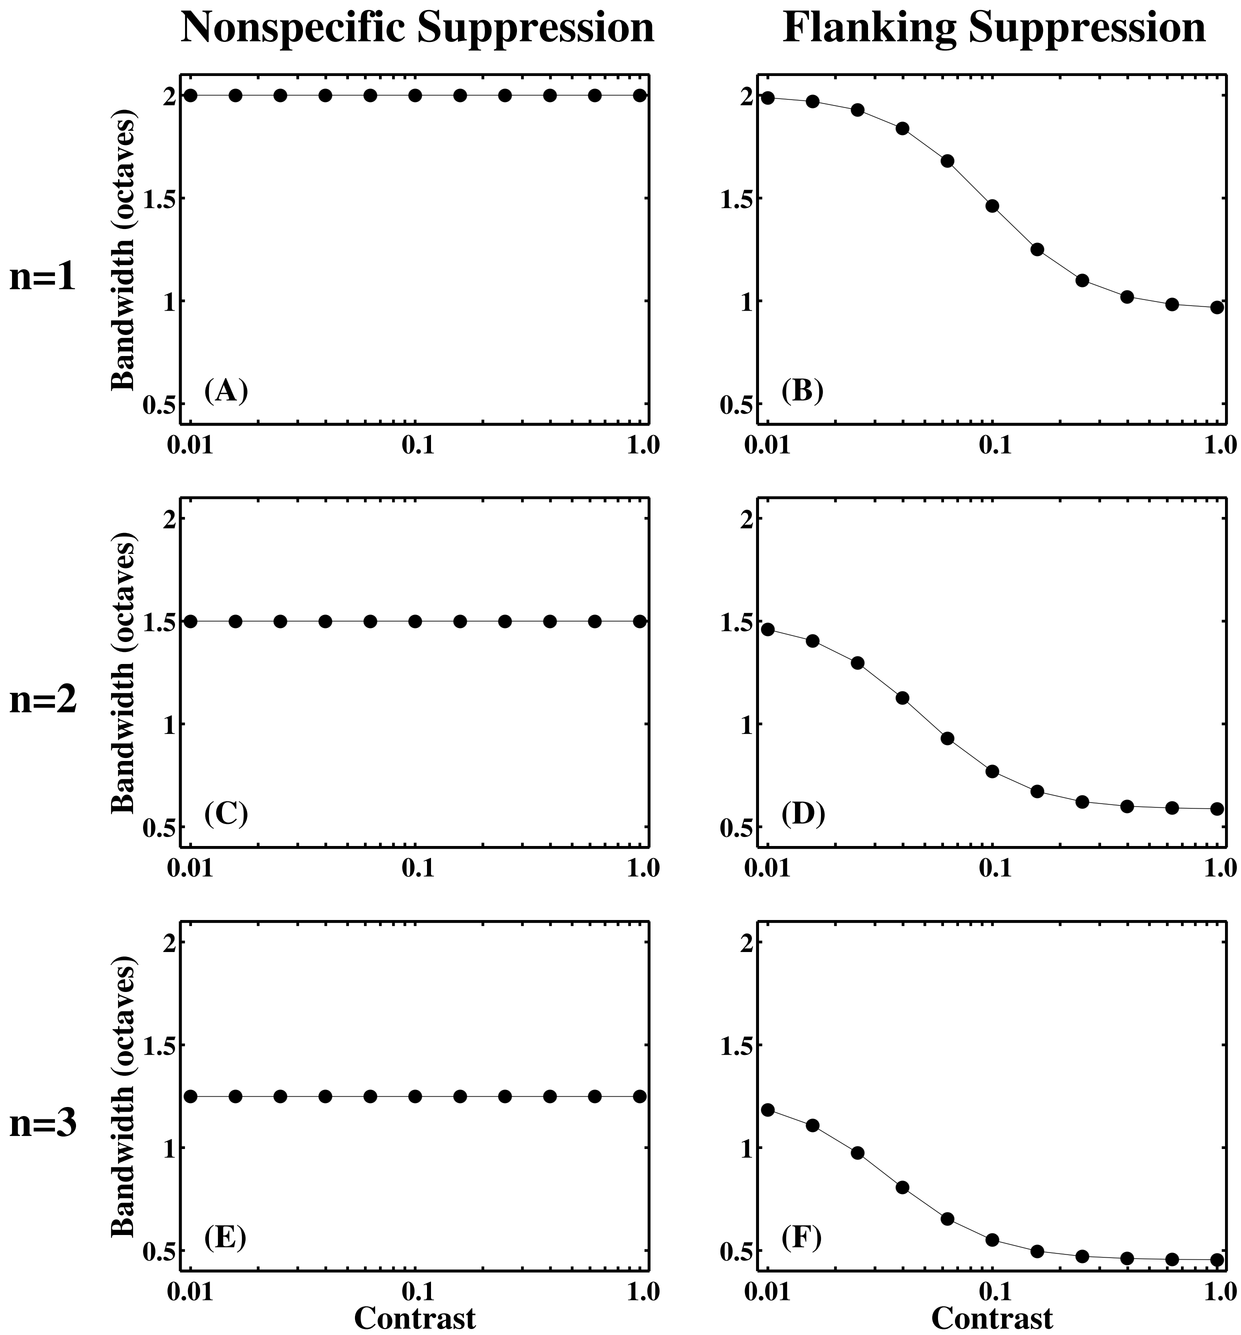
Nonspecific Suppression
Flanking Suppression
n=1
n=2
n=3
Bandwidth (octaves)
(A)
(B)
(C)
(D)
(E)
(F)
0.01
0.1
1.0
0.5
1
1.5
2
Contrast

Fig. 1.

Fig. 2.

Fig. 3.