LINEARITY AND GAIN CONTROL IN V1 SIMPLE CELLS

Matteo Carandini*
Department of Neurobiology and Physiology, Northwestern University

David J. Heeger
Department of Psychology, Stanford University

J. Anthony Movshon
Howard Hughes Medical Institute and Center for Neural Science,
New York University

September 4, 1996

Submitted as a Chapter for "CEREBRAL CORTEX. vol. XII. CORTICAL MODELS", edited by E G Jones and P S Ulinski. New York: Plenum Press.

*To whom proofs should be mailed at: Department of Neurobiology and Physiology, Northwestern University, 2153 North Campus Drive, Evanston, IL 60208. Email: matteo@nwu.edu. Phone: (847) 467 1443. Fax: (847) 491 5211.

Contents

1 Introduction

The primary visual cortex (V1) is arguably the most studied area in the mammalian cortex, and one of the very few for which we can say something sensible about the computations that it performs. V1 cells are selective for the position, shape, size, velocity, color, and eye of presentation of a visual stimulus. The mechanism of this selectivity, as well as its rationale, have recently begun to be understood, although some aspects still constitute an area of intense debate.

The receptive fields of V1 cells were first mapped by Hubel and Wiesel (1962) using flashing bars. They termed *simple cells* those cells for which they could find regions that responded either to the onset or to the offset of a bright bar, but not to both. Simple cells constitute around 50% of V1 neurons (De Valois et al., 1982; Schiller et al., 1976). This Chapter is devoted to simple cells, but it also includes ideas that can be useful in the understanding of the other major V1 cell type, the complex cells.

** Figure 1 About Here **

A number of researchers have proposed that simple cells behave linearly, and that their selectivity is determined by their linear weighting functions (Figure 1A). This *linear model* of simple cells is attractive because if it were correct it would be possible to predict the responses of a simple cell to any visual stimulus, based on a limited number of measurements. For example, any image can be approximated by a number of small pixels. Measuring the cell response by lighting each pixel one by one would allow the experimenter to predict the response to any visual stimulus.

We begin the Chapter with the full definition of the linear model (Section 2). We explain its basic properties, and we summarize the vast number of studies that were devoted to testing it (Section 3). In these studies the model was found to be largely successful in explaining the selectivity of simple cells for stimulus shape, size, position, orientation and direction of motion. We then propose a biophysical implementation of the linear model, and we discuss

its plausibility (Section 4). In this implementation, simple cells receive both excitation and inhibition arranged in push-pull, so that when one increases the other decreases, and vice versa. The importance of this arrangement is that it makes it possible for the visual stimuli to result in perfect current injection in the cell, without any conductance increase. Conductance increases would result in nonlinear behavior.

While many aspects of simple cell responses are consistent with the linear model, there also are important violations of linearity (Section 5). For example, scaling the contrast of a stimulus would identically scale the responses of a linear cell. At high contrasts, however, the responses of simple cells show clear saturation (Maffei and Fiorentini, 1973). Moreover, simple cells are subject to cross-orientation inhibition: the responses to an optimally-oriented stimulus can be diminished by superimposing an orthogonal stimulus, which would be ineffective in driving the cell when presented alone (Morrone et al., 1982). These nonlinearities may be partially explained by contrast gain control mechanisms known to operate as early as in the retina (Shapley and Victor, 1978). There is however evidence suggesting that these nonlinearities have an important cortical component. As a result, some researchers have argued that selectivity of simple cells must be due to nonlinear mechanisms (for review see Ferster and Koch, 1987).

Our opinion is that the linear model is a very powerful explanation of the behavior of simple cells. Part of the nonlinear behavior of simple cells can be attributed to the rectification (threshold) in the generation of action potentials. The many nonlinearities that are not accounted for by rectification can be explained by adding to the linear model a divisive inhibition stage (Figure 1B). This stage controls the *gain* (responsiveness) of the neurons by means of intracortical feedback from a large group of other cortical cells. This extended linear model is called the *normalization* model. Response normalization was originally proposed by Robson (1988) to provide explanations for various failures of the linear model of simple cell responses. The model has been expanded and formalized by Heeger (1993, 1992a,b, 1991), by Albrecht and Geisler (1991), and by Carandini and Heeger (1994). who have shown that the normalization model is capable, in principle, of explaining a wide variety of empirical phenomena (see also Bonds, 1989; DeAngelis *et al.*, 1992; Tolhurst and Heeger, 1996a,b; Nestares and Heeger,

1996). 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).

We propose (in Section 6) that normalization operates by shunting inhibition: cells inhibit each other by increasing each other's membrane conductance. This decreases the gain of the transformation of input currents into membrane potentials. The model explains response saturation and cross-orientation inhibition because increasing stimulus contrast or adding an orthogonal grating increases the activity of the network, resulting in increased conductance and decreased gain. Increasing the conductance also decreases the time constant of the membrane so the latency of the responses and the temporal filtering properties of the cells depend on the stimulus contrast. As a consequence the model captures a number of temporal nonlinearities in the responses of V1 cells (Albrecht, 1995; Hawken et al., 1992; Reid et al., 1992; Dean and Tolhurst, 1986; Holub and Morton-Gibson, 1981).

The normalization model is intentionally based on a very simplified view of the cellular physiology. As a consequence, it makes strong quantitative predictions with very few free parameters, some of which we have rigorously tested (Section 7). We recorded from simple cells in the primary visual cortex of paralyzed, anesthetized macaques, while presenting very large sets of visual stimuli. We derived closed-form equations for the model responses to such stimuli, and we found that these equations provide good fits to the neural responses.

We conclude the Chapter with a discussion of the biophysical plausibility of the normalization model (Section 8). Is shunting inhibition really the mechanism underlying gain control in the cortex? We tested the model using extracellular data, so we have no direct proof that the overall conductance grows with stimulus contrast. There are actually reasons to doubt that this is the case: intracellular *in vivo* measurements have consistently failed to demonstrate large conductance increases related to visual stimulation (Ferster and Jagadeesh, 1992; Berman et al., 1991). As a result, the true biophysical substrate of gain control is uncertain. The main advantage of shunting inhibition is that it constitutes the simplest possible way for the normalization pool to control both the gain and the dynamics of a cell's response. Until further

data are available the model should be considered to lie between a phenomenological model and a correct biophysical explanation.

2 The Linear Model of Simple Cells

2.1 Visual Stimuli In Space-Time

** Figure 2 About Here **

As a visual stimulus is projected on the retina it can be described by its intensity distribution $I(x, y, t)$, that varies in the two spatial dimensions x,y and in time t. This representation ignores the color of the stimulus and assumes monocular viewing, but is in all other respects complete. Consider for example a stimulus consisting of a dark vertical bar drifting from left to right, on a white background. Figure 2A shows the bar at a particular instant in time. Panel B shows that as the bar drifts from left to right, it can be considered as a solid in the x-y-t space. Panel C shows a snapshot of the volume taken from above, a space-time (x-t) plot which ignores the y dimension.

Different velocities result in different orientations in space-time (Adelson and Bergen, 1985; van Santen and Sperling, 1985; Watson and Ahumada, 1985; Fahle and Poggio, 1981). For example, if the bar in Figure 2A were going faster, its space-time (x-t) representation (Panel C) would have been more tilted towards the horizontal. Had the bar been motionless, its x-t representation would have been vertical. Had the bar been going from right to left, the orientation of its x-t representation would have been opposite to the one in Figure 2C.

2.2 Spatiotemporal Weighting Functions

A defining property of linearity is that of *superposition*: if L_1 is the response to stimulus I_1, and L_2 is the response to stimulus I_2, then the response of a linear system to the sum of the stimuli $I_1 + I_2$ is just the sum of the responses, $L_1 + L_2$. While the property of superposition

may sound a little abstract, there is an equivalent statement that will make it concrete: simple cells are linear if and only if their responses are *a weighted sum of the light intensities* falling on their receptive fields.

Figure 2D-F shows a schematic of a spatiotemporal weighting function. Panel D shows a space-space (x-y) section of the weighting function. Panel F shows the space-time (x-t) projection of the weighting function, i.e. a snapshot of the weighting function taken from above. The relation between the two is shown in Panel E. As we will see, real simple cell weighting functions don't look too different from this idealization (Figure 7).

The response of a linear cell is simply obtained by weighting the stimulus intensity I at each location and time by the value of the cell's weighting function W at that location and at that time, and by summing the results:

$$L(t) = \int\int\int W(x, y, T)\ I(x, y, t - T)\ dx\ dy\ dT. \tag{1}$$

The cell travels in time t from past to future (as we all do), while its retinal (x-y) location remains fixed. The weighting function is zero for any time that lies in the future, because the responses of the cell cannot depend on future events.

The spatiotemporal weighting function W of a linear cell determines its selectivity (e.g., for orientation or direction of motion). In particular, several researchers have pointed out that a linear cell is direction selective if and only if the subregions of its weighting function are tilted along an oblique axis in space-time (Adelson and Bergen, 1985; van Santen and Sperling, 1985; Watson and Ahumada, 1985; Fahle and Poggio, 1981).

** Figure 3 About Here **

Figure 3 illustrates how this selectivity arises, by showing the responses of a linear cell with a space-time oriented weighting function to a drifting grating stimulus. In Panels A-D the grating drifts from left to right, and the resulting space-time orientation is very similar to the space-time orientation of the weighting function. This results in strong responses (Panel A). In

Panels E-H the grating drifts in the opposite direction, and the resulting space-time orientation is almost orthogonal to that of the weighting function. This results in very small responses (Panel E) because the weighting function averages out the variations in intensity present in the stimulus.

** Figure 4 About Here **

If a linear cell's weighting function is not tilted along an oblique axis in space-time, then the cell will not have a preference for direction of motion. Figure 4, for example, shows the space-time projections of three different weighting functions. The weighting function in Panel A prefers stationary objects, since it is vertical in the x-t plane. The weighting function in Panel B prefers moving or flickering objects but has no preference for the direction of motion. These two weighting functions cannot be direction selective because they are *space-time separable*, i.e. their weighting functions $W(x, y, t)$ can be expressed as the product of a function of space x, y and a function of time t. The weighting function in Panel C is clearly tilted along an oblique axis in space-time, and is direction selective. Note that it is not space-time separable.

2.3 A Nonlinearity: Light Adaptation

In characterizing simple cells as spatiotemporal linear neurons, we have neglected an important (retinal) nonlinearity: light adaptation (Shapley and Enroth-Cugell, 1984). We can, however, safely ignore light adaptation by restricting our choice of visual stimuli to luminance distributions $l(x, y, t)$ that modulate (transiently) about a fixed mean/background luminance $\bar{l}$. Examples are drifting grating patterns and drifting or briefly flashed bars that are either brighter or darker than the mean. In these conditions the retina can be considered to be in a fixed state of adaptation, and its output is proportional to the "local contrast" $I(x, y, t) = [l(x, y, t) - \bar{l}]/\bar{l}$ of the stimulus (Shapley and Enroth-Cugell, 1984).

To avoid confusion we (improperly) use the term *intensity* to refer to the local contrast, and we reserve the term *contrast* for the maximum absolute value of the local contrast of a

grating stimulus. The maximum contrast of a grating is 1, which is attained when the lowest intensity is zero and the highest intensity is twice the mean. Finally, we use the term *local energy* to denote the variance of the stimulus over local space and recent time, and within a band of spatiotemporal frequencies.

2.4 Another Nonlinearity: Rectification

Spatiotemporal linear weighting functions are intended to be models the *intracellular* responses of simple cells. Most of the data discussed in this Chapter, however, were obtained *extracellularly*. To model this, one is forced to consider also the transformation of membrane potentials into firing rates.

This transformation is bound to introduce a nonlinearity. The responses of a linear cell would assume both positive and negative values. Likewise, the membrane potential fluctuates above and below a cell's resting potential. Firing rates, on the other hand, are by definition positive. A linear cell with a high maintained firing rate could encode the positive and negative values by responding either more or less than the maintained rate. This is for example typical of retinal ganglion cells (Enroth-Cugell and Robson, 1966). Simple cells, however, have very little maintained discharge. Since their negative responses cannot be encoded in their firing rate, simple cells cannot act truly linearly.

** Figure 5 About Here **

As we discuss in Section 2.4, the transformation of membrane potential into spike rate can be approximated by *rectification*, that is by a function that is zero for membrane potentials below a threshold, and that grows linearly from there on. Figure 5 shows some examples of rectification. The three solid lines depict cases in which the firing threshold is respectively 0, 5 and 10 mV away from the resting potential. Technically the first example of rectification, with a threshold at V_{rest}, is called "half-rectification". The other two are called "over-rectification", since their threshold is above the resting potential. We use the term "rectification" to include

all these cases.

Rectification is a static nonlinearity, that is one that depends only on the instantaneous value of its input and not on its past history. Adding rectification after a spatiotemporal linear weighting function does not substantially alter the selectivity or other basic properties of the responses (Heeger, 1992a). In the following, whenever we refer to the *linear model* of simple cells, we tacitly assume it to be *a spatiotemporal linear weighting function followed by rectification*, as shown in Figure 1A.

3 Some Linear Properties of Simple Cells

This Section describes some experimental results that provide strong evidence in favor of the linear model of simple cells (see Heeger, 1993,1992a, for a more thorough review). Most of the nonlinearities that are mentioned in this Section are explained by the rectification stage that transforms intracellular responses into firing rates.

3.1 Responses to Impulses

** Figure 6 About Here **

When Hubel and Wiesel (1962) first mapped the receptive fields of V1 cells, their stimuli were bright flashing bars. Depending on whether a region responded positively to the onset or to the offset of a bright bar, they termed that region an ON or an OFF subregion. The linear model predicts the existence of these subregions (Heeger, 1992a; Emerson, 1988). This can be understood by considering the full space-time representation of a flashed bar, which is shown in Figure 6A. Since the bar does not change position in time, its space-time (x-t) projection is vertical. The top and bottom ends of the rectangle are respectively the times at which the bar is turned on and off. The responses of a linear weighting function to such a stimulus are depicted in Figure 6B-C. Figure 6B shows the case in which the bar is flashed in an OFF subregion. As the weighting function travels down in time, the first subregion of the

weighting function that hits the stimulus is inhibitory; this gives a negative response. Later the stimulus overlaps both the excitatory and inhibitory subregions, so the response is about zero. Finally, when the stimulus overlaps only the excitatory subregion (right panel of Figure 6B) the weighting function gives a positive response. In sum, the response is negative just after the bar is turned on and positive just after it is turned off. The opposite will happen when the bar is flashed in and ON subregion (Figure 6C): the response is positive just after the bar is turned on and negative just after it is turned off. After the responses in the left panels of Figure 6B and C are passed through a rectification stage that shows only their positive parts (shaded areas), they closely resemble the spike rate responses of a real simple cell.

** Figure 7 About Here **

Since Hubel and Wiesel's original work, the method for mapping a receptive field has been made more quantitative by having a computer show sequences of bars in many different positions and recording the correlation between firing rate and light intensity. Such a correlation depends on space and time: for each location x, y and *delay time* T, one can measure the correlation between the spike train $R(t)$ and the sequence of stimulus intensities that occurred T seconds before at that x, y location, $I(x, y, t - T)$. The value of the correlation, which can be positive or negative, is taken as the strength of the weighting function at that position and time, $W(x, y, T)$. This method is called *reverse correlation*, and allows the measurement of full space-time (x-y-t) weighting functions (McLean et al., 1994; DeAngelis et al., 1993a,b; Shapley et al., 1991; McLean and Palmer, 1989; deBoer and Kuyper, 1968). Figure 7 shows the full space-time weighting function of a simple cell, measured with the reverse correlation technique. The four upper panels represent x-y snapshots of the weighting function measured at different times T in the past. A large number of snapshots like these are stacked to build a full space-time weighting function, whose x-t structure (averaged over the y axis) is shown at the bottom of the figure.

The reverse correlation method can be applied to any visual cell, linear or nonlinear, and it will always give a result, i.e. a full space-time weighting function. For a linear cell, however,

such a weighting function could then be used to predict the cell's responses to any visual stimulus (using Equation 1). This property of linear systems can be used to test whether simple cells are linear. For example, one can ask whether direction selectivity in simple cells is fully explained by an underlying linear stage.

The linear model predicts that simple cells are direction selective only if their weighting functions are oriented in space-time, and thus nonseparable (Figure 4). This prediction was tested by McLean and Palmer (1994, 1989), Shapley, Reid and Soodak (1991), Emerson and Citron (1992) and DeAngelis et al. (1993b). Their findings are mostly consistent with the linear model. They found simple cells with weighting functions tilted along an oblique axis (inseparable) in space-time, like the one in Figure 7. These cells were all direction selective, and the preferred direction of motion was always correctly predicted from the orientation of the weighting function. For example, the cell of Figure 7 was highly direction selective and preferred stimuli moving from right to left in the x-y plane. A number of simple cells were found to have space-time separable weighting functions, like the ones depicted in Figure 4A and B. Consistent with the linear model, most of these cells were not direction selective*.

3.2 Responses to Drifting Gratings

The stimulus of choice for linear systems analysis of the visual system is a stimulus whose luminance varies sinusoidally in space and time, the *sine grating*. There are many advantages to using sine gratings (reviewed in Enroth-Cugell and Robson, 1984), the most important being that linear systems are guaranteed to respond to sinusoidal modulation with a sinusoid. For example, had the modulation in luminance of the grating in Figure 3 been sinusoidal, the responses in Panels A and E would have been perfect sinusoids. The deviation of the responses from pure sinusoids can provide a quantitative measure of nonlinearity (Hochstein and Shapley, 1976). Sine gratings were first used to study the neurophysiology of the visual system by Enroth-Cugell and Robson (1966), who demonstrated the linearity of cat retinal X ganglion cells. One of their tests of linearity involved comparing the responses to gratings with the responses to luminance edges. The logic of this experiment is straightforward: since an

edge is composed of the sum of a number of gratings, the responses of a linear cell to an edge would be predictable from its response to gratings.

** Figure 8 About Here **

A similar experimental paradigm was applied to the study of simple cells by Movshon et al. (1978a). They measured the sensitivity of the cells to drifting gratings of different spatial frequencies, and the sensitivity of different receptive field regions to flashing bars. Since a grating is composed of the sum of a number of bars, the response of a linear cell to a grating is predictable (via Fourier transform) from its response to bars. Likewise, since a bar can be thought of as the sum of a number of gratings, the response to a bar is predictable (via inverse Fourier transform) from the response to gratings. Movshon et al. found good agreement between the weighting function predicted by inverse Fourier transform of the grating sensitivity and the weighting function obtained from the flashing bar data (Figure 8). This supports the linear model of simple cell responses.

Many other studies have compared grating responses to impulse responses (DeAngelis et al., 1993b; Shapley et al., 1991; Tadmor and Tolhurst, 1989; Jones and Palmer, 1987b; Jones et al., 1987; Jones and Palmer, 1987a; Field and Tolhurst, 1986; Dean and Tolhurst, 1983; Kulikowski and Bishop, 1981b,a; Glezer et al., 1980; Andrews and Pollen, 1979; Maffei et al., 1979).In many cases, the inverse transform of the response to gratings gives a weighting function with additional side bands beyond those measured directly. In addition, the measured response to gratings is often more narrowly tuned for spatial frequency than predicted from the Fourier transform of the response to impulses. This discrepancy between the grating responses and the impulse responses can be explained by over-rectification, which conceals the impulse responses of the weaker receptive field regions, so it is consistent with the linear model (Tolhurst and Heeger, 1996b; Heeger, 1992a; Tadmor and Tolhurst, 1989).

Some of the above mentioned studies, on the other hand, unveiled a serious failure of linearity: a discrepancy between the predicted and actual *sizes* of the responses. For example, when Movshon et al. (1978a) compared the observed and predicted weighting functions (Figure 8),

they did so only up to an arbitrary amplitude scaling factor. This scaling factor should not be necessary according to the linear model. We will see later on that the normalization model predicts this failure of linearity. In particular, it predicts that the cell's gain is different when it is stimulated with flashed bars from when it is stimulated with drifting gratings.

3.3 Responses to Contrast-Modulated Gratings

A contrast-modulated grating is a standing sine grating whose intensity is modulated sinusoidally over time. Simple cell responses to drifting and contrast-modulated gratings are quite similar to rectified sinusoids (see e.g. Figure 12A). This is obviously consistent with the linear model: the spatiotemporal linear weighting function responds with a sinusoid and the rectification hides everything that is below threshold.

** Figure 9 About Here **

A number of researchers (e.g. Reid et al. 1991, 1987; Tolhurst and Dean, 1991; Kulikowski and Bishop, 1981b; Movshon et al., 1978; Maffei and Fiorentini, 1973) measured the responses of simple cells while varying the spatial phase of contrast-modulated gratings. Since these responses can be reasonably fit by a sinusoid, they can be described by just two numbers, the amplitude and phase of the sinusoid. A useful way to display both response amplitude and response phase at the same time is given by a *polar plot* like the one in Figure 9. Every point in the polar plot corresponds to a sinusoid, whose amplitude is given by the distance from the origin, and whose phase is given by the angle with the horizontal axis.

The linear model predicts that as the spatial phase of the contrast-modulated grating varies between 0° and 180°, the responses should describe a "wasp-waisted" ellipse in the polar plot (Movshon et al., 1978a). In particular, for a linear cell a polar plot of the responses would be elliptical in shape. Over-rectification distorts the ellipse, producing a wasp-waist: if the neuron has to reach a certain level of excitation before any activity is seen, there will be a disproportionate decrease in small responses (Tolhurst and Heeger, 1996a; Heeger, 1993;

Albrecht and Geisler, 1991; DeAngelis et al., 1993b). The physiological results are in line with this prediction. Figure 9, for example, shows the responses of a simple cell to a contrast-modulated grating positioned at 8 different phases over the receptive field, spanning the range from 0° to 180°. The raw data (filled symbols) describe a wasp-waisted ellipse since the amplitudes near the minor axes are smaller than they should be to fit an ellipse. When the distortion introduced by rectification is removed (in this case assuming a resting "firing rate" of minus 8 spikes per second), the data fall on an ellipse (open symbols).

Because of superposition, the responses of a linear cell to contrast-modulated gratings would be easily predictable from the responses to drifting gratings, and vice versa. Several researchers tested whether this was the case for simple cells. Ferster and collaborators performed intracellular *in vivo* recordings and found that the membrane potential responses were consistent with the output of a spatiotemporal linear weighting function (Jagadeesh et al., 1993). Other researchers performed extracellular recordings (Albrecht and Geisler, 1991; Tolhurst and Dean, 1991; Reid et al., 1991, 1987). These studies are generally consistent with the linear model in that a cell's preferred direction of motion for drifting gratings can be correctly predicted from its responses to contrast-modulated gratings. They however uncovered two nonlinearities in simple cell responses. First, the linear prediction from contrast-modulated grating responses underestimates the degree of directional selectivity observed with drifting gratings. Second, the linear prediction overestimates the responses to gratings drifting in the nonpreferred direction. Albrecht and Geisler (1991) and Heeger (1993, 1991) showed that the first phenomenon can be explained by the rectification stage (which acts as an expansive nonlinearity), so it is consistent with the linear model of simple cells. The second phenomenon instead is not consistent with the linear model, but it can in most cases be explained by a gain control mechanism like that postulated by the normalization model. The normalization model predicts that the cells are less responsive in the presence of drifting gratings than in the presence of contrast-modulated gratings of equal contrast. This difference in gain can be shown to yield the observed discrepancy in the predicted and actual responses to gratings drifting in the nonpreferred direction (Heeger, 1993; Tolhurst and Heeger, 1996a).

3.4 Responses to Compound Stimuli

According to the linear model of simple cells, knowing a cell's responses to gratings would make it possible to predict the responses to any visual stimulus. This is because any visual stimulus can be expressed as the sum of many different gratings.

An important test of this prediction was performed by DeValois et al. (1979). They measured the orientation tuning of cat simple cells using individual gratings as well as checkerboard stimuli. The motivation for their use of checkerboards is very interesting. At the end of the 1970s the issue of whether simple cells were better modeled as linear weighting functions or as all-or-none edge detectors was the object of heated debate (see e.g. Marr, 1982, Maffei and Fiorentini, 1973 and Schumer and Movshon, 1984). DeValois et al. (1979) reasoned that the two models made very different predictions of a cell's response to checkerboards. In checkerboards the strongest sine grating components are oriented along the diagonals, whereas the sharp edges are oriented along the rows and columns. According to the linear model the cells will respond best when one of the diagonals is oriented in the cell's preferred orientation for gratings. According to the edge-detector model, on the other hand, a cell will respond best when either the rows or the columns are oriented in the cell's preferred orientation for gratings. The results of DeValois et al. (1979) were consistent with the linear model, and falsified the edge-detector model. The responses of the cells could be predicted by having knowledge of the location and orientation of the main sine gratings that compose the checkerboard. The precise location and orientation of the sharp edges was not relevant in predicting the cells' responses.

A similar approach was followed by other researchers (e.g. Gizzi et al., 1990, Pollen et al., 1988, DeValois and Tootell, 1983; Pollen et al., 1982; Maffei et al. 1979), who tested linearity by comparing responses to single sine gratings with responses to sums of sine gratings of different spatial frequencies or orientations. All of these results are qualitatively explained by the linear model (Heeger, 1992a).

The quantitative predictions of the linear model, however, are not always correct, and once again the discrepancy points to the existence of a gain control mechanism. For example, Gizzi

et al. (1990) found that simple cell responses to plaids composed of two sine gratings with different orientations were on average only 2/3 of the linear predictions based on the responses to the individual gratings. As we will see, the normalization model explains this behavior because it predicts that the gain of the cells is lower in the presence of plaids than in the presence of single gratings.

4 Biophysics of the Linear Model

We have been considering the linear model as little more than a mathematical abstraction. This Section describes how the model might be implemented physiologically.

4.1 Linearity of the LGN

Unless some complicated linearization mechanism is invoked, simple cells can only be as linear as the inputs they get. Since the input to the visual cortex is constituted by the activity of LGN cells, we must begin our task of modeling simple cell linearity by assuming that *the responses of LGN neurons are linear functions of the stimulus intensity distribution.*

This assumption is a better approximation in the monkey than in the cat. In the cat there are no LGN cells that are perfectly linear: the X cells are spatially and temporally linear, but they have a (retinal) contrast gain control mechanism, which violates linearity (Victor, 1987; Enroth-Cugell et al., 1983; Enroth-Cugell and Robson, 1966). The Y cells are extremely nonlinear (Victor, 1988; Troy, 1983; Hochstein and Shapley, 1976), but may not contribute any input to the primary visual cortex (Ferster, 1990b,a).

In the monkey, on the other hand, there is a geniculocortical channel, the P pathway, which is substantially linear. The other channel, the M pathway, is instead quite nonlinear, and its nonlinearity might be due to a gain-control mechanism. The substantial linearity of the P pathway and the nonlinearity of the M pathway have been observed in the responses of retinal ganglion cells (Benardete and Kaplan, 1995; Lee et al., 1994; Benardete et al., 1992), and are

reflected in the properties of LGN cells (Movshon et al., 1994; Carandini et al., 1993; Sherman et al., 1984; Derrington and Lennie, 1984). Even though P cells constitute around 90% of the monkey LGN (Dreher et al., 1976), many simple cells also receive M inputs (Malpeli et al., 1981). Indeed, while the two streams are segregated in layer 4C (Blasdel and Lund, 1983; Hendrickson et al., 1978; Hubel and Wiesel, 1972), they are not segregated at all in the upper layers (Nealey and Maunsell, 1994; Yoshioka et al., 1994; Lahica et al., 1992). In particular, for those neurons that do receive M input, the first 7-10 ms of activation are due exclusively to the M signal (Maunsell and Gibson, 1992).

Besides the fact that monkey simple cells may receive some M input, there is another phenomenon that makes our assumption of linearity of the geniculate input an imperfect approximation: at high contrasts the responses of LGN cells show evidence of rectification. Resting firing rates in the LGN have been reported to be around 18 sp/s in monkey, and respectively around 6 sp/s and 16 sp/s in cat X and Y LGN cells (Kaplan et al., 1987). An analysis of the 136 monkey LGN cells recorded by Sherman et al. (1984) reveals that their average resting firing rate was only around 7 spikes/s. On average the modulation in spike rate due to a 50% contrast grating was twice as large as the resting firing rate for P cells (median: 2.1), and three times as large for M cells (median: 3.3). Rectification in the input will be ignored in the following but should be considered in more detailed models of the visual cortex.

4.2 Building Simple Cell Receptive Fields

In their 1962 paper, Hubel and Wiesel hypothesized that the receptive fields of simple cells are the result of an orderly arrangement of LGN inputs. Geniculate cells have center-surround receptive fields. When stimulated with a spot stimulus in their center they respond either to the onset (ON-center cells) or to the offset (OFF-center cells) of the stimulus. According to the scheme proposed by Hubel and Wiesel an ON subregion of a simple cell receptive field would result from the sum of an aligned series of LGN ON-center inputs. Similarly, an OFF subregion would result from the sum of an aligned series of OFF-center inputs. This arrangement has been recently confirmed by Reid and Alonso (1995), who recorded simultaneously from simple

cells and from LGN cells.

If some of the LGN inputs reach a simple cell before some others, the resulting receptive field can also display direction selectivity. In the cat LGN, for instance, there is evidence for the existence of two classes of cells, "lagged" and "nonlagged", whose responses have different latencies (Saul and Humphrey, 1992, 1990; Mastronarde, 1987). Figure 4 illustrates how the outputs of these two classes of cells could be summed to yield a direction selective simple cell. The weighting functions in Panels A and B were drawn to resemble respectively that of idealized lagged and nonlagged LGN cells. These functions are space-time separable and have the same center-surround spatial structure, but slightly different spatial positions. They have different temporal structures, which result in different response latencies. Panel C depicts the weighting function obtained by summing the two LGN weighting functions. This weighting function is oriented in space-time, so it is direction selective. It is similar to that of a direction selective V1 simple cell. In practice, V1 simple cell weighting functions would be the result of many, not just two, LGN inputs. In general, a nonseparable (direction selective) weighting function can be obtained by simple addition of separable (not direction selective) weighting functions (Adelson and Bergen, 1985; Watson and Ahumada, 1985; Fahle and Poggio, 1981), as long as the latter differ in their temporal structure.

4.3 Push-Pull Arrangement of Inputs

The linear combination of LGN inputs involves both sums and subtractions. Indeed, besides excitatory responses, simple cell receptive fields also exhibit inhibitory responses, elicited by the onset of a light on an OFF region, or the offset of a light on an ON region. Hubel and Wiesel (1962) pointed out that these inhibitory responses could originate either from the withdrawal of excitation or from actual inhibition.

There is now evidence that both mechanisms are at work. Inhibitory postsynaptic potentials (IPSPs) do appear in intracellular recordings of simple cells (Creutzfeldt and Ito, 1968), and their interaction with the excitatory postsynaptic potentials (EPSPs) is subtractive (Fer-

ster and Jagadeesh, 1992; Berman et al., 1991). Ferster (1986) measured the selectivity of the IPSPs and found it to be identical to that of the EPSPs. Moreover, he found that in ON regions a light increase results in EPSPs and a light decrease results in IPSPs, while in OFF regions a light decrease results in EPSPs and a light increase results in IPSPs (Ferster, 1988). In other words, EPSPs and IPSPs are spatially overlapping. The inhibitory responses result from both withdrawal of excitation and actual inhibition, just as the excitatory responses result from both withdrawal of inhibition and actual excitation.

It is thus plausible that the inputs to a simple cell are arranged in *push-pull*, i.e. they come from pairs of cells with opposite signed receptive fields, one of which provides excitation and the other inhibition. For example, an ON subregion would be the result of excitatory ON-center inputs as well as of inhibitory OFF-center inputs.

This complementary arrangement of excitation and inhibition is also consistent with extracellular recording studies in cat V1 (Tolhurst and Dean, 1990, 1987; Heggelund, 1986, 1981; Palmer and Davis, 1981; Glezer et al., 1982, 1980). A similar push-pull arrangement might also be used by ganglion cells to integrate bipolar signals (Gaudiano, 1992).

For reasons of simplicity we assume that both the excitation and the inhibition are contributed by feed-forward connections. In this we differ from a number of recent models that consider intracortical feedback crucial in sharpening the selectivity conferred by the inputs from the lateral geniculate nucleus (Ben-Yishai et al., 1995; Somers et al., 1995; Suarez et al., 1995). While the feed-forward view is supported by recent evidence (Ferster et al., 1996; Reid and Alonso, 1995), the linear model should not necessarily be identified with a feed-forward arrangement inputs. A linear receptive field could, in principle, be constructed with pure feed-forward connections, pure feed-back connections, or a combination of feed-forward and feedback.

4.4 Linearity of Excitation and Inhibition

We now make the linking assumption that *simple cell synaptic conductances depend linearly on the responses of LGN neurons.* This assumption is quite realistic as far as excitatory conductances are concerned, because there is evidence for direct excitatory inputs from the LGN to simple cells (Ferster and Lindstrom, 1983), and direct synaptic transmission can be well approximated by a linear transformation of the presynaptic firing rate into the postsynaptic conductance*.

There is, however, conflicting anatomical evidence of direct geniculocortical inhibition (Einstein et al., 1987; Garey and Powell, 1971), and none of the physiological evidence supports its existence (Reid and Alonso, 1995; Ferster and Lindstrom, 1983; Tanaka, 1983; Toyama et al., 1977b,a; Toyama and Takeda, 1974; Toyama et al., 1974; Watanabe et al., 1966). Since, most inhibitory inputs from the LGN to simple cells are disynaptic (Ferster and Lindstrom, 1983), the linearity of inhibition would seem to require an inhibitory cortical interneuron that performs a linear integration of LGN inputs and encodes them linearly into firing rate.

4.5 Simplified Model of a Cortical Cell

** Figure 10 About Here **

We adopt a very simplified model of a cortical cell (Figure 10): a single compartment circuit with only passive conductances. In particular we consider a leak conductance g_{leak} and two synaptic conductances, one excitatory (g_e) and one inhibitory (g_i). The membrane potential of a model cell then obeys

$$-C\frac{dV}{dt} = g_e(V - V_e) + g_i(V - V_i) + g_{leak}(V - V_{leak}), \tag{2}$$

where C is the membrane capacitance, and V_{leak}, V_e and V_i are respectively the equilibrium potentials of the leak, excitatory and inhibitory channels.

** Figure 11 About Here **

This view of the cellular physiology deliberately ignores many known aspects of neuronal biophysics, such as voltage- and calcium-dependent channels, the possible nonlinear interactions between inputs caused by the dendritic structure, and the possible effects of electrotonic distance from the soma (Koch and Segev, 1989). Our model of a cortical cell is however in many respects a reasonable approximation. For example, Figure 11A shows the responses of an intracellularly recorded cortical neuron to sinusoidal current injection at different temporal frequencies. The membrane potential responses are dominated by their sinusoidal (*first harmonic*) component. This means that if the generation of spikes is ignored, the membrane of cortical neurons can be reasonably modeled by passive conductances, which endow it with a linear behavior. In particular, when the first harmonic responses are plotted against the temporal frequency of the stimulus (filled circles in Figure 11B), they are well fit by the predictions of a single-compartment model of the cell (dashed curve in Figure 11B).

4.6 Linear Integration of the Synaptic Inputs

The push-pull arrangement of the LGN inputs to a simple cell can lead (through a balance of excitation and inhibition) to a perfectly linear integration of the synaptic conductances by the cell membrane (Carandini and Heeger, 1994).

For the sake of simplicity, consider the steady-state behavior of the membrane ($dV/dt = 0$). At steady state Equation 2 can be rewritten as

$$V = \frac{g_e V_e + g_i V_i + g_{leak} V_{leak}}{g_e + g_i + g_{leak}}. \tag{3}$$

The push-pull arrangement (Section 4.3) guarantees that every increase in excitation will correspond to a decrease in inhibition, and vice-versa. In particular, we assume that g_e and g_i are balanced so that the total conductance of the cell is constant:

$$g_e(t) + g_i(t) + g_{leak} = g_0. \tag{4}$$

Equation 3 can then be rewritten as $V = [g_e V_e + g_i V_i + g_{leak} V_{leak}]/g_0$, which is a linear function of g_e and g_i. In words, *the membrane potential V is a linear function of the synaptic*

conductances.

If our model of the cell membrane as a single compartment is a good approximation, the exact balance of excitation and inhibition expressed in Equation 4 is an essential condition for the linear integration of the synaptic conductances. If, on the other hand, there is substantial electrotonic distance between synaptic sites on the membrane, there are other conditions in which linear integration of the synaptic inputs is possible. For example, Blomfield (1974), showed that if inhibition is located on the soma, and excitation is electrotonically remote from it, there is a range of synaptic activations in which the membrane potential will be approximately a linear combination of the excitatory and inhibitory synaptic conductances. This approach would not require the strict balance of excitation and inhibition (Equation 4), but it would require additional assumptions about the dendritic structure of the cell, the sites of the inputs, and the range of the synaptic conductances.

4.7 Spike Rate Encoding

If simple cells integrate their synaptic inputs linearly, if those inputs depend linearly on LGN activity, and if LGN activity is a linear function of the stimulus intensity distribution, then simple cells will integrate the stimulus intensity distribution linearly. This Section discusses the final, nonlinear stage of the model, which is responsible for the encoding of the input-driven membrane potential responses into spike trains.

Many characteristics of firing rate encoding are consistent with the view that the firing rate responses are a rectified copy of the membrane potential responses. An example of this can be seen in Figure 11A. The spike responses closely mirror the membrane potential responses, and there is a clear threshold below which no spikes are generated. Once above threshold, the firing rate grows with the amplitude of the membrane potential modulation. There is in fact a large literature pointing to a linear or bilinear relation between injected current and firing rate, once the current is above a threshold level (see Stafstrom et al., 1984, and references therein).

There is however an additional experimental result that is not consistent with the view

of firing rate encoding as rectification: the spike rate encoder has notable dynamic properties (Carandini et al., 1994). For example, cortical neurons typically exhibit *spike frequency adaptation*, meaning that the firing rate response to steady depolarization decreases with time (Stafstrom et al., 1984). Dynamic properties of spike encoding are also evident in Figure 11. Figure 11B plots the temporal frequency tuning of the first harmonic of the spike train taken from records like the ones in Panel A. It is clear that spike rate encoding is not at all independent of the temporal frequency of the stimulus, as would be the case for rectification. The middle temporal frequencies are transmitted much better than the low temporal frequencies, and the very high temporal frequencies are completely cut off.

Rectification cannot account for these behaviors because it is a *static* nonlinearity, i.e., it depends only on the instantaneous value of its argument. Strictly speaking, then, rectification is incorrect, because it would predict that the spike encoding properties would not depend on the past history of stimulation. Hence, we are forced to adopt a slightly more complicated model of spike rate encoding. In particular, we model the spike rate encoder as a band-pass filter followed by rectification*.

The behavior of this spike encoder model is actually quite simple when incorporated into the linear model of simple cells. In the full model, the spike rate encoder comes after a (synaptic) linear spatiotemporal weighting function. Since a chain of linear systems is itself a linear system, we can treat the band-pass (spike encoder) linear filter and the (synaptic) spatiotemporal linear weighting function as a single linear system. The weighting function of this final system is partly due to the synaptic inputs and partly to the band-pass properties of the firing rate encoder.

5 Some Nonlinear Properties of Simple Cells

Having described the linear model of simple cells, and having discussed its numerous successes and its possible biophysical implementation, it is now the time to discuss its failures. We have already encountered a number of occasions in which the linear model fails to yield precise

quantitative predictions. Indeed, the behavior of simple cells is in many ways nonlinear (see Heeger, 1993, 1992b for a review). This Section describes some of these nonlinearities, which will be discussed more quantitatively once we have introduced the normalization model.

5.1 Contrast Responses

** Figure 12 About Here **

Presented with a change in contrast, a linear neuron would scale its response by the same amount. The responses of a simple cell, instead, are often not proportional to stimulus contrast. An example of this is illustrated in the central column of Figure 12A. The central column shows the spike histograms of a simple cell in response to a drifting grating. The different rows correspond to different contrasts. As the contrast goes from 50% to 100%, the response does not double. Instead it grows very little. This phenomenon is known as *response saturation* (Sclar et al., 1990; Albrecht and Hamilton, 1982; Ohzawa et al., 1982; Dean, 1981; Maffei and Fiorentini, 1973). Cells can even exhibit "supersaturation", in which increasing the contrast of the stimulus reduces the amplitude of the responses (Bonds, 1991; Li and Creutzfeldt, 1984).

Response saturation is not due to the high firing rates. This can also be seen in Figure 12. The three columns in Panel A show the responses of a cell to three gratings of different spatial frequency. Even though the left column and the right column stimuli elicit fewer spikes than the central column stimulus, there clearly is response saturation. This phenomenon thus depends on the contrast of the stimulus *per se*, not on the amplitude of the responses it elicits in the cell. This property of the contrast responses can be more precisely observed in Panel B, which plots the amplitude of the responses shown in Panel A, as a function of contrast. In spite of the amplitude saturation the three contrast responses are vertical shifts of each other. Since the vertical axis is logarithmic, a vertical shift means that *the ratio of the responses to any two different spatial frequencies is constant, irrespective of the stimulus contrast.*

** Figure 13 About Here **

Another way to express this property is to say that the shape of the spatial frequency tuning curve is independent of the contrast at which it is measured. Changing the contrast of the stimuli just scales the tuning curve. This has been observed for both spatial frequency tuning and orientation tuning (Skottun et al., 1987; Li and Creutzfeldt, 1984; Albrecht and Hamilton, 1982; Sclar and Freeman, 1982; Movshon et al., 1978b). An example of the contrast invariance of the orientation tuning is shown in Figure 13.

The contrast independence of the tuning curve shapes would be easy to explain if the contrast responses of simple cells were linear. The responses of a linear cell to two stimuli S_1, S_2 with the same contrast c could be written as $cL(S_1)$ and $cL(S_2)$. Their ratio would be $L(S_1)/L(S_2)$, independent of the contrast c. We have seen, however, that the contrast responses of real simple cells are nonlinear, since they often saturate at high contrasts. The contrast independence of the tuning curves is thus by no means a trivial property.

5.2 Nonspecific Suppression

The response to a preferred stimulus can be suppressed by superimposing an additional stimulus that would not elicit any response when presented alone. This phenomenon is a violation of superposition, a defining property of linearity. We call it *nonspecific suppression*, as it has been found to be independent of direction of motion, largely independent of orientation and broadly tuned for spatial and temporal frequency (Carandini and Heeger, 1995; Geisler and Albrecht, 1992; DeAngelis et al., 1992; Nelson, 1991; Bonds, 1989; Gulyas et al., 1987; Kaji and Kawabata, 1985; De Valois et al., 1985; Li and Creutzfeldt, 1984; De Valois and Tootell, 1983; Morrone et al., 1982; Hammond and MacKay, 1981; Dean et al., 1980; Bishop et al., 1973). After some debate, there is now consensus that that cross-orientation inhibition can be driven dichoptically (with one grating in each eye), although monoptic suppression (with both gratings in the same eye) is typically stronger than dichoptic suppression. (Walker et al., 1996; Sengpiel et al., 1995; Sengpiel and Blakemore, 1994; DeAngelis et al., 1992; Freeman et al., 1987; Ohzawa and Freeman, 1986b,a; Ferster, 1981)

The origins of suppression are most likely cortical, as it is completely absent in monkey P LGN and cat LGN (Movshon et al., 1994; Bonds, 1989). Moreover, the temporal properties of nonspecific suppression are consistent with the view that it originates from complex cells or from a large pool of simple cells. Indeed, the suppression elicited by a drifting grating is not modulated in time (Bonds, 1989; Morrone et al., 1982), and the suppression elicited by a contrast-modulated modulates at twice the frequency of the stimulus (Morrone et al., 1982).

** Figure 14 About Here **

Figure 14 shows an example of nonspecific suppression. The stimulus was a plaid made of two gratings. One (the "test") drifted in the cell's preferred direction and evoked a large response when presented on its own. The other grating (the "mask") drifted at right angles with the test grating, and was ineffective in driving the cell. Its presence, however, clearly suppressed the responses. For example, when the mask contrast was 50% the cell responded only when the test grating had high contrast.

From Figure 14B one can see that the presence of the mask shifts the contrast response to the right (Bonds, 1989). This corresponds to a scaling of contrast (Heeger, 1992b). We will see later that the contrast responses shift to the right only when the cell is *completely unresponsive to the mask*. If the cell gives even a minimal response to the mask the effect is more complicated than just a rightward shift.

5.3 Temporal Nonlinearities

** Figure 15 About Here **

Simple cells display prominent temporal nonlinearities. Figure 15A provides a good example of this. As the stimulus contrast increases, the responses occur earlier in time. This is called *phase advance* (Albrecht, 1995; Carandini and Heeger, 1994; Dean and Tolhurst, 1986). It is a nonlinearity because for a linear cell scaling the input would just scale the output, not

change its timing. Phase advance is not entirely cortical in origin, but has a strong cortical component. Phase advance in cat LGN (measured at 2-5 Hz) is on average less than 20 ms (Carandini et al., 1993), whereas phase advance in cat V1 (measured at 2 Hz) is on average around 47 ms (Dean and Tolhurst, 1986). In the monkey LGN, phase advance is present in M cells, but it is completely absent in P cells (Sherman et al., 1984).

Another temporal nonlinearity of simple cells was uncovered by Reid et al. (1992). They measured the responses of cat simple cells to eight different stimuli and to the compound stimulus obtained by summing the eight stimuli. They found that the responses to the compound stimulus occur earlier in time than the linear prediction obtained from the responses to the individual stimuli. This decrease in "integration time" is quite prominent, in the range of 5-60 ms, and there is evidence that its origin is cortical (Reid et al., 1992).

Finally, a third temporal nonlinearity of simple cell responses is given by the contrast dependence of their temporal frequency tuning (Holub and Morton-Gibson, 1981). In particular, increasing stimulus contrast increases the cell's responsivity to the high temporal frequencies (Hawken et al., 1992). An example of this is shown in Figure 15B. According to the linear model increasing the contrast should just scale the responses, with no effect on their temporal frequency tuning. The origins of this nonlinear behavior are partially subcortical, since it was observed in the cat retina (Shapley and Victor, 1978) and in the monkey M LGN (Benardete et al., 1992). There is however evidence that in the monkey this behavior is much stronger in V1 than in the LGN. Preliminary results by M. J. Hawken et al. (personal communication) indicate that on average the high-cutoff frequency of V1 cells changes from around 10 Hz at 8-16% contrast to around 30 Hz at 64% contrast. By contrast, the average change in high-cutoff frequency of LGN cells is negligible.

6 The Normalization Model of Simple Cells

We have now seen many cases in which the linear model is an inadequate description of simple cells. Given its numerous successes, however, it would be unwise to dispense with the linear

model altogether. The linear model can be extended by endowing it with a mechanism that controls the gain (sensitivity to input) and the time course of the cell responses (Carandini and Heeger, 1994; Heeger, 1993, 1992b, 1991; Albrecht and Geisler, 1991; Robson, 1988). This mechanism ensures that the gain decreases when the contrast of a stimulus is increased, or when another stimulus is superimposed to it.

This is done in the *normalization* model (Heeger, 1992b), depicted in Figure 1B. First comes a spatiotemporal linear weighting function. This is followed by a normalization stage, where each cell's linear response is divided by a quantity proportional to the pooled activity of a large number of other cells (the *normalization pool*). The normalization stage is followed by rectification.

Normalization is a nonlinear operation; one input (a cell's underlying linear response) is divided by another input (the activity of the normalization pool). The normalization pool is assumed to include cells tuned to all orientations, spatial frequencies and temporal frequencies, so that its overall response is uniform across these parameters. The effect of normalization is that the response of each cell is rescaled with respect to local stimulus energy. The name "normalization" is due to the assumption that the gain of every cell in the pool is rescaled by the same amount ("normalized").

The normalization model explains all the nonlinearities described in the previous Section, while retaining the main features of the linear model. According to the normalization model, a cell's selectivity is attributed to summation (the linear stage) and its nonlinear behavior is attributed to division (the normalization stage). For example the model explains cross-orientation inhibition because a given cell is suppressed by many other cells including those with perpendicular orientation tunings. It explains response saturation because the divisive suppression increases with stimulus contrast.

The rest of this Section is devoted to the description of a biophysical implementation of the normalization model. Once the details of the model are laid out, one can derive closed-form equations for the responses of the model to some types of visual stimulus. Section 7 will show

how these closed-form equations perform in fitting the experimental data.

6.1 Shunting Inhibition

** Figure 16 About Here **

We hypothesize that normalization (division) acts by controlling the overall conductance of the cell's membrane (Carandini and Heeger, 1994). Figure 16 shows the biophysical structure of the model. The circuit depicted is identical to the linear model shown in Figure 10, with the addition of a *shunting* synaptic conductance, g_{shunt}. The shunting conductance has the property that its equilibrium potential V_{shunt} is the same as the resting potential of the cell, V_{rest} (Coombs et al., 1955; Fatt and Katz, 1953). For ease of notation we pick this value as the origin of the membrane potential measurements:

$$V_{shunt} = V_{rest} = 0. \tag{5}$$

The membrane potential in the circuit shown in Figure 16 obeys the following equation:

$$\begin{aligned} -C\frac{dV}{dt} &= g_i(V - V_i) + g_e(V - V_e) + g_{shunt}(V - V_{shunt}) + g_{leak}(V - V_{leak}) \\ &= gV - I_d, \end{aligned} \tag{6}$$

where g is the *total conductance*, and I_d is what we call the *driving current**:

$$\begin{aligned} g &\equiv g_e + g_i + g_{shunt} + g_{leak} \\ I_d &\equiv g_e V_e + g_i V_i + g_{shunt} V_{shunt} + g_{leak} V_{leak} \end{aligned}$$

Since g_e and g_i are in push-pull arrangement (Equation 4, in Section 4.6), and the equilibrium potential of the shunt is equal to the resting potential of the cell (Equation 5), the above equations can be rewritten as

$$g = g_0 + g_{shunt} \tag{7}$$

$$I_d = g_e V_e + g_i V_i + g_{leak} V_{leak} \tag{8}$$

As a result, *current and conductance are decoupled*: the driving current I_d is a linear function of the excitatory and inhibitory inputs g_e and g_i and does not depend on the shunt conductance g_{shunt}. The overall conductance g of the cell does not depend on the excitatory and inhibitory inputs g_e and g_i and is completely controlled by the shunt conductance g_{shunt}.

6.2 Role of the Membrane Conductance

** Figure 17 About Here **

We now make a crucial hypothesis: that *the shunt conductance g_{shunt} grows with the activity of the normalization pool.* The exact relation that we hypothesize between the two is depicted in Figure 17A. This relation is entirely *ad hoc*: it was chosen so that it would allow us to solve the model mathematically, and it is not based on experimental data. Its principal consequence is that the conductance of a model cell will depend on the energy of the visual stimuli. Increasing the energy of a stimulus, e.g. by increasing its contrast, will increase the activity of the cells belonging to the normalization pool, and thus cause an increase in cell conductance.

The membrane conductance g, in turn, determines the cell's *gain* which is the relation between input I_d and output V. At steady state the gain is $V/I_d = 1/g$, inversely proportional to the conductance. The conductance also has an effect on the time course of the response. The membrane capacitance takes time to charge and discharge, and this time is proportional to the membrane *time constant* $\tau = C/g$, which is also inversely proportional to the conductance. Figure 17B illustrates these concepts. It shows the responses of the membrane to a current step, for three values of the conductance g. If the conductance is very small, the response is slow and there is high gain (that is, the voltage response to a given current is high). If the conductance g is very large (the membrane is very leaky), it has small gain and it is fast in charging and discharging the capacitance.

To summarize, increasing the stimulus energy increases the conductance of model cells. This has two effects: (1) it provides divisive inhibition, by reducing the cell's gain; and (2) it shortens the latency of the responses, by reducing the cell's time constant.

7 Testing the Normalization Model

The variables in the model depend on each other in a circular way: (1) each cell's firing rate depends on its membrane potential (Figure 5); (2) each cell's membrane potential depends on its driving current and on its conductance (Equation 6); (3) each cell's conductance depends on the firing rates of the cells in the normalization pool (Figure 17A). The model is a nonlinear neural network (Grossberg, 1988), and is in general quite complicated because both the driving current and the conductance vary over time. As a consequence, in order to predict its responses one is often forced to resort to computer simulations (Heeger, 1993, 1992b).

For drifting sine grating and plaid stimuli, however, we were able to solve the model and derive closed-form equations for its responses (Carandini and Heeger, 1994). In fact the model was designed with these equations in mind. The advantage of closed-form equations is that we can fit them to the data and see how the model performs.

7.1 Fitting the Grating Responses

We derived two equations for the first harmonic responses to drifting gratings, one for response amplitude and one for response phase. These equations are detailed in Appendix A, which also contains a sketch of their derivation.

Here we concentrate on the expression for the response amplitude, which is:

$$\text{amplitude}(R) \propto \left[\text{amplitude}(L)\ \frac{c}{\sqrt{\sigma(\omega)^2 + c^2}}\right]^n, \tag{9}$$

where c is the contrast of the grating, and ω is its temporal frequency. The role of the quantities L, $\sigma(\omega)$, and n is easy to understand if one keeps in mind the structure of the model (Figure 1B). $L(t)$ is the output of the cell's linear weighting function (Equation 1) when the grating has unit contrast. It is a sinusoid; here we are concerned only with its amplitude, which for a stimulus of contrast c is [amplitude(L) c]. The normalization stage divides that by a quantity that depends on the activity of a large number of neurons. Appendix A shows that for drifting grating stimuli this quantity is $\sqrt{\sigma(\omega)^2 + c^2}$, where the function $\sigma(\omega)$ is related to

the low-pass properties of the cell membrane. It grows with ω, the temporal frequency of the stimulus. Finally, the exponent n is a constant and is related to the rectification stage that encodes the membrane potentials into firing rates.

Equation 9 is very similar to one that was empirically found to fit V1 contrast responses (Sclar et al., 1990; Albrecht and Hamilton, 1982). The model however does more than that: it predicts the dependence of the responses on all the other stimulus parameters, besides contrast: spatial frequency, temporal frequency and orientation. To test these predictions we used the model to fit large data sets obtained by stimulating simple cells with drifting gratings of a variety of contrasts, spatial frequencies, temporal frequencies and orientations (Carandini and Heeger, 1994). Some of our results are shown in Figures 12, 13 and 15. We found that the model provides good fits to the data, and accounts quantitatively for all the linear and nonlinear behaviors described in this Chapter.

Figure 12 illustrates the dependence of the responses of a simple cell on contrast and spatial frequency. The fits of the model capture the shape of the contrast responses and the fact that changing the stimulus spatial frequency simply scales these responses. This means that when the amplitude is plotted on a logarithmic scale, as in Panel B, changes in spatial frequency shift the contrast responses vertically, without affecting their shape. Very similar results are obtained when the gratings are varied in orientation instead of spatial frequency, as shown for example in Figure 13A (Section 5.1).

The reason for this behavior can be understood by considering the expression for the response amplitude predicted by the model (Equation 9). The expression can be seen as the product of two factors, $[\text{amplitude}(L)]^n$ and $(c/\sqrt{\sigma(\omega)^2 + c^2})^n$. The first factor depends on L, the response of the cell's linear receptive field to the grating at unit contrast. The second factor depends only on the contrast c and on the temporal frequency ω of the grating. For a fixed temporal frequency the shape of the contrast responses is entirely controlled by the second factor; hence, varying stimulus orientation or spatial frequency leads to a vertical shift of the contrast responses (Figs. 12B and 13A). Likewise, the shape of the orientation and spatial frequency tuning curves is entirely controlled by the first factor, hence, varying

stimulus contrast leads to a vertical shift of the orientation and spatial frequency tuning curves (Fig. 13B).

Figure 15 shows that the normalization model also predicts the temporal nonlinearities of simple cell responses (Section 5.3). Panel A shows that the model captures the phase advance behavior observed with increasing stimulus contrast. Panel B shows that the normalization model correctly predicts the dependence of the temporal frequency tuning on contrast. As the contrast increases the cell is more responsive to the high temporal frequencies. The reason for the contrast-dependence of response phase and of the cell's temporal frequency tuning can be understood in terms of the basic properties of the normalization model (Figure 17). In model cells, increasing stimulus contrast increases the membrane conductance, which decreases its time constant. A decreased time constant leads to a shorter response latency, thus explaining phase advance. Decreasing the time constant also enables the membrane potential to better follow the high temporal frequencies. This explains the fact that at high contrast the cell responds to higher temporal frequencies than at low contrast. This last argument can be made more quantitative by examining Equation 9. Appendix A shows that the quantity $\sigma(\omega)$ grows with the temporal frequency ω of the stimulus. This is a consequence of the low-pass properties of the membrane. At low contrasts c, σ has a strong effect, considerably scaling down the responses. At high contrasts, when $c \gg \sigma(\omega)$, the effect is much weaker, so the high frequency responses are relatively enhanced.

7.2 Fitting the Plaid Responses

Appendix B sketches the derivation of approximate equations for the amplitude and phase of the first harmonic response to plaid stimuli. In particular the expression for response amplitude is

$$\text{amplitude}(R) \propto \left[\frac{\text{amplitude}(c_1 L_1(t) + c_2 L_2(t))}{\sqrt{\sigma(\omega)^2 + c_1^2 + c_2^2}} \right]^n , \tag{10}$$

where c_1 and c_2 are the contrasts of the two gratings, $L_1(t)$ and $L_2(t)$ are the sinusoidal responses of the linear weighting function to the individual gratings at unit contrast, and the

remaining symbols have the same meaning as in the expression for the response to individual gratings (Equation 9). Since the spatiotemporal receptive field is linear, its response to the plaid is just a linear combination of its responses to the individual gratings, $c_1 L_1(t) + c_2 L_2(t)$. The normalization stage divides that by a quantity that depends on the activity of a large number of neurons. For plaids composed of two gratings this quantity is $\sqrt{\sigma(\omega)^2 + c_1^2 + c_2^2}$ (Appendix B). Finally, the rectification stage is responsible for the exponent n (Figure 5).

We tested the predictions of the model by recording the responses of simple cells to stimuli composed of two gratings whose contrasts c_1 and c_2 assumed a variety of different values (Carandini and Heeger, 1995). The two gratings always had the same temporal frequency, but could differ in orientation and/or in spatial frequency.

Figure 14 shows an example of our results. In this case we chose the gratings so that one of them (the "test") would strongly drive the cell, while the other (the "mask") would not elicit any response when presented alone. The Figure shows the responses of the cell for 5 different test contrasts, and three different mask contrasts. The cell was actually tested with a wider variety of conditions, and the predictions of the model (Equation 10 for response amplitude, and an equation for response phase described in Appendix B) were fit to all the responses at once. When we discussed nonspecific suppression (Section 5.2), we used this Figure as an example of the rightward shift of the contrast responses that results from masking. It is clear from the quality of the fits that the normalization model captures this behavior. Indeed, the effect of a mask on the responses of a model cell can be seen directly in Equation 10. Let for example grating 2 be the "mask". If, as in Figure 14, the mask alone does not elicit any response ($L_2 \approx 0$), then the suppressive effect of the mask is due to the fact that the mask contrast, c_2, appears only in the denominator. In these conditions, the effect of an increase of c_2 in the denominator is to shift the contrast response to the right (Heeger, 1992b).

** Figure 18 About Here **

Things become more complicated when both gratings that compose the plaid are able to drive the cell (even minimally) when presented alone. In these conditions each component of

the plaid acts both as a "test" and as a "mask". An example of this is shown in Figure 18. In this experiment, the "mask" was not as effective as the "test" in driving the cell, but both component gratings did elicit some spikes when presented alone. Panel B shows that depending on the contrast of the stimulus, increasing the contrast of the mask can either enhance or suppress the responses. This behavior is predicted by the normalization model (continuous curves), as the contrasts of the two gratings, c_1 and c_2, appear both in the numerator and in the denominator of Equation 10. Increasing one of the two can result either in an enhancement or in a reduction in the response, depending on the amplitude and phases of L_1 and L_2, the responses of the linear receptive field to the individual gratings, and on the size of the term $\sigma(\omega)$ in the denominator.

In summary, masking will cause the rightward shift in the contrast responses shown in Figure 14 only if the mask elicits negligible responses when presented alone. The normalization model predicts that in the general case masking will deform the contrast responses in more complicated ways, like that shown in Figure 18. Some researchers have indeed reported that suppression shifts the contrast responses downward on a logarithmic scale (Morrone et al., 1982; Dean et al., 1980), which is in apparent contradiction with the rightward shift reported by others (Section 5.2). The normalization model might explain this discrepancy: if the mask elicits even minimal (subthreshold) responses when presented alone, it may very well cause a downward shift in the contrast responses. The importance of "crosstalk' between different components of a stimulus has been previously recognized (Bonds, 1992; Bauman and Bonds, 1991), and the normalization model may provide a quantitative framework to understand these interactions.

** Figure 19 About Here **

The nature of the responses of simple cells to plaid stimuli is further illustrated by the results of another experiment (Figure 19). When presented alone the gratings elicited the responses shown in Panels A and C. Had the neuron been linear, it would have obeyed superposition, and its plaid response (Panel B) would have been well fitted by the linear prediction (dashed

curve) obtained by summing the responses to the individual gratings. The actual response to the plaid instead was smaller than that, and it was well fit by the normalization model (continuous curve in Panel B), which takes into account the fact that the gratings also provide divisive suppression. A useful way to picture this is with a polar plot (Figure 19D). Each data point in Figure 19D corresponds to a first harmonic sinusoid fitted to the responses of the simple cell. The amplitude of the sinusoid is given by the distance from the origin, and the phase of the sinusoid is given by the angle with the horizontal axis. The data points going in the 1 o'clock direction correspond to the responses to a single drifting grating whose contrast was varied between 0 and 50%. The data points going in the 9 o'clock direction correspond to the responses to the other drifting grating. The data points going in the 11 o'clock direction correspond to the responses to the "plaid" obtained by superimposing the two gratings. The dashed curve close to the plaid data is obtained by summing (vectorially) the curves fitting the individual grating responses. The actual plaid responses are smaller (closer to the origin) than this linear prediction. They also occur earlier (their angle with the horizontal axis increases, phase advance) than the linear prediction. This trend is captured by the continuous curve fitted by the normalization model, and is due to the fact that the local stimulus energy of the plaid is greater than that of the individual gratings. In the model this results in higher membrane conductance, which causes a decrease in the gain and time constant of the cell.

8 Biophysical Plausibility of the Normalization Model

The two models that we have described in this Chapter call for the existence of two different kinds of inhibitory inputs to V1 simple cells. The linear model postulates hyperpolarizing inhibition with the same selective tuning properties as the excitation. This inhibition is needed to ensure the linearity of the simple cell responses (Section 4). The normalization model in addition postulates shunting inhibition originating from a large number of cortical units. This nonselective shunting inhibition is responsible for controlling the cell's gain (Section 6). While there is strong evidence for the existence of the first (selective hyperpolarizing) kind of inhibition (Ulinski and Fowler, 1995; Ferster and Jagadeesh, 1992; Berman et al., 1991; Ferster,

1986), the second (nonselective shunting) kind of inhibition is for now mostly speculation. This Section discusses the effects of blocking inhibition in V1, and the plausibility of nonselective shunting inhibition.

8.1 Blocking Inhibition

Sillito (Sillito, 1984; Sillito et al., 1980; Sillito, 1975), blocked inhibition by iontophoresing a GABA antagonist on a region of cortex in the vicinity of the recording electrode. Blocking GABA strongly enhances the responses of V1 cells and dramatically broadens their orientation tuning curves. Based on these results, Sillito concluded that the orientation tuning of a cell is substantially sharpened by cross-orientation inhibition. Blocking GABA also largely eliminates direction selectivity (Sillito, 1977). In the conditions of Sillito's experiments, however, inhibition is blocked simultaneously in a large number of cortical cells. This greatly increases the responsiveness of all of these neurons, not just the one being recorded, potentially resulting in epileptogenesis (Chagnac-Amitai and Connors, 1989). It is not entirely clear, therefore, how to interpret Sillito's results.

Nelson et. al. (1994) have been able to perform a much more focused blockade of inhibition. They recorded intracellularly from a neuron while blocking inhibition *to only that one cell*. Blocking inhibition under these circumstances does not lead to a substantial change in orientation tuning. The normalization model is generally consistent with this result. According to the normalization model, blocking inhibition in a simple cell would destroy both the underlying linearity and the gain control (normalization), but it would not have much of an effect on selectivity. This is because neither the shunting (normalization) inhibition nor the hyperpolarizing (linear) inhibition are critical for selectivity.

Removing the shunting inhibition corresponds to setting $g_{shunt} = 0$ in the model. This would turn a cell described by the normalization model into one described by the linear model. It would result in the complete loss of the gain control mechanism. But the cell's selectivity would be largely unaffected because the selectivity is set up by the underlying linear spatiotem-

poral weighting function.

Removing the hyperpolarizing inhibition corresponds to setting $g_i = 0$. This would interfere with the linearity of the cell. Indeed, the exact balance of excitation and inhibition (Equation 4) is an essential condition in our model for the linear integration of the synaptic inputs. Removing the hyperpolarizing inhibition would, however, have little influence on the cell's tuning curves. Note that the excitatory inputs alone would still provide differential responses for stimuli in the ON and OFF subregions of a simple cell's receptive field; flashing a light in an ON subregion would evoke increased excitation and flashing a light in an OFF subregion would cause a withdrawal of excitation. Instead of Equation 3 the membrane potential would obey $V = (g_e V_e + g_{leak} V_{leak})/(g_e + g_{leak})$ which is a nonlinear (saturating) function of the synaptic excitation g_e. The cell's tuning would be a distorted version of that provided by the excitatory inputs. The distortion would only be noticeable if the excitatory inputs were large enough to approach saturation. But even then the distortion would be subtle; saturation would result in tuning curves with broader/flatter tops and steeper flanks.

8.2 Evidence For Nonselective Shunting Inhibition

Shunting inhibition is a widely cited proposal for how neurons might perform division (Coombs et al., 1955; Fatt and Katz, 1953). Its defining property is that it does not introduce any current when the cell is at rest, thus affecting only the cell's overall conductance. Shunting inhibition is usually thought to operate through $GABA_A$ synaptic channels, permeable to Chloride ions, because the equilibrium potential of Chloride is close to the resting potential of a typical cell.

There is evidence that there are strong inhibitory circuits in the cortex, and that these circuits operate through GABA-mediated shunting inhibition. A seminal study by Krnjević and his colleagues (Dreifuss et al., 1969) showed that electrical stimulation of the cortical surface produces very large (up to 300%) increases in membrane conductance, and that similar effects are obtained by iontophoretic application of GABA. In addition, Rose (1977) showed that iontophoresing GABA over V1 cells yields divisive effects on their visual responses (Rose,

1977). Cortical neurons are known to have $GABA_A$ receptors, which may be under a constant barrage of synaptic input originating from nearby cortical cells (Salin and Prince, 1996b,a).

In a study of cat area 18, Crook et al. (1992) found that a block of inhibition from cross-orientation sites broadens the orientation tuning, while a block of inhibition from iso-orientation sites increases response magnitude. This suggests that under normal conditions the cells receive inhibitory inputs from sites tuned to a wide variety of orientations. This nonspecific inhibition could be mediated by basket cells, which are GABAergic, and are considered to provide inhibition to cells tuned for all orientations (Kisvarday and Eysel, 1993; Kisvarday et al., 1993). In addition, there is direct physiological evidence for inhibition between V1 neurons when their orientation tuning differs by less than 45° (Hata et al., 1988).

8.3 Evidence For Conductance Increases

The normalization model predicts that a cell's conductance should increase with contrast. In order to evaluate how substantial these increases should be, we measured responses of a sample of simple cells for grating stimuli of a variety of temporal frequencies, contrasts and orientations and/or spatial frequencies. We then fit the data with the normalization model and used the best-fit parameters to estimate the time constants predicted by the model for zero contrast (uniform gray field) and full contrast gratings.

** Figure 20 About Here **

Figure 20 shows our results. The estimated time constants for a zero contrast stimulus generally varied between 5 and 100 ms. These time constants were estimated to drop by around a factor of three for a full contrast grating. A threefold drop in time constant corresponds to a threefold increase in conductance.

Our estimates of the conductance changes are surely inflated by our assumption that LGN cells are perfectly linear. As mentioned in Section 4.1, both cat LGN and monkey M LGN cells exhibit a gain control that is in some ways similar to what we have been ascribing entirely to

simple cells. Including the effects of known LGN nonlinearities would make our model more plausible, but it would clash with our goal of accounting for the data with a bare minimum of free parameters.

The conductance increase predicted by the model is based on extracellular data, and is at odds with intracellular *in vivo* measurements. These have consistently failed to demonstrate large conductance increases related to visual stimulation (Ferster and Jagadeesh, 1992; Berman et al., 1991). Berman et al. (1991), using sharp electrodes, measured only slight (less than 20%) conductance increases for a drifting bar stimulus. There are at least three possible explanations for this lack of conductance increase.

First, the drifting bar is a weak stimulus; our model predicts that increasing the contrast energy of the stimulus (e.g., by using a full contrast drifting grating instead of a drifting bar) would yield a larger change in conductance. Still, by our best estimates, we would have expected a conductance increase in their experiment of at least 50%, that is significantly greater than the amount they reported.

Second, it is possible that sharp electrodes may themselves introduce large shunts (Stratford et al., 1990), thus making them inappropriate for measuring conductance. On the other hand the results of Berman and collaborators were confirmed by Ferster and Jagadeesh (1992) using the whole-cell patch technique, which is not thought to introduce substantial leaks in the membrane.

Finally, it is possible that normalization is achieved by some other, as yet unknown, mechanism. This mechanism would manage to reduce the gain and the effective time constant of the cells, without recourse to shunting inhibition.

9 Conclusions

The ultimate goal of this research is to develop detailed, quantitative, predictive models of neural function in visual cortex. We will have succeeded - if we succeed at all - when we

can record from a neuron while presenting a basic set of visual stimuli (to measure model parameters) and then be able to predict the neuron's response to any visual stimulus. With such models, we can hope to understand the neural basis of perceptual appearance.

We have presented two models of V1 simple cells. The first, the linear model, accounts for the cells' selectivity but fails when the stimulus energy (or contrast) is varied. The second, the normalization model, overcomes this limitation by endowing the linear model with a feedback gain control mechanism. According to the normalization model the gain of each cell is rescaled by a quantity that grows with the overall activity of a pool of cortical neurons. This effectively normalizes the outputs with respect to stimulus energy. We have tested the normalization model, using grating and plaid stimuli, with encouraging results.

Hence it appears that response normalization is one of the primary roles of short-range suppression in cat and monkey V1. V1 neurons have a limited dynamic range, a limit to how strong an output signal they can generate and, thus, a limit to the range of contrasts over which they can respond differentially. The normalization operation preserves essential features of the neurons' responses (spatiotemporal-frequency tuning, orientation, and direction-selectivity in V1) while limiting the dynamic range of their outputs. For example, the response curves (for preferred and non-preferred stimuli) in Figures 12B and 13A are vertically shifted copies of one another; since the data are plotted on a logarithmic response scale, this means that the ratio of responses is about the same at all stimulus contrasts, even in the face of response saturation. This invariance, that we attribute to normalization, is critical for encoding visual information (e.g., about motion, orientation, binocular disparity, etc.) independent of contrast.

The issue of limited dynamic range is, of course, not restricted to V1 neurons. Gain control has been measured and modeled in a variety of other neural systems including: turtle photoreceptors (Baylor and Hodgkin, 1974), the vestibulo-ocular reflex (Lisberger and Sejnowski, 1992), and the velocity-selective neurons in the middle temporal (MT) area of the primate cortex (Simoncelli and Heeger, 1994). In particular, the normalization model of simple cell responses is analogous to models of retinal adaptation/normalization (Grossberg and Todorovic, 1988; Shapley and Enroth-Cugell, 1984; Tranchina et al., 1984; Sperling and Sondhi, 1968),

in which the stimulus intensity at a particular point is normalized with respect to the mean stimulus intensity. This makes the retinal response largely independent of the overall level of illumination, and allows the brain to proceed to process visual information without having to attend to the overall light level. Similarly, the normalization mechanism that we propose for V1 simple cells allows the brain to process visual information without having to attend further to contrast; the perceived orientation or motion direction of a stimulus is indeed largely invariant with respect to contrast.

However, the biophysical implementation of the normalization model is uncertain. We have presented a simple way in which normalization could be implemented physiologically, by means of shunting inhibition that increases the overall membrane conductance. This seems to us to be the simplest possible way that the normalization pool could control both the gain and the dynamics of a cell's response. It is, however, important to remember that we tested the model using extracellular data, so we have no proof that the overall conductance grows with stimulus energy. Indeed, there is evidence to the contrary (Berman et al., 1991). It is likely, therefore, that we will have to consider alternative mechanisms for normalization that do not require overall conductance increases. Such alternatives include conductance increases localized to the axon hillock and other mechanisms that control firing rate encoding.

Another problem with our model is that there are some important aspects of simple cell responses that we have largely ignored. First, practically all V1 cells adapt to prolonged stimulation (Saul, 1995; Sclar et al., 1989; Ohzawa et al., 1985; Albrecht et al., 1984; Movshon and Lennie, 1979; Maffei et al., 1973). Second, some V1 neurons are end-inhibited and/or side-inhibited, that is, stimulation outside of the classical receptive field suppresses responses to stimuli placed within the classical receptive field (Born and Tootell, 1991; DeAngelis *et al.*, 1994; and references therein). Third, some V1 neurons exhibit center-surround phenomena that are significantly more complicated; for some very specific stimulus configurations, introducing a stimulus in the surrounding field can facilitate a neuron's response (Gilbert et al., 1996; Sillito et al., 1995; Kapadia et al., 1995; Gilbert and Wiesel, 1990; van Essen et al., 1989; Nelson and Frost, 1985; Maffei and Fiorentini, 1976)

To some extent, adaptation and surround inhibition can be framed within the context of the normalization model (Heeger, 1992b). Some aspects of adaptation can be explained by assuming that the normalization mechanism has a long memory, and some aspects of end- and side-inhibition can be explained by assuming that neurons with surrounding receptive fields contribute to the normalization pool

In some cells, however, surround inhibition is arguably much stronger than nonspecific suppression and it is not clear that its nature is divisive. And there are some adaptation results that cannot be explained simply by changing a cell's gain (see (Heeger, 1992b) for review). Finally, the normalization model provides no explanation of surround facilitation. It may very well be that there are additional neural circuits mediating these phenomena.

10 Acknowledgements

We thank M. J. Hawken, C. S. Leonard, F. Mechler, L. P. O'Keefe, A. B. Poirson, D. Ringach, R. M. Shapley and C. Tang for help and suggestions. Supported by a CNR (Italy) grant and an NYU Dissertation Fellowship to MC, by a NIMH grant (MH50228) and a an Alfred P. Sloan Fellowship to DJH, and by a NIH grant (EY2017) and a Howard Hughes Medical Institute Investigatorship to JAM.

A Appendix: Predicted Responses to Gratings

This Appendix summarizes the derivation of closed form equations for the firing rate responses of model cells to drifting sinusoidal gratings.

The model predicts that the amplitude and phase of the response R to a drifting grating of contrast c and temporal frequency ω are given by:

$$\text{amplitude}(R) \propto \left[\text{amplitude}(L)\ \frac{c}{\sqrt{\sigma(\omega)^2 + c^2}}\right]^n, \tag{11}$$

$$\text{phase}(R) = \text{phase}(L) - \text{atan}\frac{\omega\tau_0}{\sqrt{1 + ((g_1/g_0)^2 - 1)c^2}}, \tag{12}$$

where

$$\sigma(\omega)^2 = \frac{1 + \omega^2\tau_0^2}{(g_1/g_0)^2 - 1}. \tag{13}$$

The parameters are: the amplitude and phase of the linear receptive field response L, the exponent n, the membrane time constant at rest τ_0, and the maximal increase in conductance that normalization can achieve, g_1/g_0.

The rest of the appendix describes the essential steps in the derivation of Equations 9, 13 and 12.

As we already mentioned, we model the firing rate R as a rectified copy of the membrane potential V, $R \propto max(0, V - V_{thresh})$, where $V_{thresh} > V_{rest}$. Rectification is however not very easily handled in mathematical derivations. We thus approximate rectification ($V_{thresh} > V_{rest}$)

with half-rectification ($V_{thresh} = V_{rest} = 0$) followed by elevation to the power n:

$$R \propto max(0, V)^n. \tag{14}$$

The quality of this approximation is shown in Figure 5. The value of the exponent n grows with the distance of the threshold V_{thresh} from the resting potential V_{rest}. If the threshold is very close to rest, then $n \approx 1$ ("half-rectification"). If the threshold is a bit above the rest, *e.g.* 5 mV higher, then $n \approx 2$ ("half-squaring"). If the threshold is much above rest then $n \approx 3$ or more. We assume the average exponent for the cells in the normalization pool to be $n = 2$ (Heeger, 1992a). Nevertheless, when fitting the model to individual cell responses we let the exponent be a free parameter.

For drifting sine grating stimuli the driving current $I_d(t)$ to each cell in the normalization pool is a sinusoid. This follows from the fact that I_d is a linear function of the stimulus intensity distribution. The amplitude of I_d thus depends linearly on contrast:

$$I_d(t) = c\ L(t) \tag{15}$$

For drifting sine grating stimuli, moreover, the sum of the squared linear inputs to all the units in the normalization pool is constant over time and is a measure of the *local stimulus energy*: $\sum(I_d)^2 = k^2c^2$. Here, k is a measure of the grating effectiveness in driving the normalization pool. This follows from the fact that the receptive fields of adjacent simple cells tend to exhibit either 90° or 180° phase relationships (Liu et al., 1992; Foster et al., 1983; Palmer and Davis, 1981; Pollen and Ronner, 1981). The normalization pool thus contains quadruples of cells with the same amplitude response but with phases 90° apart.

As a consequence, we find that the model reaches a steady state condition for drifting gratings, in which the conductance is constant and is given by

$$g = \sqrt{g_0^2 + (g_1^2 - g_0^2)c^2}, \tag{16}$$

where $g_1 = \sqrt{k^2 + g_0^2}$ is the maximal conductance, achieved when the stimulus has contrast $c = 1$.

Whenever the conductance does not change in time, the membrane of the cell behaves as a linear system. Stimulation with gratings of fixed contrasts thus results in a sinusoidal membrane potential. It is easy then to show, from Equation 6, that the amplitude and phase of such a sinusoid are given by

$$\begin{aligned} \text{amplitude}(V) &= \frac{\text{amplitude}(I_d)}{\sqrt{g^2 + (\omega C)^2}}, \\ \text{phase}(V) &= \text{phase}(I_d) - \text{atan}(\omega C/g). \end{aligned} \tag{17}$$

Equations 11, 13 and 12 can now be obtained by putting together the dependence of firing rate on membrane potential (Equation 14), the dependence of membrane potential on the driving current and on the conductance (Equation 17), and the dependence of the driving current and the conductance on stimulus contrast (Equations 15 and 16).

B Appendix: Predicted Responses to Plaids

The expressions derived in Appendix A for the firing rate of simple cells to drifting sinusoidal gratings can be approximately extended to stimuli composed of more than one grating. This Appendix describes the extension to stimuli composed of two drifting gratings. We restrict our attention to the case in which the two gratings have the same temporal frequency ω.

Let c_1 and c_2 be the contrasts of the two gratings. Let L_1 and L_2 (sinusoids) be the responses of the linear weighting function to the individual gratings. The driving current is just the sum of the linear responses weighted by the contrasts:

$$I_d(t) = c_1 L_1(t) + c_2 L_2(t). \tag{18}$$

The quantity $\sum I_d(t)^2$ is not in general constant in time, since it contains a component at twice the temporal frequency of the stimulus. Nonetheless, if one assumes that the relation between population firing rate and membrane conductance described by Figure 17A is preceded or followed by an averaging stage, then the resulting conductance change will be approximately

constant in time. The same result would be obtained by assuming that the normalization pool is very large so it effectively performs a local spatial average.

If the conductance is approximately constant over time, the same arguments as in the last Appendix apply, yielding:

$$\text{amplitude}(R) \propto \left[\frac{\text{amplitude}(c_1 L_1(t) + c_2 L_2(t))}{\sqrt{\sigma(\omega)^2 + c_1^2 + c_2^2}}\right]^n, \tag{19}$$

$$\text{phase}(R) = \text{phase}(c_1 L_1(t) + c_2 L_2(t)) - \text{atan}\frac{\omega\tau_0}{\sqrt{1 + ((g_1/g_0)^2 - 1)(c_1^2 + c_2^2)}}, \tag{20}$$

where σ is defined in Equation 13.

Footnotes

- Page 10. Some cells, however, are direction selective even though they have space-time separable weighting functions (McLean et al., 1994; Emerson and Citron, 1992; McLean and Palmer, 1989). The behavior of these cells cannot be accounted for by models like those advocated in this Chapter, in which direction selectivity is due to an underlying spatiotemporal linear stage.

- Page 19. Indeed, it is widely held that in the absence of synaptic plasticity and for reasonable presynaptic firing rates, each presynaptic spike results in a stereotyped postsynaptic conductance increase (Koch and Poggio, 1987; Jack et al., 1975). Synaptic transmission can then be considered to be a linear transformation whose impulse response is given by the shape of an isolated postsynaptic conductance increase.

- Page 22. An even better model for the spike rate encoder of visual cortical cells is given by the opposite arrangement, in which rectification is *followed* by a band-pass linear filter (Carandini et al., 1994). This arrangement however would make the model very hard to deal with analytically. The two arrangements give identical first harmonic responses to sinusoidal stimulation at a fixed temporal frequency.

- Page 28. The driving current depends on the cell's synaptic inputs, but it is independent of V, the cell's membrane potential. The driving current could only be measured by voltage clamping the cell. It is not the actual synaptic current, which depends on the membrane potential V.

References

Adelson, E. H. and Bergen, J. R. (1985). Spatiotemporal energy models for the perception of motion. *J. Opt. Soc. Am. A*, *2*, 284–299.

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

Albrecht, D. G., Farrar, S. B. and Hamilton, D. B. (1984). Spatial contrast adaptation characteristics of neurones recorded in the cat's visual cortex. *J. Physiol. (London)*, *347*, 713–739.

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

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

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

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

Baylor, D. A. and Hodgkin, A. L. (1974). Changes in time scale and sensitivity in turtle photoreceptors. *J. Physiol. (London)*, *242*, 729–758.

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

Benardete, E. A. and Kaplan, E. (1995). The receptive field of the primate P retinal ganglion cell. I: linear dynamics. Submitted to?

Benardete, E. A., Kaplan, E. and Knight, B. W. (1992). Contrast gain in the primate retina: P cells are not X-like, some M cells are. *Vis. Neurosci.*, *8*, 483–486.

Berman, N. J., Douglas, R. J., Martin, K. A. C. and Whitteridge, D. (1991). Mechanisms of inhibition in cat visual cortex. *J. Physiol. (London)*, *440*, 697–722.

Bishop, P. O., Coombs, J. S. and Henry, G. H. (1973). Receptive fields of simple cells in the cat striate cortex. *J. Physiol. (London)*, *231*, 31–60.

Blasdel, G. G. and Lund, J. S. (1983). Termination of afferent axons in macaque striate cortex. *J. Neurosci.*, *3*, 1389–1413.

Blomfield, S. (1974). Arithmetical operations performed by nerve cells. *Br. Res.*, *69*, 115–124.

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

Bonds, A. B. (1991). Temporal dynamics of contrast gain in single cells of the cat striate cortex. *Vis. Neurosci.*, *6*, 239–255.

Bonds, A. B. (1992). Spatial and temporal nonlinearities in receptive fields on the cat striate cortex. In R. B. Pinter and B. Nabet (Eds.), *Nonlinear vision*. CRC Press, Boca Raton, FL.

Born and Tootell (1991). Single-unit and 2-deoxyglucose studies of side inhibition in macaque striate cortex. *Proc. Natl. Acad. Sci.*, *88*, 7071–7075.

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

Carandini, M. and Heeger, D. J. (1995). Summation and division in V1 simple cells. In J. M. Bower (Ed.), *The neurobiology of computation: proceedings of the third annual Computation and Neural Systems conference* (pp. 59–65). Norwell, MA: Kluwer.

Carandini, M., Heeger, D. J. and Movshon, J. A. (1993). Amplitude and phase of contrast responses in LGN and V1. *Soc. Neurosci. Abs.*, *19*, 628.

Carandini, M., Mechler, F., Leonard, C. S. and Movshon, J. A. (1994). Firing rate encoding by visual cortical neurons in vitro. *Soc. Neurosci. Abs.*, *20*, 624.

Chagnac-Amitai, Y. and Connors, B. W. (1989). Horizontal spread of synchronized activity in neocortex and its control by GABA-mediated inhibition. *J. Neurophysiol.*, *61*, 747–758.

Coombs, J. S., Eccles, J. C. and Fatt, P. (1955). The inhibitory suppression of reflex discharges from motoneurones. *J. Physiol. (London)*, *130*, 396–413.

Creutzfeldt, O. D. and Ito, M. (1968). Functional synaptic organization of primary visual cortex neurones in the cat. *Exp. Br. Res.*, *6*, 324–352.

Crook, J. M. and Eysel, U. T. (1992). Gaba-induced inactivation of functionally characterized sites in cat visual cortex (area 18): effects on orientation tuning. *J. Neurosci.*, *12*, 1816–1825.

De Valois, K. and Tootell, R. (1983). Spatial-frequency-specific inhibition in cat striate cortex cells. *J. Physiol. (London)*, *336*, 359–376.

De Valois, K. K., De Valois, R. L. and Yund, E. W. (1979). Responses of striate cortex cells to grating and checkerboard patterns. *J. Physiol. (London)*, *291*, 483–505.

De Valois, R. L., Albrecht, D. G. and Thorell, L. G. (1982). Spatial frequency selectivity of cells in macaque visual cortex. *Vis. Res.*, *22*, 545–559.

De Valois, R. L., Thorell, L. G. and Albrecht, D. G. (1985). Periodicity of striate-cortex-cell receptive fields. *J. Opt. Soc. Am. A*, *2*, 1115–1123.

Dean, A. F. (1981). The relationship between response amplitude and contrast for cat striate cortical neurones. *J. Physiol. (London)*, *318*, 413–427.

Dean, A. F., Hess, R. F. and Tolhurst, D. J. (1980). Divisive inhibition involved in direction selectivity. *J. Physiol. (London)*, *308*, 84p–85p.

Dean, A. F. and Tolhurst, D. J. (1983). On the distinctiveness of simple and complex cells in the visual cortex of the cat. *J. Physiol. (London)*, *344*, 305–325.

Dean, A. F. and Tolhurst, D. J. (1986). Factors influencing the temporal phase of response to bar and grating stimuli for simple cells in the cat striate cortex. *Exp. Br. Res.*, *62*, 143–151.

DeAngelis, G. C., Freeman, R. D. and Ohzawa, I. (1994). Length and width tuning of neurons in the cat's primary visual cortex. *J. Neurophysiol.*, *71*, 347–374.

DeAngelis, G. C., Ohzawa, I. and Freeman, R. D. (1993a). The spatiotemporal organization of simple cell receptive fields in the cat's striate cortex. I. General characteristics and postnatal development. *J. Neurophysiol.*, *69*, 1091–1117.

DeAngelis, G. C., Ohzawa, I. and Freeman, R. D. (1993b). The spatiotemporal organization of simple cell receptive fields in the cat's striate cortex. II. Linearity of temporal and spatial summation. *J. Neurophysiol.*, *69*, 1118–1135.

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. *J. Neurophysiol.*, *68*, 144–163.

deBoer, E. and Kuyper, P. (1968). Triggered correlation. *IEEE Trans. Biomed. Eng.*, *15*, 169–179.

Derrington, A. M. and Lennie, P. (1984). Spatial and temporal contrast sensitivities of neurones in lateral geniculate nucleus of macaque. *J. Physiol. (London)*, *357*, 219–240.

Dreher, B., Fukada, Y. and Rodieck, R. W. (1976). Identification, classification and anatomical segregation of cells with X-like and Y-like properties in the lateral geniculate nucleus of old-world primates. *J. Physiol. (London)*, *258*, 433–452.

Dreifuss, J. J., Kelly, J. S. and Krnjević, K. (1969). Cortical inhibition and gamma-amminobutirric acid. *Exp. Br. Res.*, *9*, 137–154.

Einstein, G., Davis, T. L. and Sterling, P. (1987). Ultrastructure of synapses from the A-laminae of the lateral geniculate nucleus in layer IV of the cat striate cortex. *J. Comp. Neurol.*, *260*, 63–75.

Emerson, R. C. (1988). A linear model for symmetric receptive fields: Implications for classification test with flashed and moving images. *Spatial Vision*, *3*, 159–177.

Emerson, R. C. and Citron, M. C. (1992). Linear and nonlinear mechanisms of motion selectivity in simple cells of the cat's striate cortex. In R. B. Pinter and B. Nabet (Eds.), *Nonlinear vision*. CRC Press, Boca Raton, FL.

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

Enroth-Cugell, C. and Robson, J. G. (1984). Functional characteristics and diversity of cat retinal ganglion cells. *Inv. Opht. and Vis. Science*, *25*, 250–267.

Enroth-Cugell, C., Robson, J. G., Schweitzer-Tong, D. E. and Watson, A. B. (1983). Spatiotemporal interactions in cat retinal ganglion cells showing linear spatial summation. *J. Physiol. (London)*, *341*, 279–307.

van Essen, D. C., DeYoe, E. A., Olavarria, J. F., Knierim, J. J., Sagi, D., Fox, J. M. and Julesz, B. (1989). Neural responses to static and moving texture patterns in visual cortex of the macaque monkey. In D. M. K. Lan and C. D. Gilbert (Eds.), *Neural Mechanisms of Visual Perception* (pp. 137–156). Woodlands, Texas: Portfolio Publishing.

Fahle, M. and Poggio, T. (1981). Visual hyperacuity: spatiotemporal interpolation in human vision. *Proc. Roy. Soc. Lon. B*, *213*, 451–477.

Fatt, P. and Katz, B. (1953). The effect of inhibitory nerve impulses on a crustacean muscle fibre. *J. Physiol.*, *121*, 374–389.

Ferster, D. (1981). A comparison of binocular depth mechanisms in areas 17 and 18 of the cat visual cortex. *J. Physiol. (London)*, *311*, 623–655.

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

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

Ferster, D. (1990a). X- and Y-mediated current sources in areas 17 and 18 of cat visual cortex. *Vis. Neurosci.*, *4*, 135–145.

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

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

Ferster, D. and Jagadeesh, B. (1992). EPSP-IPSP interactions in cat visual cortex studied with *in vivo* whole-cell patch recording. *J. Neurosci.*, *12*, 1262.

Ferster, D. and Koch, C. (1987). Neuronal connections underlying selectivity in cat visual cortex. *Trends in Neuroscience*, *10*, 487–492.

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

Field, D. J. and Tolhurst, D. J. (1986). The structure and symmetry of simple-cell receptive field profiles in the cat's visual cortex. *Proc. R. Soc. Lon. B*, *228*, 379–400.

Foster, K. H., Gaska, J. P., Marcelja, S. and Pollen, D. A. (1983). Phase relationships between adjacent simple cells in the feline visual cortex. *J. Physiol. (London)*, *345*, 22P.

Freeman, R. D., Ohzawa, I. and Robson, J. G. (1987). A comparison of monocular and binocular inhibitory processes in the visual cortex of cat. *J. Physiol. (London)*, *396*, 69p.

Garey, L. J. and Powell, T. P. S. (1971). An experimental study of the termination of the lateral geniculo-cortical pathway in the cat and monkey. *Proc. R. Soc. Lon. B*, *179*, 41–63.

Gaudiano, P. (1992). A unified neural network model of spatiotemporal processing in X and Y retinal ganglion cells I: Analytical results. *biocyb*, *67*, 11–21.

Geisler, W. S. and Albrecht, D. G. (1992). Cortical neurons: isolation of contrast gain control. *Vis. Res.*, *8*, 1409–1410.

Gilbert, C. D., Das, A., Ito, M., Kapadia, M. and Westheimer, G. (1996). Spatial integration and cortical dynamics. *Proc. Natl. Acad. Sci.*, *93*, 615–622.

Gilbert, C. D. and Wiesel, T. N. (1990). The influence of contextual stimuli on the orientation selectivity of cells in primary visual cortex of the cat. *Vis. Res.*, *30*, 1689–1701.

Gizzi, M. S., Katz, E., Schumer, R. A. and Movshon, J. A. (1990). Selectivity for orientation and direction of motion of single neurons in cat striate and extrastriate visual cortex. *J. Neurophysiol.*, *63*, 1529–1543.

Glezer, V. D., Tscherbach, T. A., Gauselman, V. E. and Bondarko, V. E. (1980). Linear and nonlinear properties of simple and complex receptive fields in area 17 of the cat visual cortex. *Biol. Cyb.*, *37*, 195–208.

Glezer, V. D., Tscherbach, T. A., Gauselman, V. E. and Bondarko, V. E. (1982). Spatio-temporal organization of receptive fields of the cat striate cortex. *Biol. Cyb.*, *43*, 35–49.

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

Grossberg, S. and Todorovic, D. (1988). Neural dynamics of 1-d and 2-d brightness perception: A unified model of classical and recent phenomena. *Percept. & Psychophys.*, *43*, 241–277.

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

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

Hata, Y., Tsumoto, T., Sato, H., Hagihara, K. and Tamura, H. (1988). Inhibition contributes to orientation selectivity in visual cortex of cat. *Nature*, *335*, 815–817.

Hawken, M. J., Shapley, R. M. and Grosof, D. H. (1992). Temporal frequency tuning of neurons in macaque V1: effects of luminance contrast and chromaticity. *Inv. Opht. and Vis. Sci. (Suppl.)*, *33*, 955.

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

Heeger, D. J. (1992a). Half-squaring in responses of cat simple cells. *Vis. Neurosci.*, *9*, 427–443.

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

Heeger, D. J. (1993). Modeling simple cell direction selectivity with normalized, half-squared, linear operators. *J. Neurophysiol.*, *70*, 1885–1897.

Heggelund, P. (1981). Receptive-field organization of simple cells in cat striate cortex. *Exp. Br. Res.*, *42*, 89–98.

Heggelund, P. (1986). Quantative studies of enhancement and suppression zones in the receptive field of simple cells in cat striate cortex. *J. Physiol. (London)*, *373*, 293–310.

Hendrickson, A. E., Wilson, J. R. and Ogren, M. P. (1978). The neuroanatomical organization of pathways between the dorsal lateral geniculate nucleus and visual cortex in Old World and New World primates. *J. Comp. Neurol.*, *182*, 123–136.

Hochstein, S. and Shapley, R. M. (1976). Quantitative analysis of retinal ganglion cell classifications. *J Physiol*, *262*, 237–264.

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

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

Hubel, D. H. and Wiesel, T. N. (1972). Laminar and columnar distribution of geniculo-cortical fibers in macaque monkeys. *J. Comp. Neurol.*, *146*, 421–450.

Jack, J. J. B., Noble, D. and Tsien, R. W. (1975). *Electric current flow in excitable cells.* Oxford, UK: Oxford University Press.

Jagadeesh, B., Wheat, H. S. and 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. and Palmer, L. A. (1987a). An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. *J. Neurophysiol.*, *58*, 1233–1258.

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

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

Kaji, S. and Kawabata, N. (1985). Neural interactions of two moving patterns in the direction and orientation domain in the complex cells of cat's visual cortex. *Vis. Res.*, *25*, 749–753.

Kapadia, M. K., Ito, M., Gilbert, C. D. and Westheimer, G. (1995). Improvement in visual sensitivity by changes in local context: Parallel studies in human observers and in V1 of alert monkeys. *Neuron*, *15*, 843–856.

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

Kisvarday, Z. and Eysel, U. T. (1993). Functional and structural topography of horizontal inhibitory connections in cat visual cortex. *Eur. J. Neurosci.*, *5*, 1558–1572.

Kisvarday, Z. F., Beaulieu, C. and Eysel, U. T. (1993). Network of gaba-ergic large basket cells in cat visual cortex (area 18): implication for lateral disinhibition. *J. Comp. Neurol.*, *327*, 398–415.

Koch, C. and Poggio, T. (1987). Biophysics of computation: neurons, synapses and membranes. In G. M. Edelman, W. E. Gall and W. M. Cowan (Eds.), *Synaptic function*. Wiley, NY.

Koch, C. and Segev, I. (1989). *Methods in Neuronal modeling*. Cambridge, MA: MIT Press.

Kulikowski, J. J. and Bishop, P. O. (1981a). Fourier analysis and spatial representation in the visual cortex. *Experimentia*, *37*, 160–163.

Kulikowski, J. J. and Bishop, P. O. (1981b). Linear analysis of the response of simple cells in the cat visual cortex. *Exp. Br. Res.*, *44*, 386–400.

Lahica, E. A., Beck, P. D. and Casagrande, V. A. (1992). Parallel pathways in macaque monkey striate cortex: Anatomically defined columns in layer III. *Proc. Natl. Acad. Sci.*, *89*, 3566–3570.

Lee, B. B., Pokorny, J., Smith, V. C. and Kremers, J. (1994). Responses to pulses and sinusoids in macaque ganglion cells. *Vis. Res.*, *34*, 3081–3096.

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

Lisberger, S. G. and Sejnowski, T. J. (1992). Motor learning in a recurrent network model based on the vestibulo-ocular reflex. *Nature*, *360*, 159–161.

Liu, Z., Gaska, J. P., Jacobson, L. D. and Pollen, D. A. (1992). Interneuronal interaction between members of quadrature phase and anti-phase pairs in the cat's visual cortex. *Vis. Res.*, *7*, 1193–1198.

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

Maffei, L. and Fiorentini, A. (1976). The unresponsive regions of visual cortical receptive fields. *Vis. Res.*, *16*, 1131–1139.

Maffei, L., Fiorentini, A. and Bisti, S. (1973). Neural correlate of perceptual adaptation to gratings. *Science*, *182*, 1036–1038.

Maffei, L., Morrone, C., Pirchio, M. and Sandini, G. (1979). Responses of visual cortical cells to periodic and nonperiodic stimuli. *J. Physiol. (London)*, *296*, 27–47.

Malpeli, J. G., Schiller, P. H. and Colby, C. L. (1981). Response properties of single cells in monkey striate cortex during reversible inactivation of individual lateral geniculate laminae. *J. Neurophysiol.*, *46*, 1102–1119.

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

Marr, D. (1982). *Vision*. San Francisco: W. H. Freeman and Co.

Mastronarde, D. N. (1987). Two classes of single-input X-cells in cat lateral geniculate nucleus. I. receptive field properties and classification of cells. *J. Neurophysiol.*, *57*, 357–380.

Maunsell, J. H. R. and Gibson, J. R. (1992). Visual response latencies of striate cortex of the macaque monkey. *J. Neurophysiol.*, *68*, 1332–1344.

McLean, J. and Palmer, L. A. (1989). Contribution of linear spatiotemporal receptive field structure to velocity selectivity of simple cells in area 17 of cat. *Vis. Res.*, *29*, 675–679.

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

Morrone, M. C., Burr, D. C. and Maffei, L. (1982). Functional implications of cross-orientation inhibition of cortical visual cells. 1. Neurophysiological evidence. *Proc. R. Soc. Lon. B*, *216*, 335–354.

Movshon, J. A., Hawken, M. J., Kiorpes, L., Skoczenski, A. M., Tang, C. and O'Keefe, L. P. (1994). Visual noise masking in macaque LGN neurons. *Inv. Opht. and Vis. Sci. (Suppl.)*, *35*, 1662.

Movshon, J. A. and Lennie, P. (1979). Pattern-selective adaptation in visual cortical neurones. *Nature*, *278*, 850–852.

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. *J. Physiol. (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. *J. Physiol. (London)*, *283*, 101–120.

Nealey, T. A. and Maunsell, J. H. (1994). Magnocellular and parvocellular contributions to the responses of neurons in macaque striate cortex. *J. Neurosci.*, *14*, 2069–79.

Nelson, J. I. and Frost, B. (1985). Intracortical facilitation amoung co-oriented, co-axially aligned simple cells in cat striate cortex. *Exp. Br. Res.*, *6*, 54–61.

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

Nelson, S. B. (1991). Temporal interactions in the cat visual system i. Orientation-selective suppression in visual cortex. *J. Neurosci.*, *11*, 344–356.

Nestares, O. and Heeger, D. J. (1996). Modeling the apparent frequency-specific suppression in simple cell responses. *Vis. Res.* Submitted.

Ohzawa, I. and Freeman, R. D. (1986a). The binocular organization of complex cells in the cat's visual cortex. *J. Neurophysiol.*, *56*, 243–259.

Ohzawa, I. and Freeman, R. D. (1986b). The binocular organization of simple cells in the cat's visual cortex. *J. Neurophysiol.*, *56*, 221–242.

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

Ohzawa, I., Sclar, G. and Freeman, R. D. (1985). Contrast gain control in the cat's visual system. *J. Neurophysiol.*, *54*, 651–667.

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

Pollen, D. and Ronner, S. (1981). Phase relationships between adjacent simple cells in the visual cortex. *Science*, *212*, 1409–1411.

Pollen, D. and Ronner, S. (1982). Spatial computation performed by simple and complex cells in the visual cortex of the cat. *Vis. Res.*, *22*, 101–118.

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

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

Reid, R. C., Soodak, R. E. and Shapley, R. M. (1987). Linear mechanisms of directional selectivity in simple cells of cat striate cortex. *Proc. Natl. Acad. Sci.*, *84*, 8740–8744.

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

Reid, R. C., Victor, J. D. and Shapley, R. M. (1992). Broadband temporal stimuli decrease the integration time of neurons in cat striate cortex. *Vis. Neurosci.*, *9*, 39–45.

Robson, J. G. (1988). Linear and nonlinear operations in the visual system. *Inv. Opht. and Vis. Sci. (Suppl.)*, *29*, 117.

Rose, D. (1977). On the arithmetical operation performed by inhibitory synapses onto the neuronal soma. *Exp. Br. Res.*, *28*, 221–223.

Salin, P. A. and Prince, D. A. (1996a). Electrophysiological mapping of $GABA_A$ receptor-mediated inhibition in adult rat somatosensory cortex. *J. Neurophysiol.*, *75*, 1589–1599.

Salin, P. A. and Prince, D. A. (1996b). Spontaneous $GABA_A$ receptor-mediated inhibitory currents in adult rat somatosensory cortex. *J. Neurophysiol.*, *75*, 1573–1588.

van Santen, J. P. H. and Sperling, G. (1985). Elaborated Reichardt detectors. *J. Opt. Soc. Am. A*, *2*, 300–321.

Saul, A. B. (1995). Adaptation aftereffects in single neurons of cat visual cortex: response timing is retarded by adapting. *Vis. Neurosci.*, *12*, 191–205.

Saul, A. B. and Humphrey, A. L. (1990). Spatial and temporal response properties of lagged and nonlagged cells in cat lateral geniculate nucleus. *J. Neurophysiol.*, *64*, 206–224.

Saul, A. B. and Humphrey, A. L. (1992). Evidence for input from nonlagged cells in the lateral geniculate nucleus to simple cells in cortical area 17 of the cat. *J. Neurophysiol.*, *68*, 1190–.

Schiller, P. H., Finlay, B. L. and Volman, S. F. (1976). Quantitative studies of single-cell properties in monkey striate cortex. I. Spatiotemporal organization of receptive fields. *J. Neurophysiol.*, *39*, 1288–1319.

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

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

Sclar, G., Lennie, P. and DePriest, D. D. (1989). Contrast adaptation in striate cortex of macaque. *Vis. Res.*, *29*, 747–755.

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

Sengpiel, F. and Blakemore, C. (1994). Interocular control of neuronal responsiveness in cat visual cortex. *Nature*, *368*, 847–850.

Sengpiel, F., Blakemore, C. and Harrad, R. (1995). Interocular suppression in the primary visual cortex: a possible neural basis of binocular rivalry. *Vis. Res.*, *35*, 179–196.

Shapley, R. and Enroth-Cugell, C. (1984). Visual adaptation and retinal gain control. *Progress in Retinal Research*, *3*, 263–346.

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

Shapley, R. M. and Victor, J. D. (1978). The effect of contrast on the transfer properties of cat retinal ganglion cells. *J. Physiol.*, *285*, 275–298.

Sherman, S. M., Schumer, R. A. and Movshon, J. A. (1984). Functional cell classes in the macaque's LGN. *Soc. Neurosci. Abs.*, *10*, 296.

Sillito, A. M. (1975). The contribution of inhibitory mechanisms to the receptive field properties of neurones in the cat's striate cortex. *J. Physiol. (London)*, *250*, 304–330.

Sillito, A. M. (1977). Inhibitory processes underlying the directional specificity of simple, complex, and hypercomplex cells in the cat's visual cortex. *J. Physiol. (London)*, *271*, 699–720.

Sillito, A. M. (1984). Functional considerations of the operation of gabaergic inhibitory processes in the visual cortex. In E. G. Jones and A. Peters (Eds.), *Cerebral Cortex. vol 2. Functional Properties of Cortical Cells* (pp. 91–117). New York: Plenum Press.

Sillito, A. M., Grieve, K. L., Jones, H. E., Cudeiro, J. and Davis, J. (1995). Visual cortical mechanisms detecting focal orientation discontinuities. *Nature*, *378*, 492–496.

Sillito, A. M., Kemp, J. A., Milson, J. A. and Berardi, N. (1980). A reevaluation of the mechanisms underlying simple cell orientation selectivity. *Br. Res.*, *194*, 517–520.

Simoncelli, E. P. and Heeger, D. J. (1994). A velocity representation model for MT cells. *Inv. Opht. and Vis. Sci. (Suppl.)*, *35*, 1827.

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

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

Sperling, G. and Sondhi, M. M. (1968). Model for visual luminance discrimination and flicker detection. *J. Opt. Soc. Am. A*, *58*, 1133–1145.

Stafstrom, C. E., Schwindt, P. C. and Crill, W. E. (1984). Repetitive firing in layer V neurons from cat neocortex in vitro. *J. Neurophysiol.*, *52*, 264–277.

Stratford, K., Mason, A., Larkman, A., Major, G. and Jack, J. (1990). The modeling of pyramidal neurons in the visual cortex. In R. Durbin, C. Miall and G. Mitchison (Eds.), *The Computing Neuron* (pp. 296–321). England: Addison-Wesley.

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

Tadmor, Y. and 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. *Vis. Neurosci.*, *3*, 445–454.

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

Tolhurst, D. J. and Dean, A. F. (1987). Spatial summation by simple cells in the striate cortex of the cat. *Exp. Br. Res.*, *66*, 607–620.

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

Tolhurst, D. J. and Dean, A. F. (1991). Evaluation of a linear model of directional selectivity in simple cells of the cat's striate cortex. *Vis. Neurosci.*, *6*, 421–428.

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. *Vis. Neurosci.*, (p. In press).

Tolhurst, D. J. and Heeger, D. J. (1996b). Contrast normalization and hard threshold models of the responses of simple cells in cat striate cortex. *Vis. Neurosci.*, (p. In press).

Toyama, K., Kimura, M., Shiida, T. and Takeda, T. (1977a). Convergence of retinal inputs onto visual cortical cells: II. A study of the cells disynaptically excited from the lateral geniculate body. *Br. Res.*, *137*, 221–231.

Toyama, K., Maikawa, K. and Tanaka, T. (1977b). Convergence of retinal inputs onto visual cortical cells: I. A study of the cells monosynaptically excited from the lateral geniculate body. *Br. Res.*, *137*, 207–220.

Toyama, K., Matsunami, K., Ohno, T. and Tokashiki, S. (1974). An intracellular study of neuronal organization in the visual cortex. *Exp. Br. Res.*, *21*, 45–66.

Toyama, K. and Takeda, T. (1974). A unique class of cat's visual cortical cells that exhibit either ON or OFF excitation for stationary light slits and are responsive to moving edge patterns. *Br. Res.*, *73*, 350–355.

Tranchina, D., Gordon, J. and Shapley, R. M. (1984). Retinal light adaptation — evidence for a feedback mechanism. *Nature*, *310*, 314–316.

Troy, J. B. (1983). Spatial contrast sensitivities of X and Y type neurones in the cat's dorsal lateral geniculate nucleus. *J. Physiol. (London)*, *344*, 399–417.

Ulinski, P. S. and Fowler, M. (1995). Feedforward and feedback inhibition in visual cortex. *Inv. Opht. and Vis. Sci. (Suppl.)*, *36*, S1052.

Victor, J. (1987). The dynamics of the cat retinal X cell centre. *J. Physiol. (London)*, *386*, 219–246.

Victor, J. (1988). The dynamics of the cat retinal Y cell subunit. *J. Physiol. (London)*, *405*, 289–320.

Walker, G. A., Ohzawa, I. and Freeman, R. D. (1996). Interocular transfer of cross-orientation suppression in the cat's visual cortex. *Inv. Opht. and Vis. Sci. (Suppl.)*, *37*, S485.

Watanabe, S., Konishi, M. and Creutzfeldt, O. D. (1966). Postsynaptic potentials in the cat's visual cortex following electrical stimulation of afferent pathways. *Exp. Br. Res.*, *1*, 272–283.

Watson, A. B. and Ahumada, A. J. (1985). Model of human visual-motion sensing. *J. Opt. Soc. Am. A*, *2*, 322–342.

Yoshioka, T., Levitt, J. B. and Lund, J. (1994). Independence and merger of thalamocortical channels within macaque monkey primary visual cortex: anatomy of interlaminar projections. *Vis. Neurosci.*, *11*, 467–489.

Figure Captions

1. Schemata of the two models of simple cell function that are discussed in this Chapter. A: The *linear model*, composed of a linear *spatiotemporal weighting function* and of a *rectification* stage. The cell performs a weighted average of the light intensities over local space and recent time. The rectification stage converts the output of this linear stage into firing rate. Note that rectification is a nonlinearity, so the "linear model" is not entirely linear. B: The *normalization model*, which extends the linear model by adding a divisive stage. The output of each cell's linear stage is divided by the pooled output of a large number of other cells.

2. A stimulus and a receptive field in space-time. A: A vertical bar translating to the right. B: The space-time volume of stimulus intensities corresponding to motion of the vertical bar. C: An x-t slice through the space-time volume. Orientation in the x-t slice is the horizontal component of velocity. D: an x-y section of a spatiotemporal weighting function. Dark areas represent locations where the weighting function is negative, bright areas represent locations where it is positive. E: The same x-y section together with a projection of the weighting function on the x-t plane. F: The x-t projection, which ignores the y dimension. The receptive field travels in time t, from past to future. Panels A-C are based on an illustration by Adelson and Bergen (1985).

3. Direction selectivity in a linear cell. A: Responses to gratings drifting in the preferred direction. The stimulus elicits large responses. B-D: Relative space-time positions of weighting function and stimulus at three instants in time. When the central excitatory region of the weighting function is aligned with a dark bar (B), the response is negative. When it is aligned with a bright bar, the response is positive (C), and so on. E: Responses to gratings drifting in the opposite direction. The stimulus elicits small responses. F-H: Relative space-time positions of weighting function and stimulus at three instants in time. At any given time, each bar of the grating is covering both excitatory and inhibitory subregions of the weighting function, whose outputs are averaged out by the

cell. The filled areas in Panels A and E show the parts of the responses that would be visible after rectification.

4. Contour plots of three spatiotemporal weighting functions, averaged across one spatial dimension. Continuous curves: positive contours. Dashed curves: negative contours. A: A space-time separable weighting function, the product of a center-surround spatial weighting function and of a monophasic temporal weighting function. B: Another space-time separable weighting function, spatially displaced with respect to A and with a biphasic temporal weighting function that is less delayed than the one in A. C: Non-separable weighting function obtained by summing the ones in A and B. The result is a space-time oriented weighting function resembling that of the direction selective cell in Figure 7.

5. Three possible transformations of membrane potential into firing rate. For simplicity the resting potential is assigned the value $V_{rest} = 0$. The continuous lines represent *rectification* with threshold V_{thresh}. The thick, intermediate and thin lines represent respectively the cases in which the threshold is $V_{thresh} = 0, 5$ and 10 mV. The dashed lines represent approximations to rectification. These approximations are useful to simplify the mathematics of the normalization model (Appendix A). They are power functions of the positive deviation from resting potential, for different exponents n. The exponent n is 2 for the thicker dashed curve, 3 for the thinner dashed curve.

6. Flashed bar stimuli and ON- or OFF- subregions of a linear cell's receptive field. A: Spatiotemporal structure of a flashed bar stimulus. The x-t projection of a flashed bar is a vertical rectangle. B: Response of a linear cell when the bar is flashed on an OFF subregion of the receptive field. The response is negative when the bar is turned on, positive when it is turned off. The left panel shows the responses as a function of time, before rectification. The other three panels represent receptive field and stimulus at three instants in time. C: as in B, except that the bar is flashed on an ON subregion of its receptive field. The filled areas in the leftmost panel in B and C show the parts of the responses that would be visible after rectification.

7. The full space-time receptive field of a simple cell, as obtained with the reverse correlation method. The four upper panels represent x-y snapshots of the receptive field measured at different times T in the past. Gray levels indicate the correlation between the appearance of a bar and the firing rate T ms later. Zero correlation is indicated by mid-gray. Lighter grays indicate points of positive correlation with the appearance of a bright bar. Darker grays indicate points of positive correlation with the appearance of a dark bar. A large number of snapshots like these are stacked to build a full space-time receptive field, whose space-time projection is shown at the bottom. This previously unpublished Figure is courtesy of Greg DeAngelis. Cell is part of sample published in (DeAngelis et al., 1993a).

8. Linearity of spatial summation in four cat V1 simple cells. Spatial weighting functions as measured with flashing bars (*histograms*) and as predicted by inverse Fourier transformation of the spatial frequency tuning curves (*continuous curves*). For a linear cell the two would be identical. Plots show one spatial dimension (e.g. x), and collapse all information about the other two dimensions (y and t). Both the observed and predicted weighting functions were independently rescaled. Positive values in each weighting function represent incremental responses to the introduction of a bright bar; negative values represent incremental responses to the introduction of a dark bar. *Insets*: The spatial frequency tuning curves used to compute each predicted weighting function. The abscissa of these insets is spatial frequency (in cycles/degree) and the ordinate is contrast sensitivity, the inverse of the threshold contrast value for each spatial frequency. Reprinted with permission from (Movshon et al., 1978a).

9. Linearity of spatio-temporal summation in a cat V1 simple cell. The polar plot shows the cell's responses to standing gratings whose contrasts were modulated sinusoidally in time at different spatial phases. The amplitude of the first harmonic sinusoid of each response is represented radially, while the angular coordinate indicates the temporal phase of each response. The filled symbols represent the unaltered data from the experiment; the open symbols and the ellipse fitted to them represent the same data corrected for a resting

"firing rate" of -8 Spikes/sec. Reprinted with permission from (Movshon et al., 1978a).

10. Simplified model of a cortical cell, and possible biophysical implementation of the linear model. The cell membrane is modeled as a single compartment with passive properties and two classes of synaptic inputs, excitatory and inhibitory. In the central excitatory subregion of the receptive field the excitation is provided by ON-center cells and the inhibition by OFF-center cells with superimposed receptive fields. The flanking inhibitory subregions are obtained by the opposite arrangement of excitation and inhibition (not shown). This *push-pull* arrangement of excitation and inhibition ensures the linearity of the membrane potential V. The membrane potential is encoded into firing rate by a rectifier. See text for explanation of symbols.

11. Encoding of input current in a visual cortical cell *in vitro*. Responses of an intracellularly recorded regular spiking neuron in a slice of guinea pig cortex to sinusoidal current injection (0.8 nA). A: Time course of the responses for five different frequencies of stimulation (1, 2, 4, 8 and 16 Hz). The traces are dominated by their first harmonic. This means that apart from the presence of the spikes the membrane is acting linearly. B: Temporal frequency tuning of spike rate and membrane potential. Filled symbols: Amplitude of the first harmonic of the membrane potential obtained by fitting a sinusoid to the raw membrane potential traces in A. Scale is on right axis. The dashed line is the prediction of a single compartment with only passive conductances. Open symbols: Amplitude of the first harmonic of the firing rate, obtained by fitting a sinusoid to the spike times. This Figure was first presented in Carandini et al. (1994).

12. Responses of a monkey V1 simple cell to a drifting sine grating for three different spatial frequencies and three different contrasts. The curves are fits of the normalization model. The fits were performed on a larger data set, which included the responses to 45 different drifting gratings, that had 5 different contrasts, 3 different spatial frequencies, and 3 different temporal frequencies. These stimuli were randomly interleaved to minimize the effect of visual adaptation. A: Spike histograms of one period of the responses, averaged over many presentations. The three columns show the responses to drifting gratings with

a spatial frequency of 1, 0.6 and 0.4 cycles/degree (cpd). Each row corresponds to one of three different contrasts: 25, 50 and 100%. B: Amplitude of the responses as a function of contrast. The ordinate plots the amplitude of the first harmonic of the same responses as Panel A. Error bars indicate the standard error of the mean (N=3). The number next to each curve specifies the spatial frequency of the stimulus. Experimental methods for this and for the following figures are outlined in (Carandini and Heeger, 1994).

13. Responses of a monkey V1 simple cell to a drifting sine grating for different stimulus orientations and contrasts. The continuous curves are fits of the normalization model. The fits were performed on a larger data set, which included the responses to additional temporal frequencies. Error bars indicate the standard error of the mean (N=3). A: Contrast responses for two different stimulus orientations. Changing the orientation of a grating shifts the contrast responses up and down on a logarithmic scale. B. Effect of contrast on the orientation tuning. Data for 40° and 80° are the same as those in panel A. The orientation tuning is invariant with contrast.

14. Responses of a monkey V1 simple cell to a plaid composed of two drifting gratings ("test" and "mask"), for different contrasts of the two gratings. The orientations of test and mask differ by 90°. The curves are fits of the normalization model. A: Spike histograms of one period of the responses, averaged over many presentations. Rows correspond to a fixed test contrast, columns to a fixed mask contrast. The mask does not elicit any overt response when presented alone (top row) and strongly inhibits the responses to the test (second and third columns) B: Amplitude of the responses as a function of contrast. The ordinate plots the amplitude of the first harmonic of the same responses as Panel A. Error bars indicate the standard error of the mean (N=3). The white, gray and black circles refer respectively to mask contrasts of 6, 25 and 50%.

15. Responses of a monkey V1 simple cell to a drifting sine grating for different contrasts and temporal frequencies. The curves are fits of the normalization model. The fits were performed on a larger data set, which included the responses to an additional orientation. A: Spike histograms of one period of the responses, averaged over many presentations.

Each panel corresponds to a different contrast of the stimulus. Note the prominent phase advance with increasing contrast. B: Amplitude of the responses as a function of contrast and temporal frequency. The ordinate plots the amplitude of the first harmonic responses such as those in Panel A. Data points are joined by dashed lines. The temporal frequency tuning is strongly dependent on contrast: at low contrasts the high frequencies are more attenuated than at high contrasts.

16. A possible biophysical implementation of the normalization model. The diagram is identical to the one for the linear model (Figure 10), except for the presence of a shunting conductance g_{shunt}. The equilibrium potential for the shunt is $V_{shunt} = V_{rest}$. In the normalization model the shunting conductance grows with the activity of a large number of cortical cells, the normalization pool. The shunt conductance is assumed to be the same for all the neurons in the pool. See Sections 6.1 and 6.2 for details.

17. A. Dependence of the shunt conductance on the overall activity of the normalization pool, as postulated by the normalization model. The precise functional form is $g_{shunt} = g_0(1/\sqrt{1 - k \sum R} - 1)$, where $\sum R$ is the overall firing rate of the normalization pool. This function is entirely *ad hoc* and does not have any experimental support. It was chosen so that the model would be mathematically tractable and yield the desired equations. B: Effect of the membrane conductance on the size and time course of the membrane potential responses. The three curves show the membrane potential responses of a model cell to a current step for three different values of the total conductance g. The origin of the abscissa is set at the time of the current step. The scale of the ordinate is linear. Increasing the conductance of the cell reduces both the gain and the latency. The arrows point to twice the time constant of the cell.

18. Responses of a monkey V1 simple cell to a plaid composed of two drifting gratings (*test* and *mask*), for different contrasts of the two gratings. The orientations of test and mask differed by 90°. Error bars indicate the standard error of the mean (N=3). The *curves* are fits of the normalization model. The fits were performed on a larger data set, which included all combinations of five test contrasts and five mask contrasts (23 of the 25 data

points are shown). A: Dependence of the response amplitude on the test contrast, for three different values of the mask contrast (white: 6%, gray: 25%, black: 50%). The leftmost data points show that the mask alone was able to elicit some responses. In these conditions the normalization model correctly predicts that increasing mask contrast does not simply shift the contrast responses to the right, as was the case in Figure 14, but it also changes their shape. B: Dependence of the response amplitude on the mask contrast, for four different values of the test contrast (white: 1%, light gray: 6%, dark gray: 12% black: 50%). Many data points are represented both in A and in B. The prediction of the model for the data point in the lower right of panel B is off by 2 Spikes/sec.

19. Responses of a monkey V1 simple cell to two individual gratings of different spatial frequency (1 and 0.5 cpd), and to the "plaid" obtained by summing the two. The continuous curves are fits of the normalization model. A: Spike histogram of the response to grating 1 alone, with contrast $c_1 = 25\%$. The histogram shows one period of the responses, averaged over many presentations. B: Spike histogram of the response to the plaid, with contrasts $c_1 = c_2 = 25\%$. C: Spike histogram of the response to grating 2 alone, with contrast $c_2 = 25\%$. *Dashed curve in B*: linear prediction, obtained by summing the sinusoids fitted in A and C. The linear prediction overestimates the plaid response. D: Polar plot of the first harmonic responses, for a variety of grating and plaid contrasts. Abscissa and ordinate are the cosine and sine components of the responses at the temporal frequency of the stimulus. Response amplitude is given by the distance from the origin. Response phase is given by the angle with the abscissa. Grating responses are in black, plaid responses in gray. The data points corresponding to the responses in A, B and C are surrounded by squares. The *dashed curve* shows the linear prediction obtained by summing the two fits to the grating responses. The dashed parallelogram illustrates the vectorial sum that gives rise to the linear prediction corresponding to the dashed line in B. The actual response is smaller and occurs slightly earlier than the linear prediction. Panel D is adapted from (Carandini and Heeger, 1995).

20. Estimated time constants for a population of cells. The estimates come from the fits

of the normalization model to 34 extracellular recording experiments on 21 monkey V1 simple cells. The experiments involved randomly interleaved grating stimuli varying in temporal frequency, contrast and orientation (or spatial frequency). *Abscissa* shows τ_0, the estimated time constant at rest, i.e. for zero contrast. *Ordinate* shows τ_1, the estimated time constant when the stimulus is a drifting grating with 100% contrast. The dashed diagonal line marks the identity $\tau_0 = \tau_1$. The bars on the two axes show the medians: 29 ms for the time constants at rest, 7.6 ms for the time constants at 100% contrast. The *continuous line* is a regression line fit to the data points with the constraint that it goes through the origin. Its equation is $\tau_1 = 0.334\tau_0$, implying a *threefold average decrease in time constant.*

A The linear model

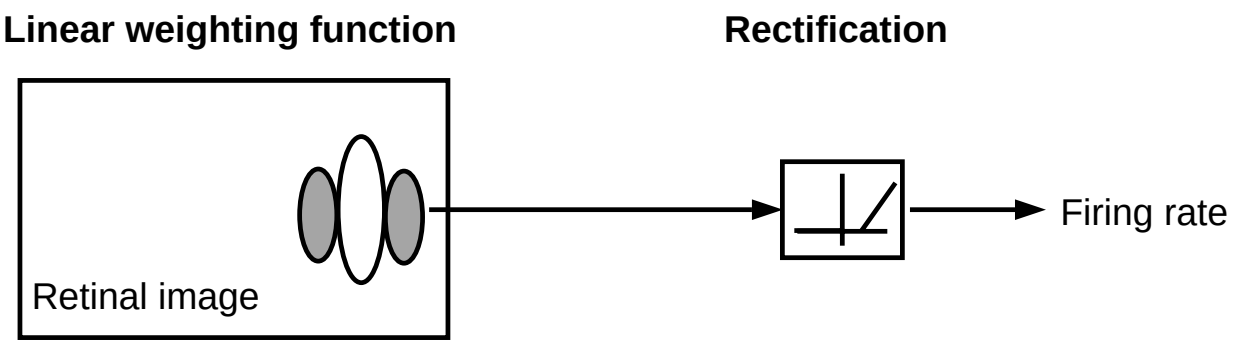

B The normalization model

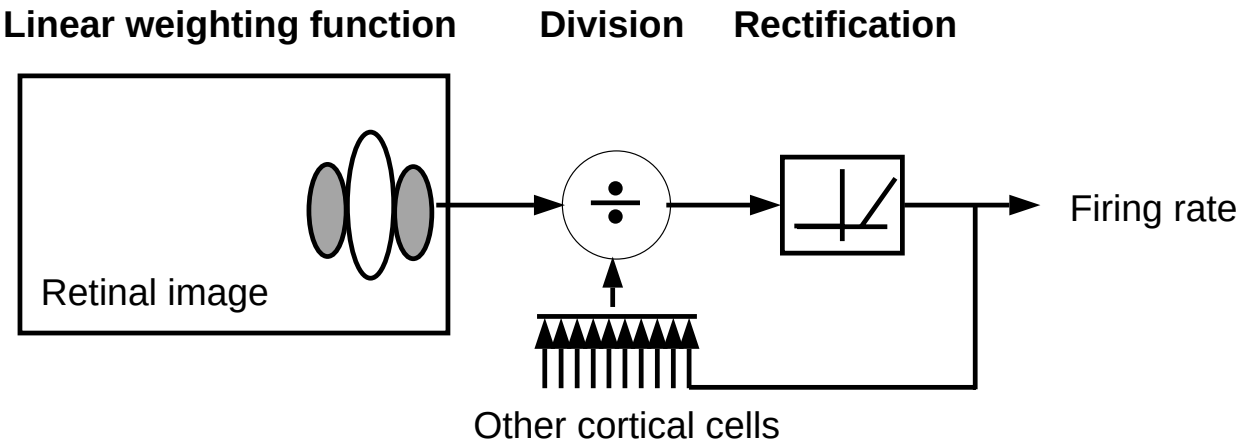

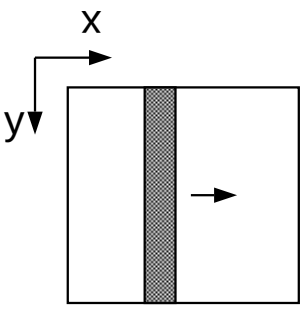

A

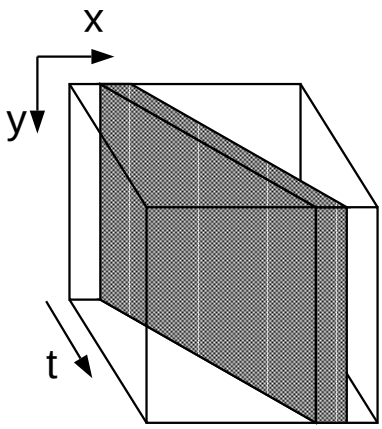

B

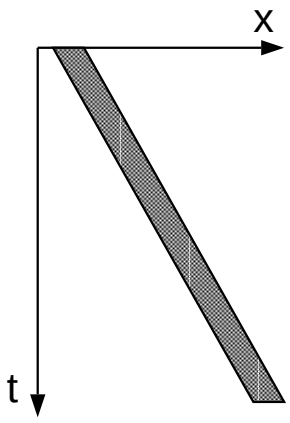

C

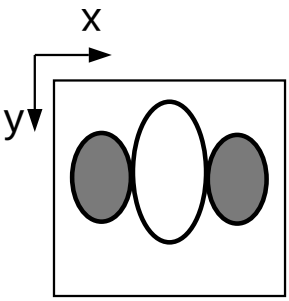

D

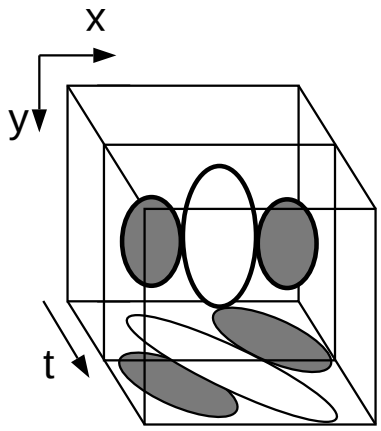

E

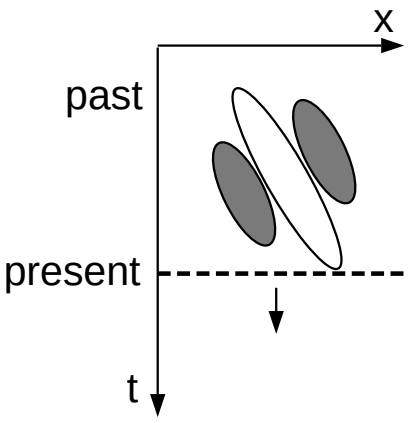

F

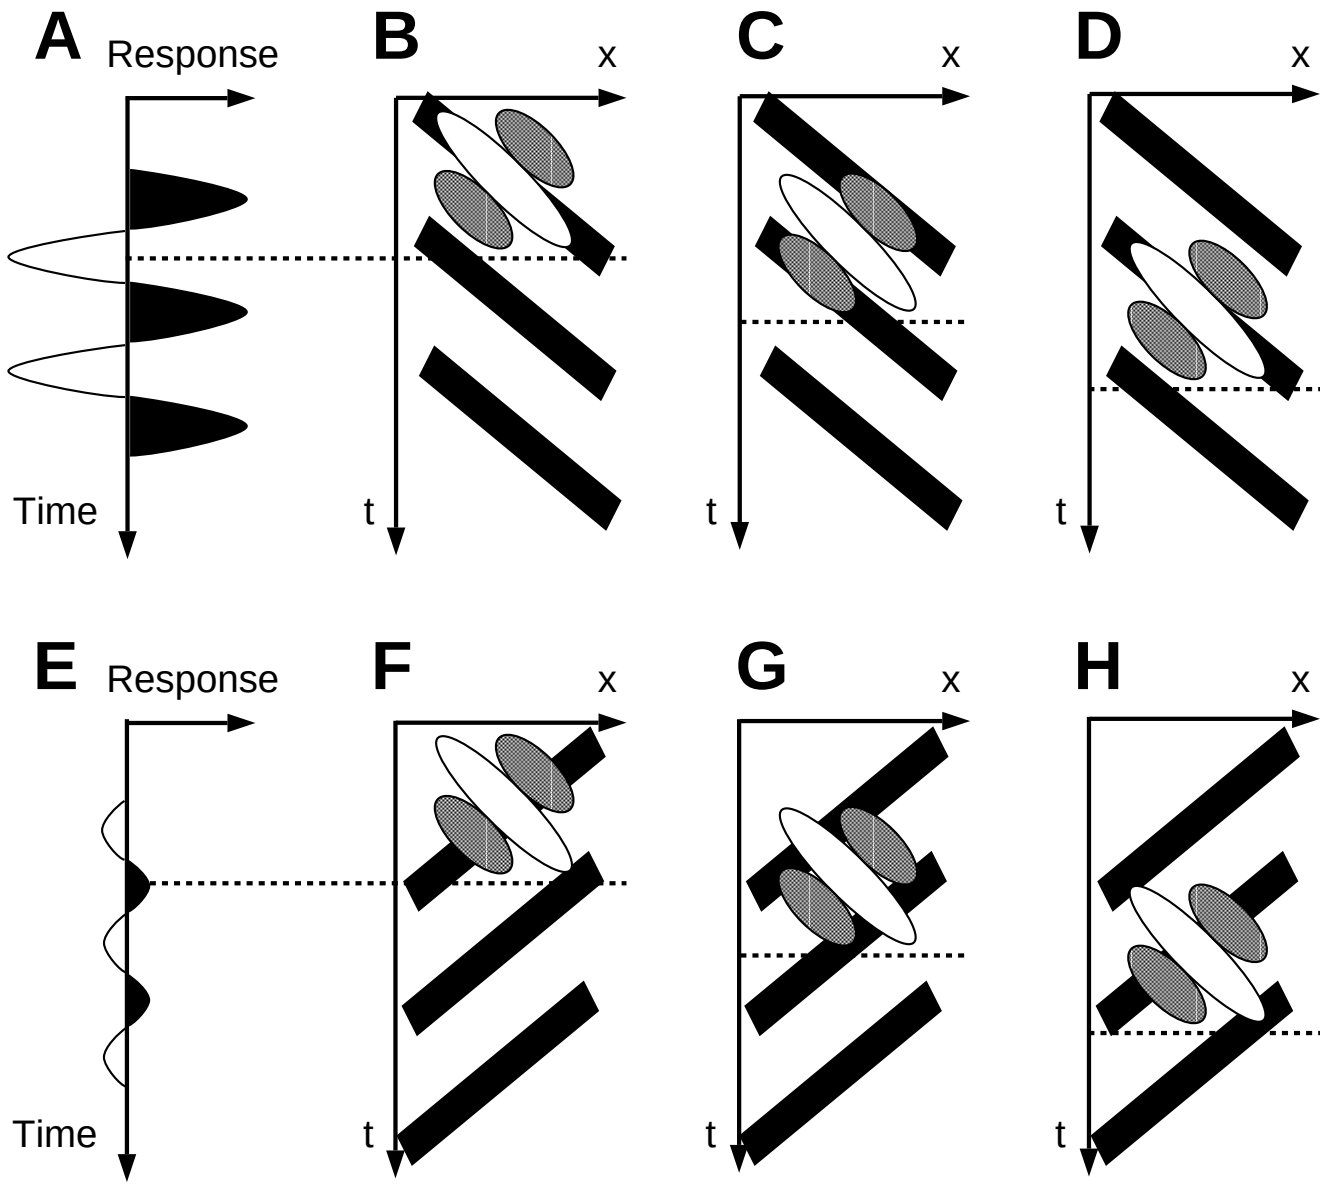
A
Response
Time
B
x
t
C
x
t
D
x
t
E
Response
Time
F
x
t
G
x
t
H
x
t

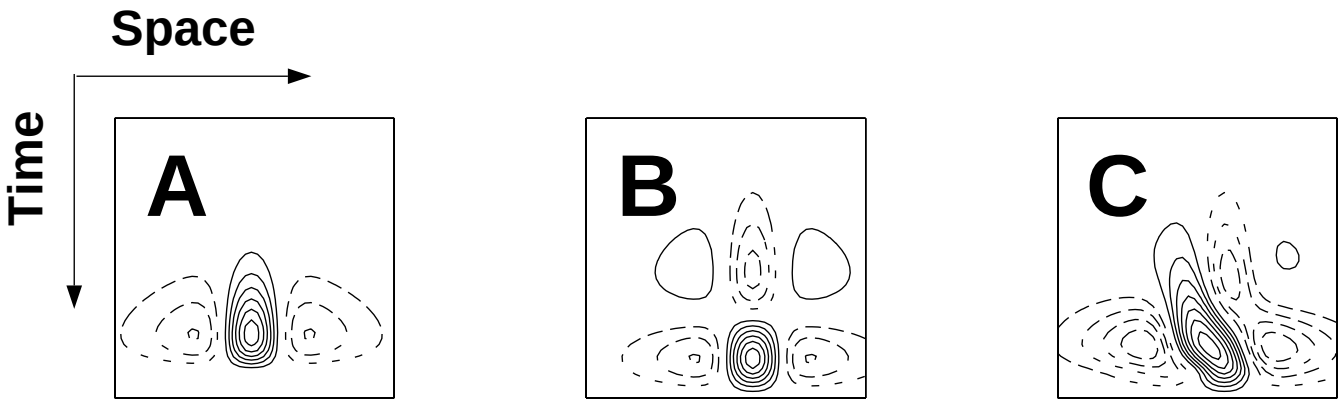
Space
Time
A
B
C

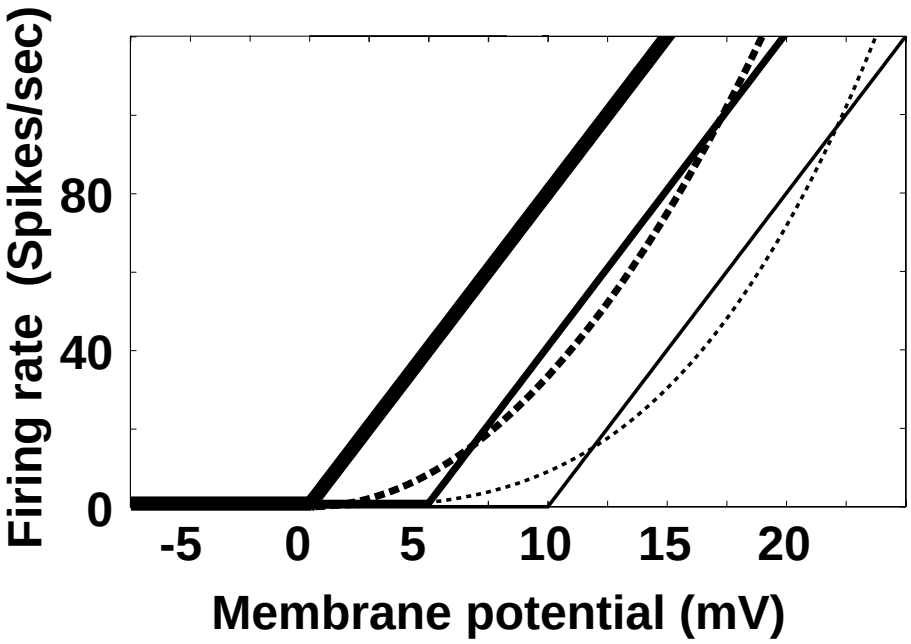
Firing rate (Spikes/sec)
80
40
0
-5
0
5
10
15
20
Membrane potential (mV)

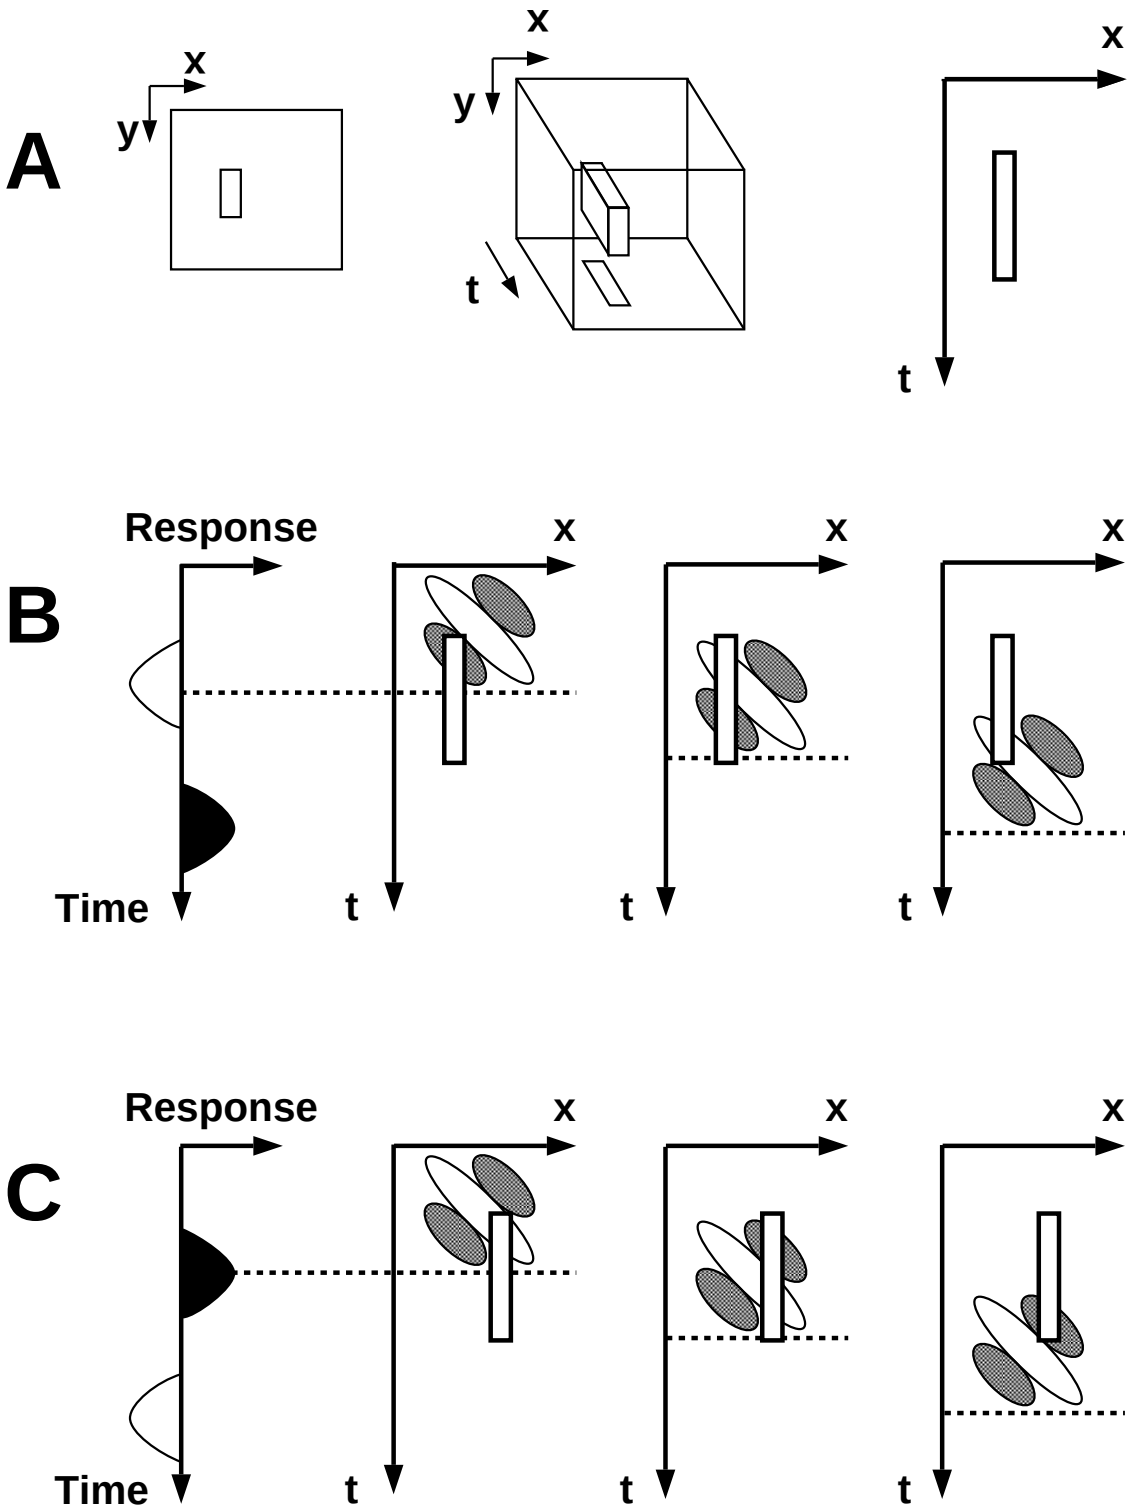
A
x
y
x
y
t
x
t
B
Response
x
Time
t
x
t
x
t
C
Response
x
Time
t
x
t
x
t

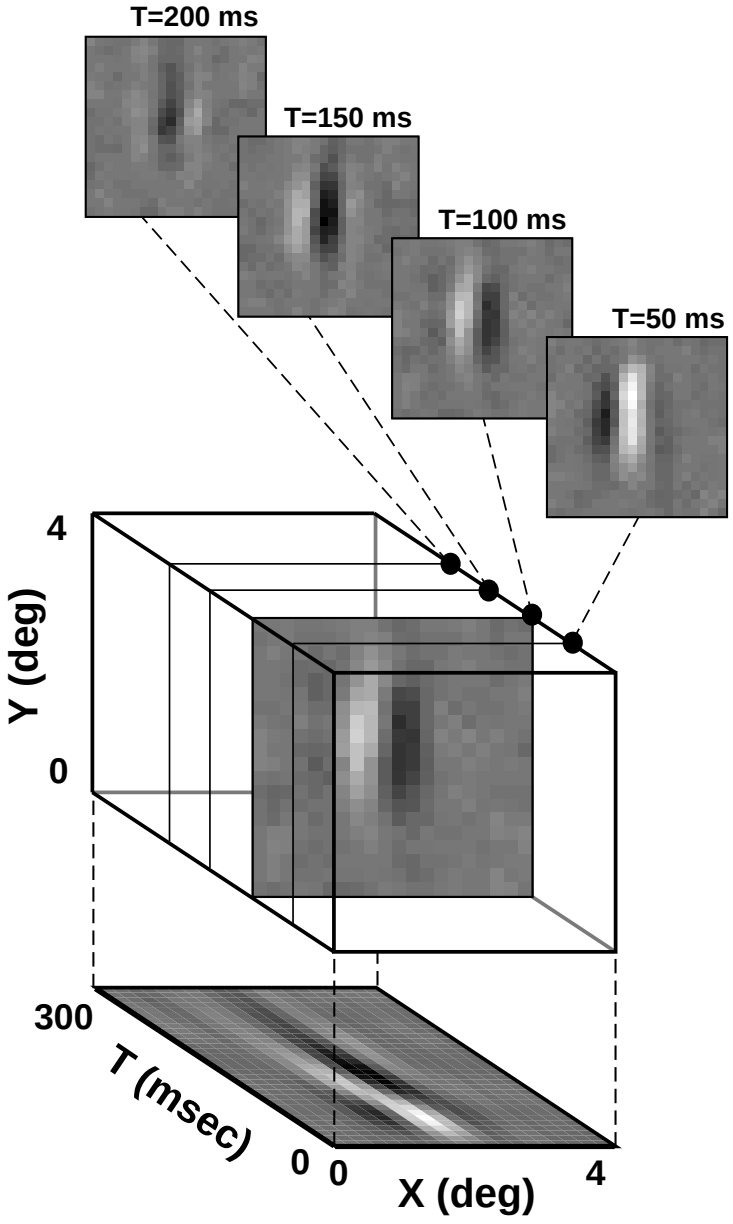
T=200 ms
T=150 ms
T=100 ms
T=50 ms
4
Y (deg)
0
300
T (msec)
0
0
4
X (deg)

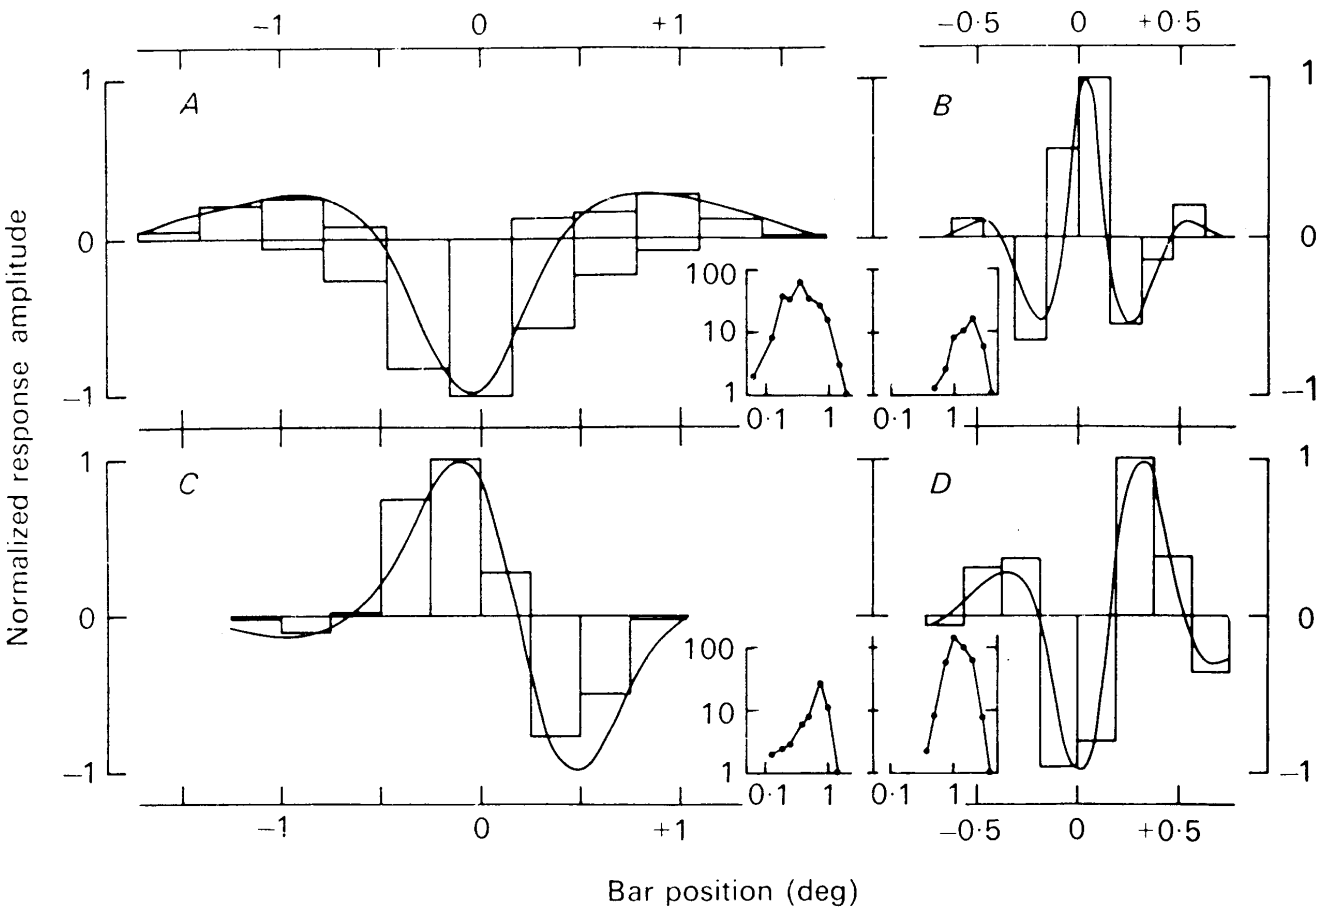
A
B
C
D
−1
0
+1
−0·5
0
+0·5
1
0
−1
100
10
1
0·1
1
Normalized response amplitude
Bar position (deg)

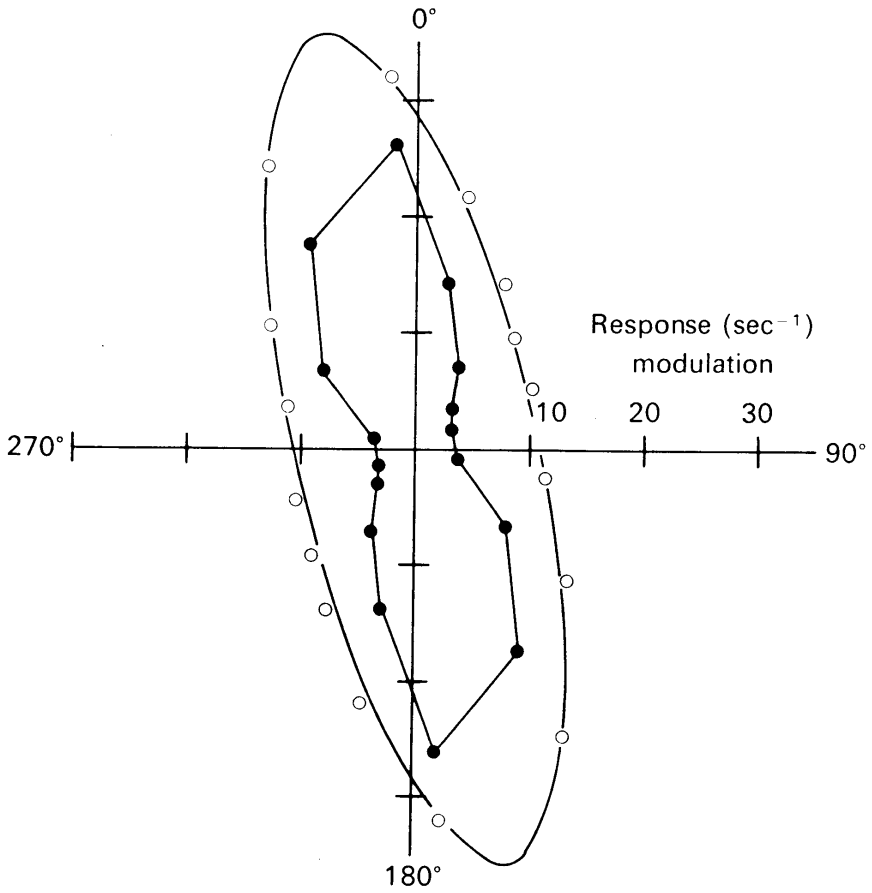
0°
Response (sec⁻¹)
modulation
10
20
30
270°
90°
180°

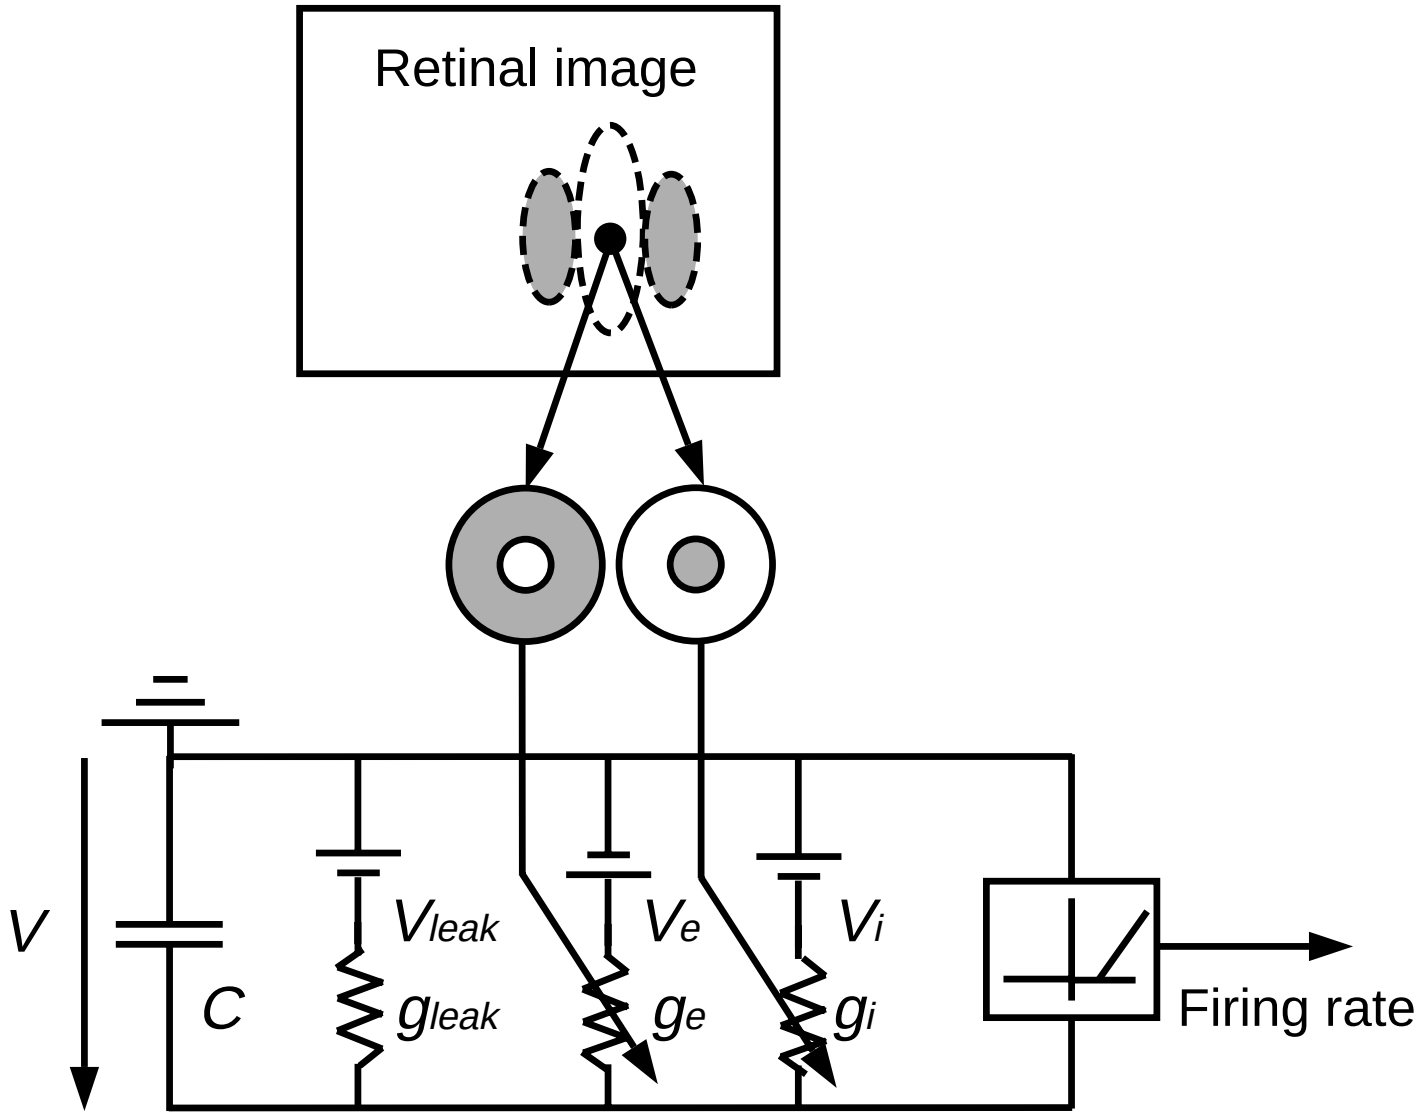
Retinal image
V
C
V_{leak}
g_{leak}
V_e
g_e
V_i
g_i
Firing rate

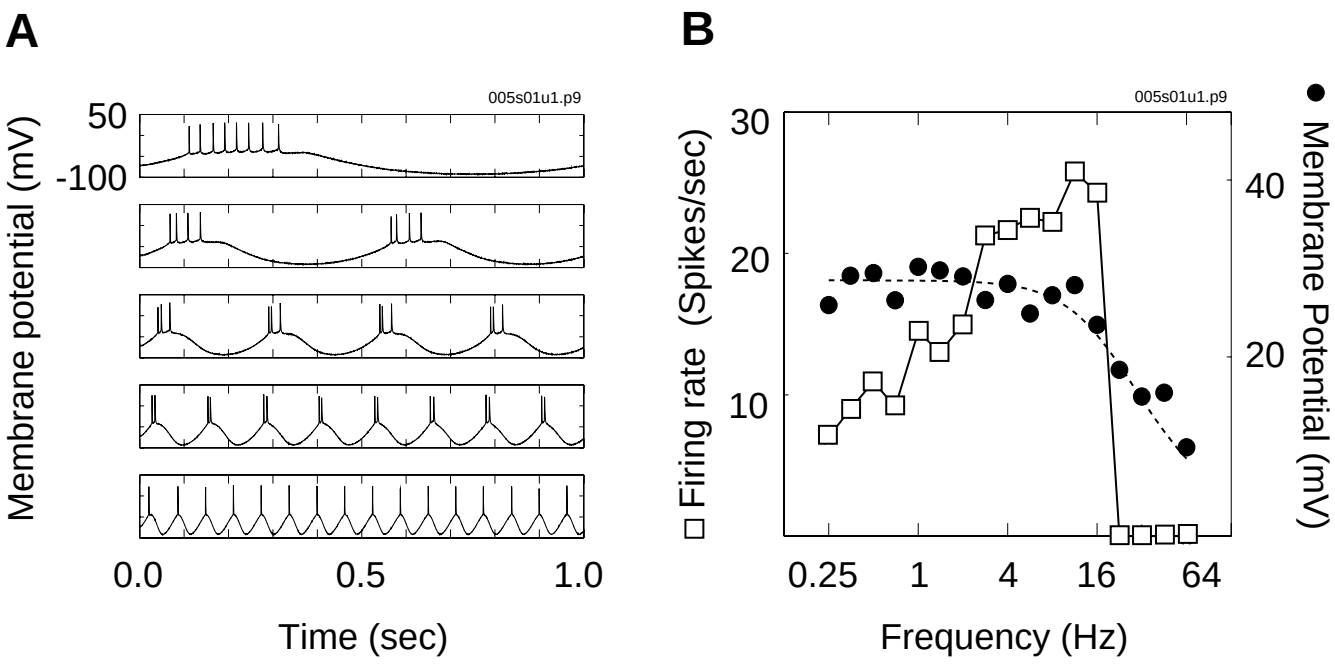
A
005s01u1.p9
50
-100
Membrane potential (mV)
0.0
0.5
1.0
Time (sec)
B
005s01u1.p9
30
20
10
□ Firing rate (Spikes/sec)
40
20
● Membrane Potential (mV)
0.25
1
4
16
64
Frequency (Hz)

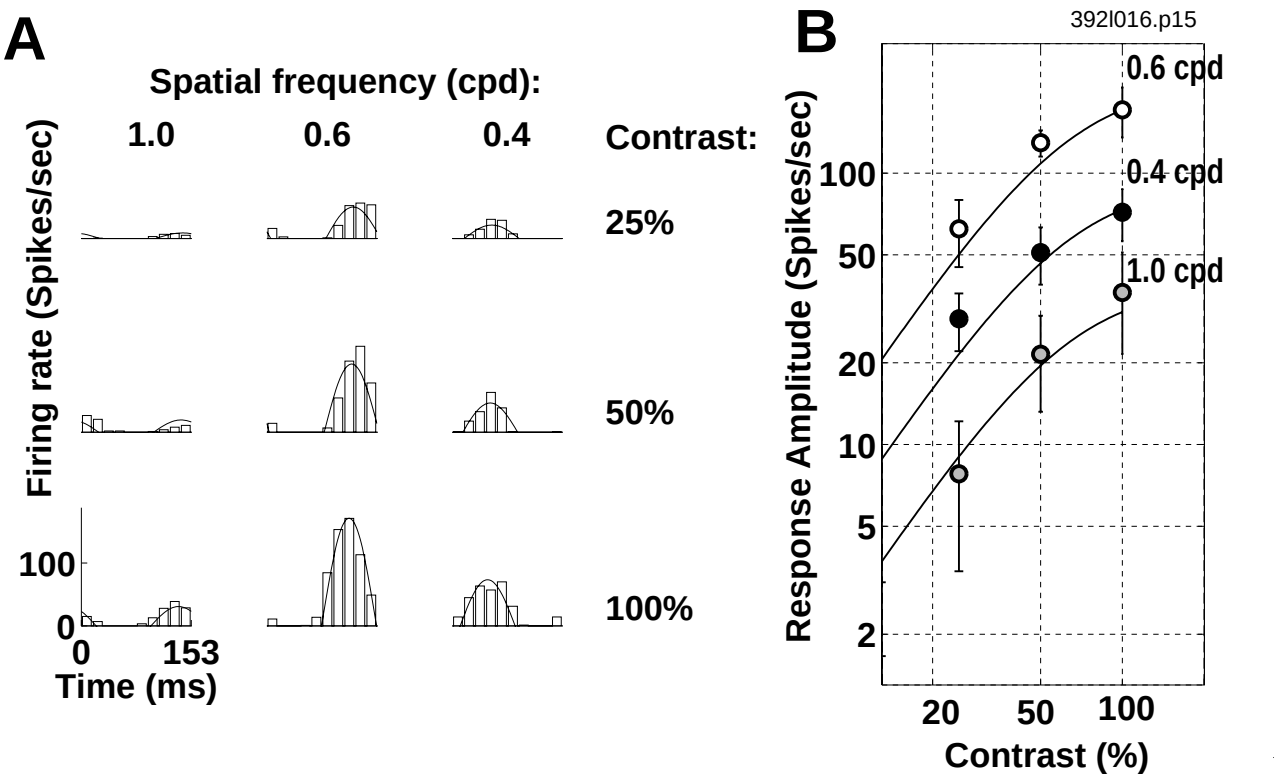
A
Spatial frequency (cpd):
1.0
0.6
0.4
Contrast:
25%
50%
100%
Firing rate (Spikes/sec)
100
0
0
153
Time (ms)
B
392l016.p15
0.6 cpd
0.4 cpd
1.0 cpd
Response Amplitude (Spikes/sec)
100
50
20
10
5
2
20
50
100
Contrast (%)

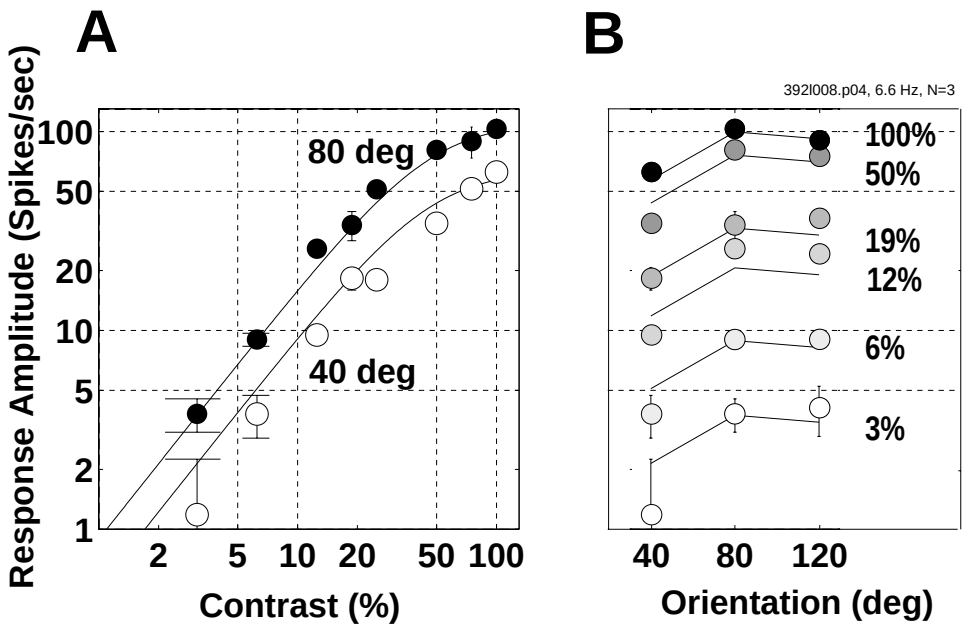
A
B
392l008.p04, 6.6 Hz, N=3
Response Amplitude (Spikes/sec)
100
50
20
10
5
2
1
80 deg
40 deg
2
5
10
20
50
100
Contrast (%)
100%
50%
19%
12%
6%
3%
40
80
120
Orientation (deg)

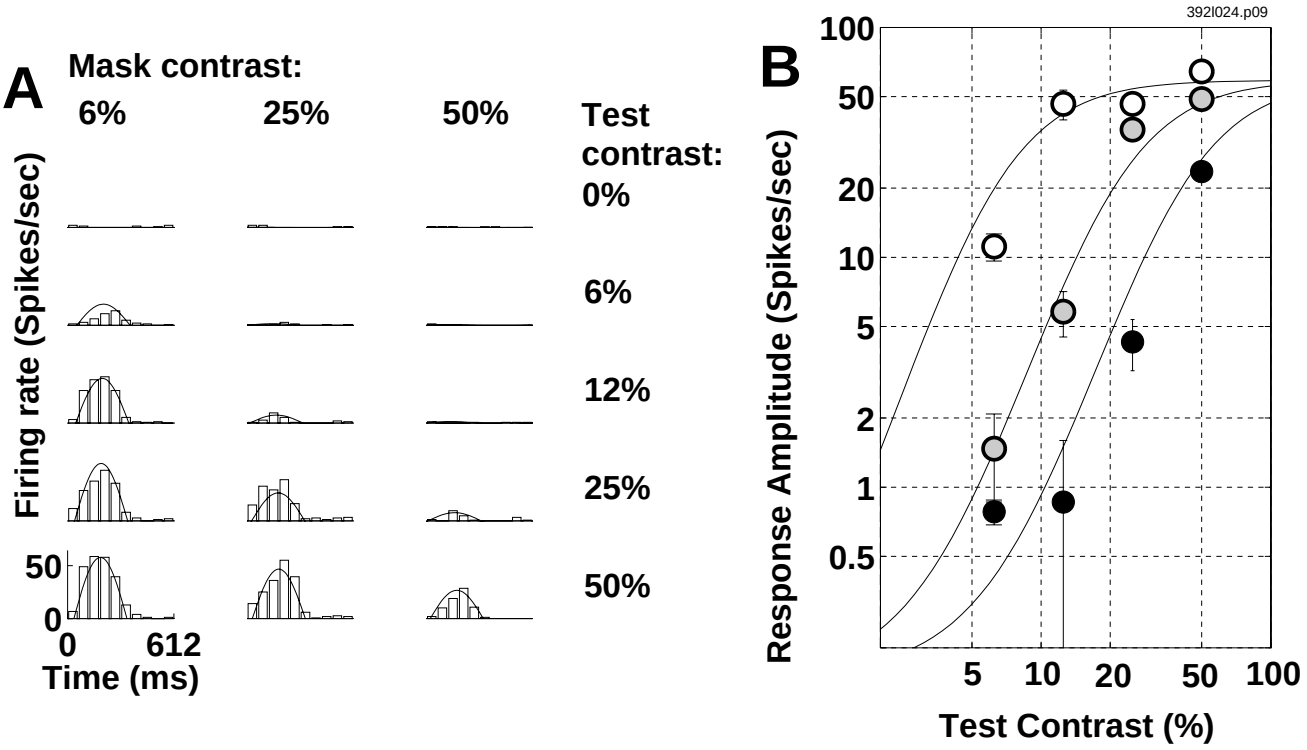
A
Mask contrast:
6%
25%
50%
Test contrast:
0%
6%
12%
25%
50%
Firing rate (Spikes/sec)
50
0
0
612
Time (ms)
B
392l024.p09
Response Amplitude (Spikes/sec)
100
50
20
10
5
2
1
0.5
5
10
20
50
100
Test Contrast (%)

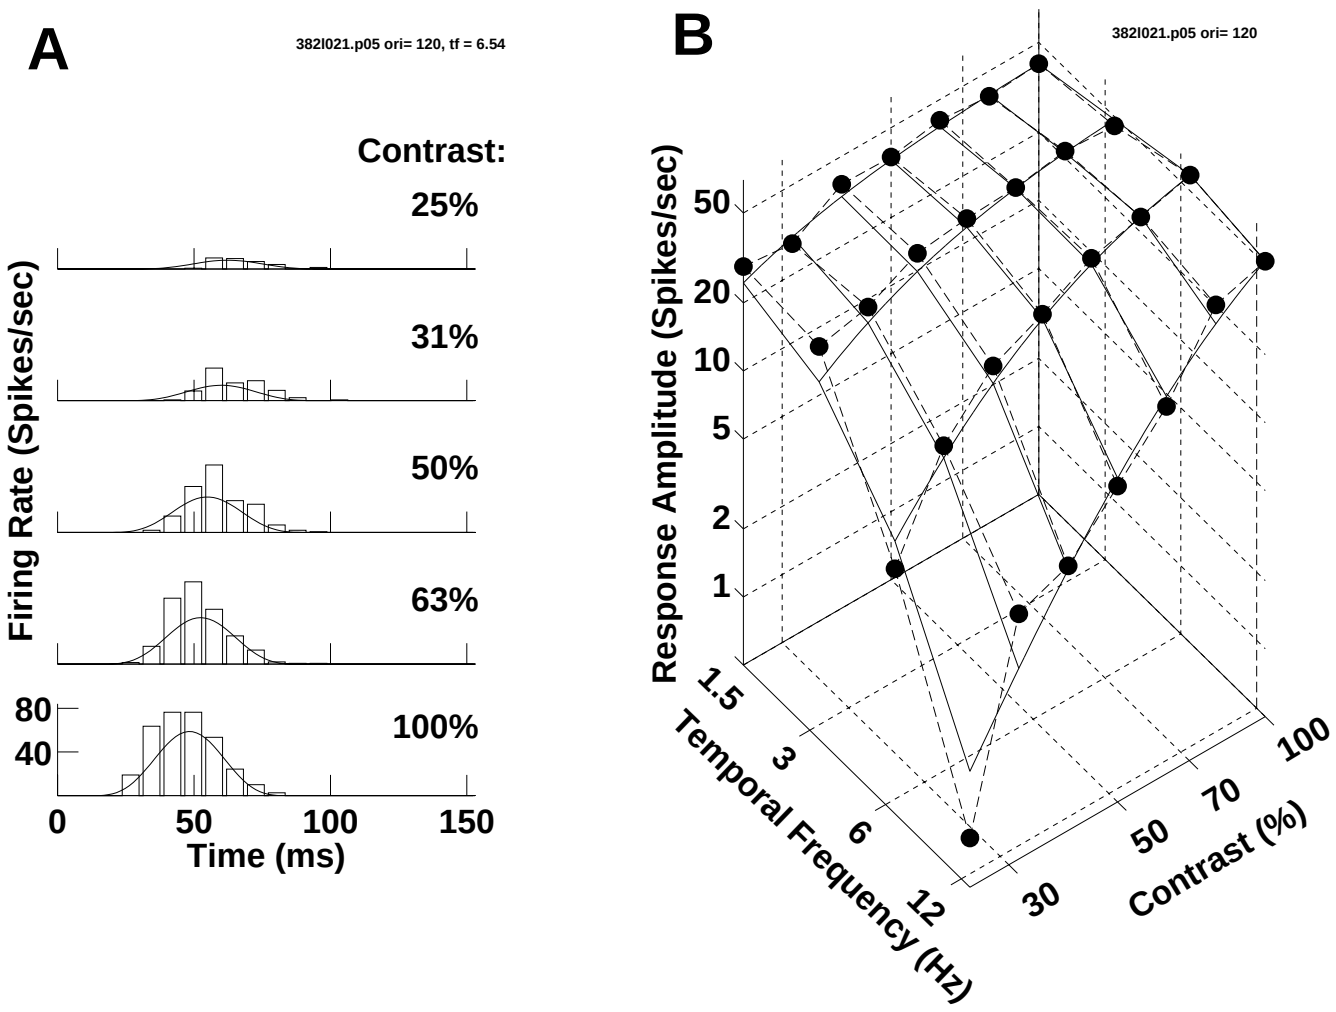
A
382l021.p05 ori= 120, tf = 6.54
Contrast:
25%
31%
50%
63%
100%
Firing Rate (Spikes/sec)
80
40
0
50
100
150
Time (ms)
B
382l021.p05 ori= 120
Response Amplitude (Spikes/sec)
50
20
10
5
2
1
1.5
3
6
12
Temporal Frequency (Hz)
30
50
70
100
Contrast (%)

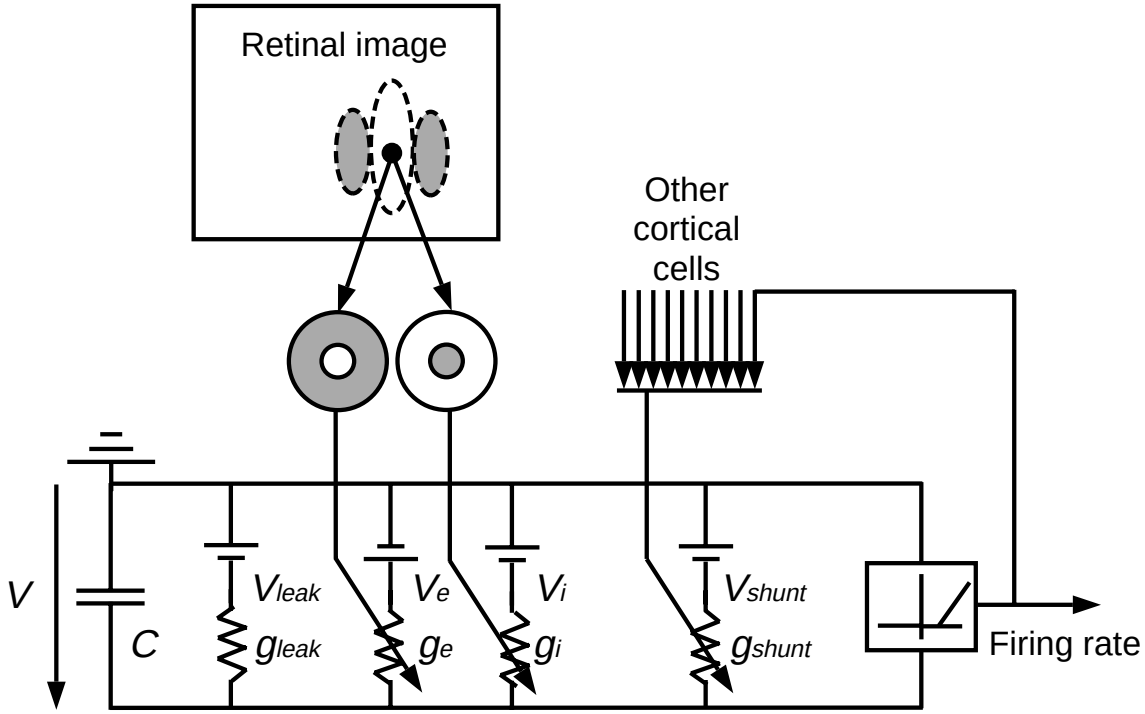
Retinal image
Other
cortical
cells
V
C
V_{leak}
g_{leak}
V_e
g_e
V_i
g_i
V_{shunt}
g_{shunt}
Firing rate

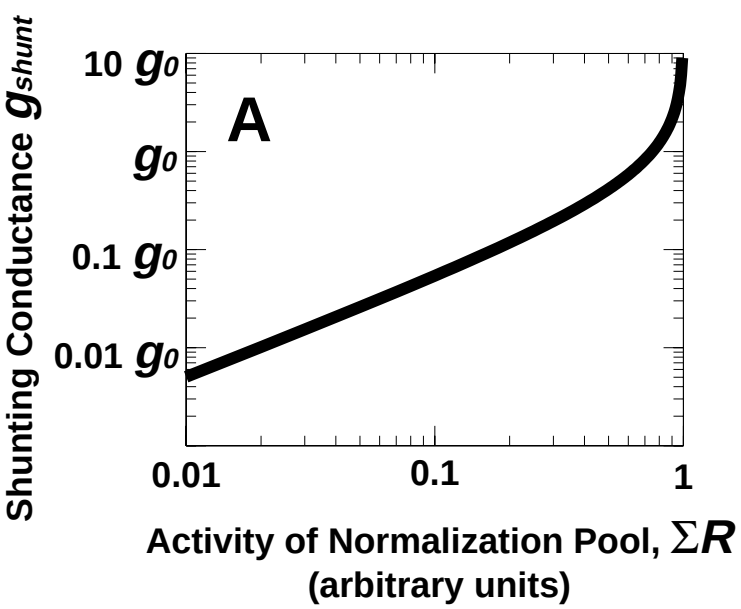
A
Shunting Conductance g_{shunt}
10 g_0
g_0
0.1 g_0
0.01 g_0
0.01
0.1
1
Activity of Normalization Pool, ΣR
(arbitrary units)

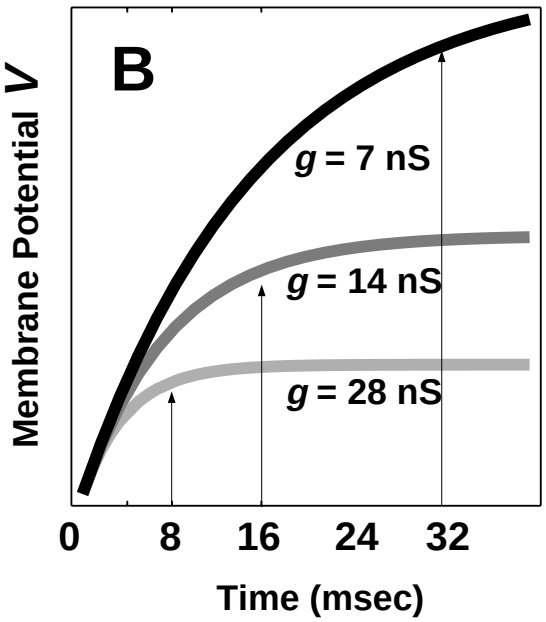
B
Membrane Potential V
g = 7 nS
g = 14 nS
g = 28 nS
0
8
16
24
32
Time (msec)

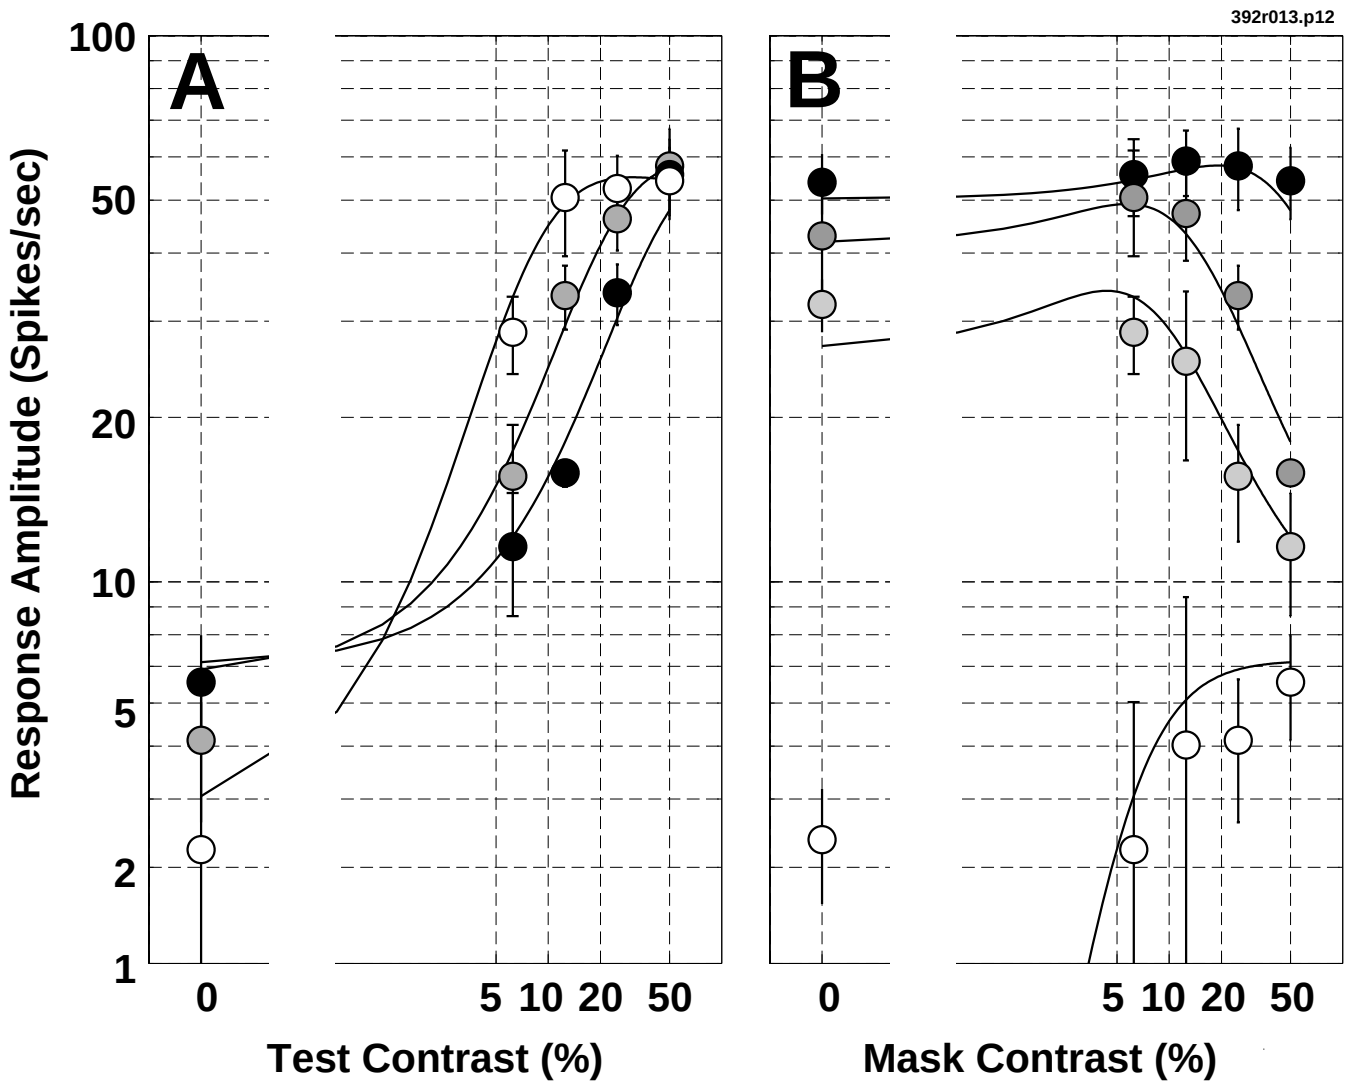
392r013.p12
A
B
100
50
20
10
5
2
1
Response Amplitude (Spikes/sec)
0
5 10 20 50
Test Contrast (%)
0
5 10 20 50
Mask Contrast (%)

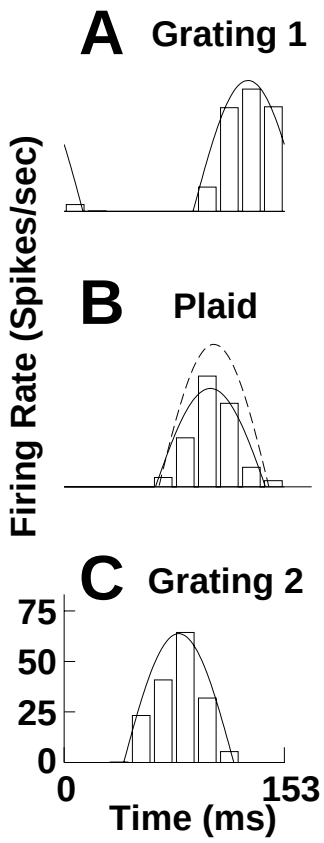
A
Grating 1
B
Plaid
C
Grating 2
Firing Rate (Spikes/sec)
75
50
25
0
0
153
Time (ms)

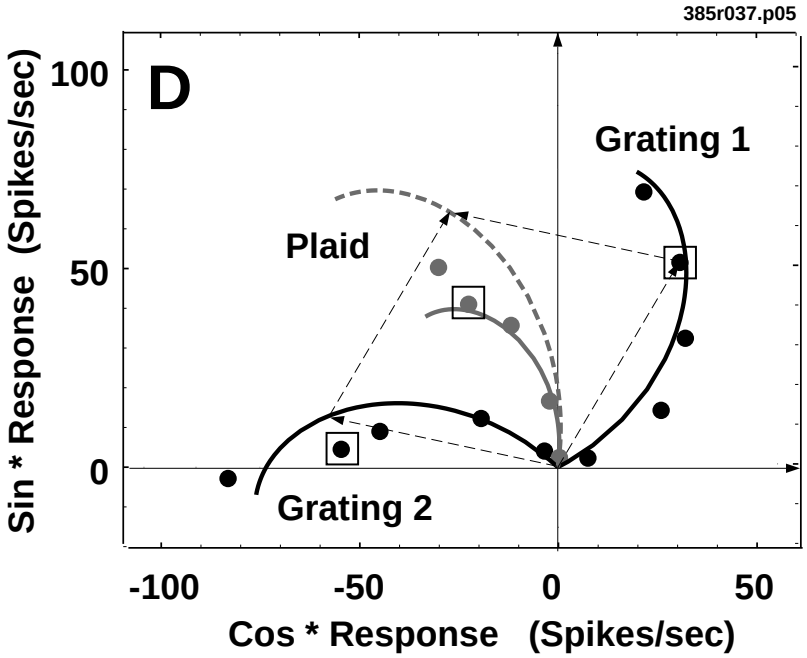
385r037.p05
D
Grating 1
Plaid
Grating 2
100
50
0
-100
-50
0
50
Sin * Response (Spikes/sec)
Cos * Response (Spikes/sec)

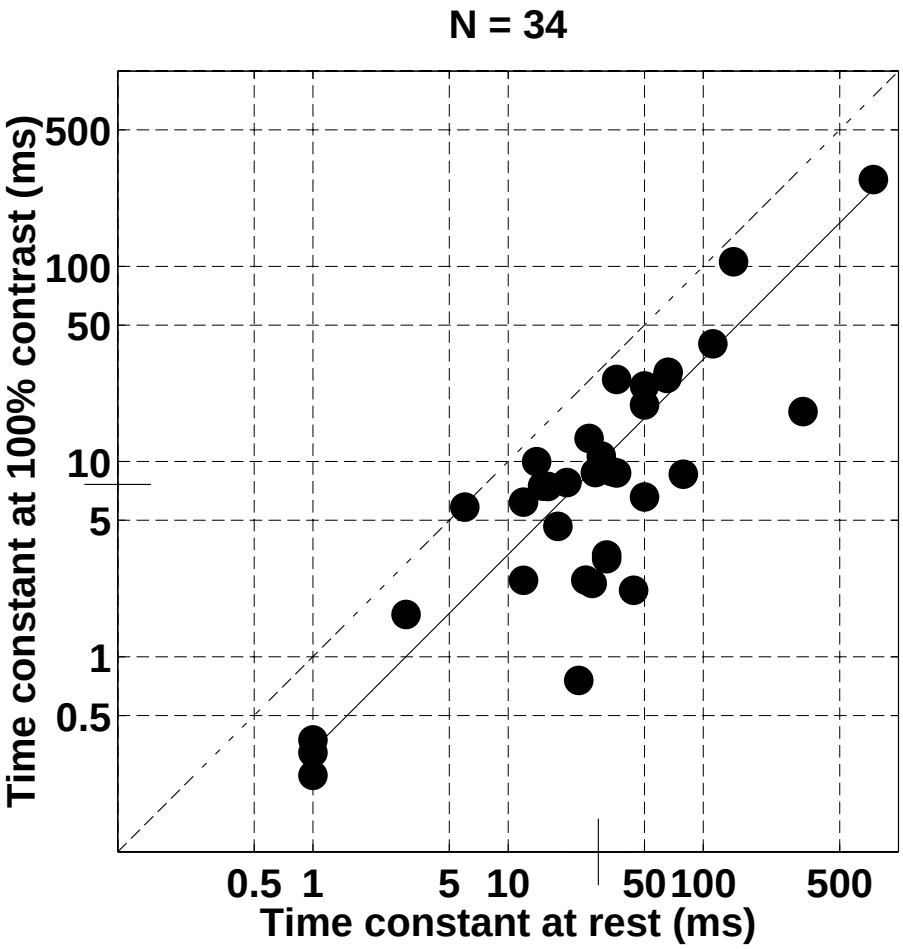
N = 34
Time constant at 100% contrast (ms)
500
100
50
10
5
1
0.5
0.5
1
5
10
50
100
500
Time constant at rest (ms)