A neurophysiologically plausible population code model for human contrast discrimination

Robbe L. T. Goris

Laboratory of Experimental Psychology, University of Leuven, Leuven, Belgium

Felix A. Wichmann

Technische Universität Berlin, FG Modellierung Kognitiver Prozesse, Berlin, Germany, & FRG and Bernstein Center for Computational Neuroscience, Berlin, Germany

G. Bruce Henning

Colour and Vision Group, The Institute of Ophthalmology, London, UK

The pedestal effect is the improvement in the detectability of a sinusoidal grating in the presence of another grating of the same orientation, spatial frequency, and phase—usually called the *pedestal*. Recent evidence has demonstrated that the pedestal effect is differently modified by spectrally flat and notch-filtered noise: The pedestal effect is reduced in flat noise but virtually disappears in the presence of notched noise (G. B. Henning & F. A. Wichmann, 2007). Here we consider a network consisting of units whose contrast response functions resemble those of the cortical cells believed to underlie human pattern vision and demonstrate that, when the outputs of multiple units are combined by simple weighted summation—a heuristic decision rule that resembles optimal information combination and produces a contrast-dependent weighting profile—the network produces contrast-discrimination data consistent with psychophysical observations: The pedestal effect is present without noise, reduced in broadband noise, but almost disappears in notched noise. These findings follow naturally from the normalization model of simple cells in primary visual cortex, followed by response-based pooling, and suggest that in processing even low-contrast sinusoidal gratings, the visual system may combine information across neurons tuned to different spatial frequencies and orientations.

Keywords: spatial vision, contrast perception, pedestal effect, computational model, population coding

Introduction

Behavioral evidence from detection experiments that involve barely visible low-contrast stimuli suggests that information in the early visual system is analyzed in spatially localized, spatial-frequency-tuned, orientation-selective "channels" of limited bandwidth (Blakemore & Campbell, 1969; Campbell & Robson, 1968; DeValois & DeValois, 1988; Graham, 1989; Graham & Nachmias, 1971; see Figure 1a). At the neural level, the multidimensional stimulus selectivity of cortical neurons makes primary visual cortex a likely substrate for these channels (DeValois & DeValois, 1988; Graham, 1989). Understanding the processing of stimuli that are more than just detectable, however, is a prerequisite for any useful model of spatial vision. To gain insight into visual processing of suprathreshold contrasts, sinusoidal contrast discrimination has been studied extensively (Bird, Henning, & Wichmann, 2002; Foley, 1994; Goris, Wagemans, & Wichmann, 2008a; Henning & Wichmann, 2007; Legge & Foley, 1980; Nachmias & Sansbury, 1974; Wichmann, 1999). One of the main findings of contrast-discrimination

studies is the pedestal effect—the improved detectability of a sinusoidal "signal" grating in the presence of a low-contrast grating with the same spatial frequency, orientation, and phase as the signal and often called "the pedestal."

Characteristics of contrast transduction and/or gain-control mechanisms believed to operate within single channels have often been inferred from the pedestal effect on the assumption that stimuli with narrowband spatial-frequency spectra are processed within a single channel tuned to the orientation and spatial frequency of the signal (Foley, 1994; Legge & Foley, 1980; Wichmann, 1999). This approach has not been without success: The contrast gain-control model, for instance, makes use of a narrowly tuned excitatory factor and a broadly tuned divisive inhibitory factor that give rise to a channel with a Mexican-hat-shaped weighting function (see Figure 2a) and explains much of the variance in sinusoidal contrast-discrimination data (Foley, 1994; Wichmann, 1999). Further, the gain-control model correctly predicts some of the ways in which contrast-discrimination performance changes in the presence of an additional sinusoidal masker having an orientation other than that of the signal (Foley, 1994).
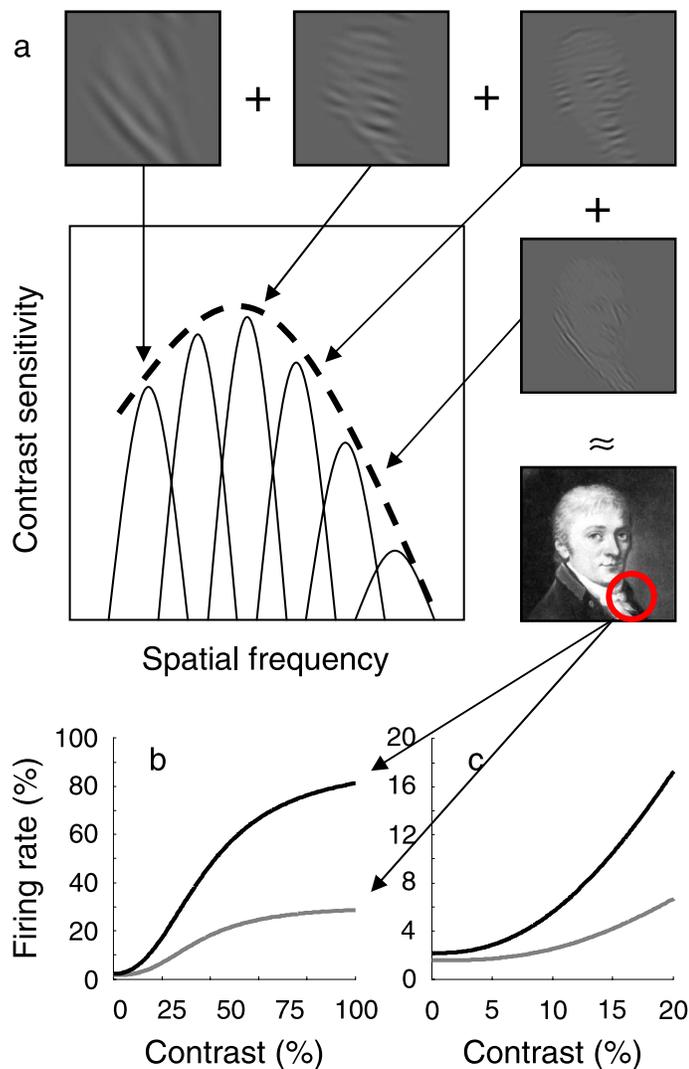
Figure 1. The multi-channel model of spatial vision. (a) Linear narrow-band filter responses to an example image. Particular filters' spatial-frequency tuning is shown in the main panel and their response illustrated in the surrounding panels. For simplicity in the illustration, the cells' orientation tuning has been ignored. The weighted sum of their responses approximately produces the characteristics of the full image. The dashed curve in the main panel of (a) depicts the overall contrast sensitivity function. (b) The simulated contrast response functions of two hypothetical cortical neurons that behave as localized spatial-frequency and orientation-selective filters. Cell firing rate is plotted as a function of image contrast, and the two filters are differentially sensitive to some narrow band stimulus. At all contrasts, the cell that is tuned to the spatial frequency and orientation spectrum of the image patch has a stronger response (upper dark line) than a less well-tuned cell (lower gray line). Note that response saturation of both cells occurs at the same contrast. (c) The same contrast response functions, expanded to show the response over the low-contrast range. At low contrasts, the contrast response function behaves as a positively accelerating non-linearity (see Finn et al., 2007; Miller & Troyer, 2002).

However, the assumption that only the most sensitive channel is monitored has recently been challenged in a series of contrast-discrimination experiments (Henning & Wichmann, 2007): The pedestal effect is somewhat reduced in the presence of broadband noise or when either low-pass or high-pass noise distant in frequency from the spatial frequency of the signal is used. But when the high-pass and low-pass noises are combined to produce a "notched noise" from which a 1.5-octave wide notch centered on the signal frequency has been removed, the pedestal effect all but disappears. As illustrated in Figure 2, single-channel models like the contrast gain-control model fail to account for these results.

Many different simulations of the gain-control model—more specifically, Foley's model 3—one of which is illustrated in Figure 2, all predict that the pedestal effect should be reduced in broadband noise, but persist in notched noise—the opposite of what in fact happens. This discrepancy may indicate that this implementation of divisive inhibition is not the appropriate computation to capture the effects of visual noise. However, in addition, there is a discrepancy between the deep dip attributed to single channels by gain-control models, and the mild dip observed in single cells of the striate cortex (Geisler & Albrecht, 1997). Thus, a single-channel explanation of the pedestal effect seems unlikely.

Bayesian models based on single-neuron characteristics that include the variability of the neurons' responses may reconcile the mild dip in single cells and the strong pedestal effect observed psychophysically, at least when a "hard" response threshold is assumed (Chirimuuta & Tolhurst, 2005). Without a response threshold, Bayesian optimal decoding of single cell responses does not produce a strong pedestal effect (Chirimuuta & Tolhurst, 2005; Geisler & Albrecht, 1997). In general, neuro-physiological studies have not demonstrated the necessity of a response threshold to describe single cell contrast response functions at a phenomenological level (Miller & Troyer, 2002). Further, because the aforementioned Bayesian models do not consider the variability in the spatial-frequency tuning of cortical neurons, they would also fail to capture the contrast discrimination in noise results of Henning and Wichmann (2007). In sum, these results appear to require a revised model for human contrast discrimination.

Henning and Wichmann (2007) conclude that the pedestal effect may stem from the use of contrast information carried by channels tuned to spatial frequencies other than that of the signal frequency, thereby raising the important question of how the response of the whole population of channels is decoded in contrast discrimination and, indeed, in detection. In this paper, we evaluate the consequences for contrast discrimination of simple weighted summation—where the contribution of each channel to detection or discrimination performance is proportional to its response. In particular, we consider possible neural correlates of "channels," where a channel
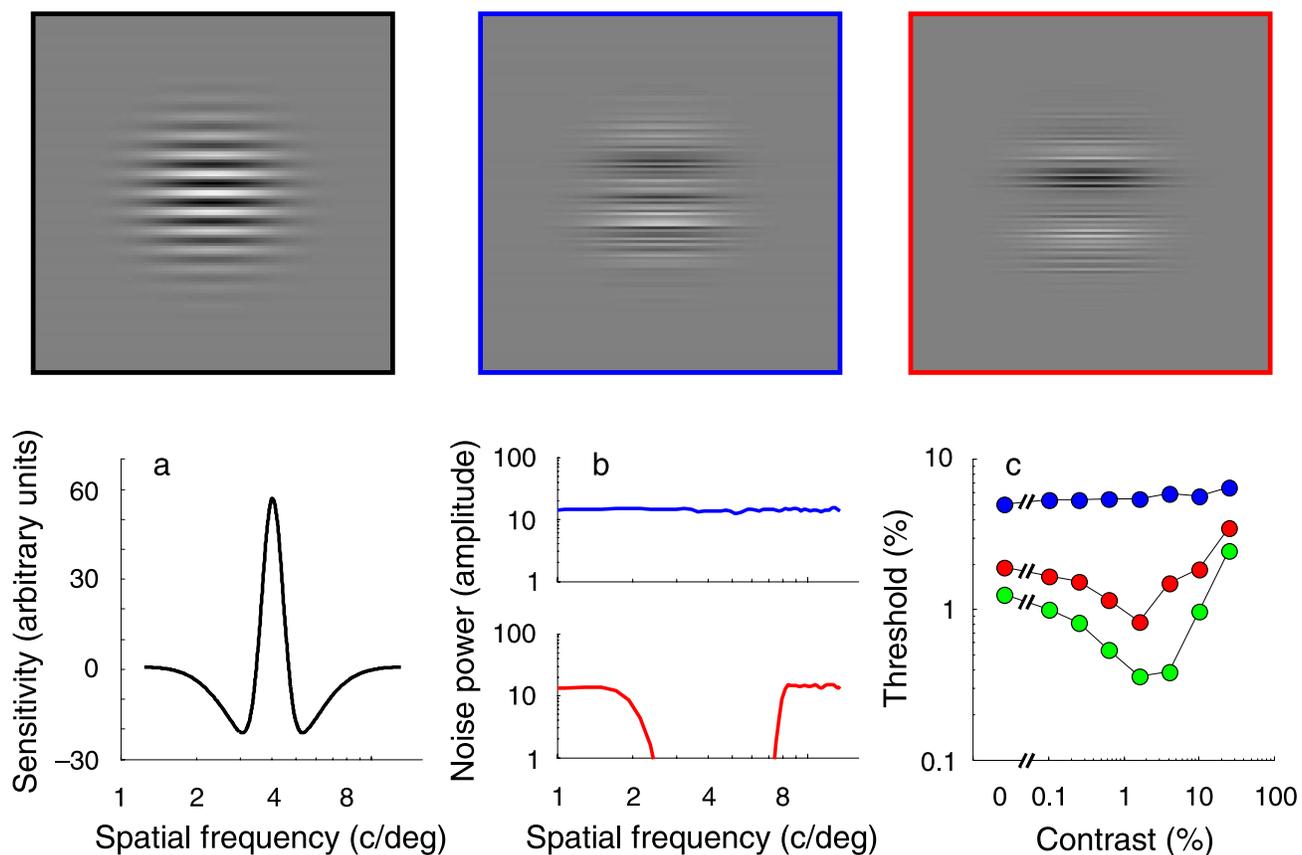
Figure 2. Simulation of the contrast gain-control model (Foley's Model 3) in the contrast-discrimination experiments of Henning and Wichmann (2007). Upper row: illustration of the signal and noise stimuli used in the simulation—from left to right: the 4 cycles/deg sinusoidal signal, broadband noise, and notched noise. Lower row: (a) The difference in contrast sensitivity of the excitatory and inhibitory channel components as a function of spatial frequency. The narrowly tuned excitatory component and broadly tuned inhibitory component give rise to a Mexican-hat-shaped channel. (b) The noise-power densities as a function of spatial frequency—top: broadband noise (blue), bottom: notched noise (red). (c) Performance predicted by the gain-control model in the contrast-discrimination experiments of Henning and Wichmann (2007). Without noise added (green symbols), the model produces a dipper-shaped threshold-versus-contrast function that closely mimics human data. Inconsistent with human performance, the dip is virtually absent in the presence of broadband noise (blue symbols). However, the improvement in performance occurs in the presence of notched noise (red symbols), again inconsistent with human behavior.

is a set of neurons with similar selectivity for orientation and spatial frequency. Weighted summation of signals from neural populations has been suggested to offer a general, biologically plausible mechanism capable of approximating ideal observers' behavior (Deneve, Latham, & Pouget, 1999; Jazayeri & Movshon, 2006; Pouget, Zemel, & Dayan, 2000). Further, under many realistic conditions, employing a simple response-based weighting heuristic even resembles optimal combination rules.

Here, we simulate a network made up of units whose contrast response functions resemble those found in cortex (see Methods). In our simulations, we additionally use the population distributions of the units' characteristics even though information about population distributions is very limited. The units are restricted to physiologically measured characteristics to obtain a neurophysiologically plausible model. Because the neurophysiologically determined data for cortical cell responses to noise are also very

limited, our implementation of noise effects is necessarily based on theory-based predictions and simulations. Nevertheless, this implementation is consistent with available physiological findings and, as we shall see, the simulated networks predict plausible contrast-discrimination data. Further, contrary to the suggestion of Henning and Wichmann (2007), notched noise in our implementation does not prevent off-frequency looking; rather it modifies cortical cell activity in such a way that, upon pooling, the pedestal effect virtually disappears.

Neurophysiological recordings show that as the contrast of a sinusoidal grating driving cortical cells increases from zero, the cells' responses first increase in an accelerating fashion, remain approximately linear over a limited contrast range, and then saturate (Albrecht, Geisler, Frazor, & Crane, 2002; Geisler & Albrecht, 1997) (see Figures 1b and 1c). The saturation level is often attributed to some form of contrast-gain control.

Both non-linearities—response expansion at low contrasts and saturation at high contrasts—are usually thought to be fully expressed at the onset of the response and may play an important role in enhancing and maintaining stimulus selectivity (Albrecht et al., 2002; Geisler & Albrecht, 1997), although the dynamics of the physiological responses are not yet fully understood (Ringach, Hawken, & Shapley, 1997). Computationally, these properties are consistent with a model in which a strictly linear spatiotemporal stage is followed by squaring, then half-wave rectification, and broadband divisive inhibition—the normalization model (Carandini, Heeger, & Movshon, 1997; Heeger, 1992a, 1992b)—although all the properties may be intrinsic to the feed-forward mechanisms of simple cells (Carandini, Heeger, & Senn, 2002; Finn, Priebe, & Ferster, 2007; Freeman, Durand, Kiper, & Carandini, 2002; Miller & Troyer, 2002). The contrast-determined response saturation allows cortical neurons to signal information about the location, spatial frequency (or size), and orientation of local image features with considerable precision, at the expense of a detailed representation of contrast information (Albrecht et al., 2002; Geisler & Albrecht, 1997).

A final and particularly important property of cortical neurons is the characteristic that the variance of their responses to a sinusoidal grating is proportional to their mean response (Albrecht et al., 2002; Geisler & Albrecht, 1997; Vogels, Spileers, & Orban, 1989). We call this crucial characteristic "the proportionality rule" and implement it in our network units as a multiplicative noise source. The simple proportionality relation implies that the ratio of a cell's mean response to the standard deviation of its responses—its signal-to-noise ratio—will increase as the response increases and this leads to the attractive idea that in combining the responses of many cells, the visual system would do well to first weight them by the strength of their response.

In the simulations to follow, we first discuss effects of unit selectivity and stimulus contrast on weight assignment in the model, making use of a pool of 100 units. We then consider, for the sake of clarity and simplicity, a situation where the pool of combined elements consists of only 12 uncorrelated units. One of these units is optimally tuned to the signal, while the tuning of the other units varies randomly from completely insensitive to highly sensitive. This somewhat artificial but illustrative situation enables an explicit comparison between the performance of an "optimal" single unit and the performance of a network and thus allows an assessment of pooling effects. We show how this simple version of the model produces a pedestal effect based on informational pooling across spatial frequencies and is also able to simulate successfully the findings of Henning and Wichmann (2007). We further consider grating *detection* (no pedestal) in noise and find that the model correctly predicts how detection performance changes in spectrally flat and filtered noise. Finally, having estimated the likely number of contributing units, we demonstrate that a more realistic pool consisting of 250 units with correlated responses and tuned to a broad range of spatial frequencies robustly produces similar results.

Methods

Model equations

The contrast response functions of the units in the network simulations (Equation 1 below) are those of the *Invariant Response Descriptive Model* described in Albrecht et al. (2002), expanded to include an explicit selectivity parameter, *Sel,* which varies between 0 (i.e., the unit is not sensitive to the signal) and 1 (i.e., the peak sensitivity of the unit's spatial-weighting function corresponds to the spatial frequency of the signal). The selectivity parameter is needed for the units' response functions to have the behavior demonstrated by cortical neurons (Albrecht et al., 2002; Geisler & Albrecht, 1997). The response functions, based on the Naka–Rushton equation, provide a good fit to the contrast response functions of striate cortex neurons to preferred (Sel = 1) and non-preferred (Sel < 1) stimuli (Albrecht et al., 2002; Geisler & Albrecht, 1997). Equation 1 shows the mean response of a unit, $\overline{R_u}$, as a function of stimulus contrast c, expressed as a fraction of the unit's maximal firing rate:

$$\overline{R_u}(c) = r_0 + \text{Sel}\left(r_{\max}\frac{c^n}{c_{50}^n + c^n}\right), \qquad (1)$$

where r_0 is a spontaneous discharge rate. In the simulations, r_0 is drawn from an exponential distribution with a mean value of 1.5% of the maximal firing rate (Olshausen & Field, 2005) [$r_0 \sim \text{Exp}(1.5)$]; $r_{\max}$ is the maximum firing rate, drawn from a normal distribution with mean 81.8 and standard deviation 12.2 [$r_{\max} \sim N(81.8, 12.2)$]; n is the response exponent [$n \sim N(2.4, 0.18)$]; c_{50} is the semi-saturation contrast [$c_{50} \sim N(0.387, 0.0351)$]. The expressions in square brackets following the definition of the terms in equations, give, where appropriate, the form and parameters of the distribution from which values for the terms were randomly selected. The parameter distributions are based on neurophysiologically determined estimates (Albrecht et al., 2002). However, the exact parameter settings are not critical to any of the claims made in the paper. Nevertheless, we shall see that these distributions produce a remarkably good approximation of psychophysical data.

In the model, as in the measured behavior of cortical cells, the variance of a unit's response is proportional to its mean value and is given by Equation 2:

$$\text{Var}(R_u(c)) = 1.5(\overline{R_u}(c)), \qquad (2)$$

where $\text{Var}(R_u(c))$ is the variance of a unit's response as a function of stimulus contrast. The scaling value of 1.5 is based on estimates provided in several papers on cortical cell response reliability (Albrecht et al., 2002; Geisler & Albrecht, 1997; Vogels et al., 1989). The particular value of the proportionality constant is not critical, but the fact that variance is proportional to mean activity is crucial.

For the network simulations in this paper, the weight of each unit, $\omega_u(c)$, is fully determined by its mean response, normalized, for convenience, by the sum of all the contributing weights and given by Equation 3:

$$\omega_u(c) = \frac{\overline{R_u}(c)}{\sum_{i=1}^{N} \overline{R_i}(c)}, \tag{3}$$

where $\omega_u(c)$ is the weight of unit u at stimulus contrast c in a network of N units. In order to make the simulations tractable, the trial-to-trial variation in the weights was ignored; we used the mean responses in calculating the weights. This is a simplification. In a real nervous system, the means would not, of course, be available from a single unit and weights would necessarily be based on responses alone. In Appendix A, we demonstrate that this simplification is immaterial with respect to the conclusions we draw.

The mean pooled response at any given contrast, $\overline{R_{\text{pooled}}}(c)$, is:

$$\overline{R_{\text{pooled}}}(c) = \sum_{u=1}^{N} \omega_u(c)\overline{R_u}(c). \tag{4}$$

The more responsive and hence, because of the proportionality rule, the more reliable units thus attract more weight. We do not wish to suggest that the decoding is necessarily as simple as our heuristic rule. But its simplicity is appealing and for most situations, it resembles an optimal combination rule without the computational burden of determining the covariance matrix for the units in the network or, indeed, knowing anything about the precision of any unit—because the variance is proportional to the responsiveness, only the strength of the response to any given stimulus matters.

Finally, to compute the variance of the pooled responses, we made use of the pooling formula (Shadlen & Newsome, 1998):

$$\text{Var}(R_{\text{pooled}}(c)) = \sum_{u=1}^{N} \omega_u^2(c)\text{Var}(R_u(c)) + \sum_{u=1}^{N}\sum_{v\neq u}^{N} r_{uv}$$
$$\cdot \sqrt{\omega_u^2(c)\text{Var}(R_u(c))\omega_v^2(c)\text{Var}(R_v(c))}, \tag{5}$$

where r_{uv} is the correlation coefficient between the uth and the vth unit. Correlation among units is, of course, notoriously difficult to determine. Nevertheless, it is an important—indeed a crucial—factor in some models of MT pooling (Shadlen, Britten, Newsome, & Movshon, 1996).

Contrast-discrimination performance of the network, expressed as d', is fully determined by the mean and variances of the pooled responses to the pedestal and to the pedestal-plus-signal (Green & Swets, 1966) and given by Equation 6:

$$d'_{\text{pooled}}(c_{\text{ped}}, c_{\text{ped+sig}}) = \frac{\overline{R_{\text{pooled}}}(c_{\text{ped+sig}}) - \overline{R_{\text{pooled}}}(c_{\text{ped}})}{\sqrt{\text{Var}(R_{\text{pooled}}(c_{\text{ped+sig}})) + \text{Var}(R_{\text{pooled}}(c_{\text{ped}}))}}. \tag{6}$$

Cortical neurons and visual noise

Early neurophysiological work demonstrated some heterogeneity among cortical neurons in cat striate cortex in response to noise. Simple cells have been reported to be unresponsive to a broadband noise stimulus while their response to an otherwise optimal stimulus is reduced when this stimulus is embedded in noise (Hammond & MacKay, 1977; Maffei, Morrone, Pirchio, & Sandini, 1979). Other neurons (mainly complex cells) have been reported to respond to (some kinds of) broadband noise, although their response to an otherwise optimal stimulus is also reduced when this stimulus is embedded in noise (Hammond & MacKay, 1977; Maffei et al., 1979). For both kinds of neurons, the main effect of spectrally flat noise is thus that single units' contrast–response functions shift toward higher contrasts and lower response rates.

More recent work by Carandini et al. (1997) demonstrated that the effects of binary noise on the contrast–response functions of simple cells in macaque primary visual cortex are well captured by the normalization model mentioned briefly in the Introduction. In their paper, Carandini et al. make the simplifying assumption that the noise would be unable to drive the linear receptive field of the cells, so that the sole effect of the noise would be to provide divisive normalization. To fit their data, Carandini et al. introduced an additional parameter α, controlling the effectiveness of noise in driving the normalization pool and reported that the values of α resulting from the fits to 22 simple cells were equally spread (on logarithmic coordinates) between 0.1 and 10, indicating that the noise provided very strong inhibition for some cells but only weak inhibition for others.

In this paper, we largely follow the implementation of Carandini et al. (1997) to capture the effects of noise on units in our model. However, we do not assume that broadband noise is unable to drive the linear receptive field for several reasons. First, Carandini et al. reported a
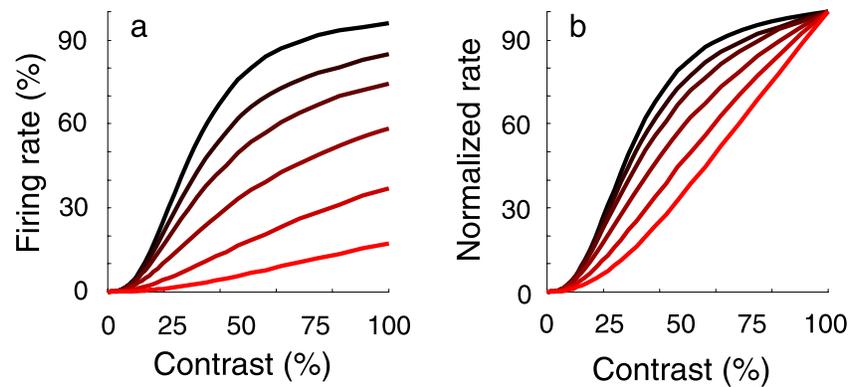
Figure 3. The effects of notched noise. (a) The contrast response function for one particular unit—tuned to the center frequency of the notch, with $r_0 = 0$, $n = 2.6$, and $c_{50} = 0.3$—as a function of signal contrast with different levels of notched noise added to the signal. As the noise power increases (the lower, more reddish curves), the inhibitory response increases and becomes more variable, which suppresses and effectively linearizes the contrast response function. (b) The same response functions as in a, normalized by their respective response at maximal signal contrast.

3-fold elevation of the maintained discharge in noise. This is unlikely to be explained by suppression and implies activation; second, noise activation is crucial to explain some psychophysical observations—contrast detection performance improves slightly when weak noise is added to a low-contrast signal (Goris, Zaenen, & Wagemans, 2008b); third, in the logic of the normalization model, the crucial difference between notched noise and flat noise for cells tuned to frequencies in the notch is that notched noise provides only non-specific suppression, while flat noise provides both excitatory and inhibitory activation.

The equations used to describe unit responses in broadband and notched noise are derived from Carandini's implementation combined with simulations and explained in detail in Appendix B (Equations B3 and B4). In summary, in the model developed here, addition of broadband noise increases the maintained discharge, modifies the response exponent, and leads to an effective shift of the contrast response function toward higher contrasts and lower response rates. The strength of this shift varies across units, as has been reported for simple cells. On the other hand, while the addition of notched noise produces similar effects, its effects are modified by the presence of the notch and by units' selectivity to frequencies both within and outside the notch. For units tuned to the frequencies outside the notch, notched noise is effectively equivalent to broadband noise, but for units tuned to frequencies in the notch, notched noise produces strong non-specific suppression. This is illustrated in Figure 3. Figure 3a shows how the mean response of one particular unit—tuned to the signal frequency in the center of the notch—varies with signal contrast when different levels of notched noise are added to the signal. As the noise power increases (the more reddish curves), the inhibitory response increases, because the normalization pool is broadly tuned, and becomes more variable, which suppresses and effectively linearizes the contrast response

function. This is shown in Figure 3b where the response functions shown in panel are replotted, normalized by the response at maximal signal contrast. Note that the response acceleration decreases with noise power. However, in the presence of notched noise, it is not this unit, as we shall see, that attracts the greatest weighting in the pool.

The characteristics of the noise effects are based on a relatively small number of physiological observations combined with theory-based simulations, thus it is at best an approximation. Nevertheless, the simulated networks robustly predict plausible contrast-discrimination data, as we shall see.

Results

In the simulations described below, we first discuss the effects of unit selectivity and stimulus contrast on weight assignment in the model, making use of a pool of 100 units. We then consider a situation where the pool consists of 12 uncorrelated units. We show how this simple version of the model produces a pedestal effect based on weighted pooling across spatial frequencies and successfully simulates the findings of Henning and Wichmann (2007). Finally, we demonstrate that a more realistic version of the model produces similar results.

Weight assignment: Effects of unit selectivity and stimulus contrast

In our network, each element's contribution depends simply on the magnitude of its response. Since the reliability of the unit is proportional to the mean response,
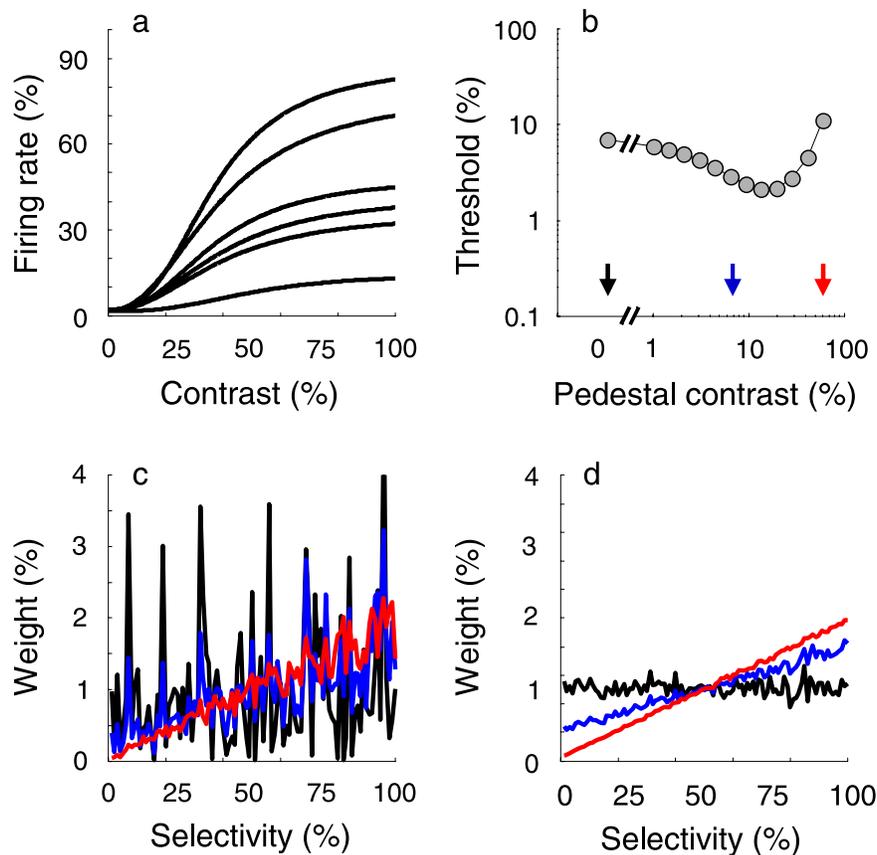
Figure 4. Illustration of the effects of stimulus contrast on the distribution of weights. In this simulation, the network consisted of 100 units, with stimulus selectivity, indicated by the relative saturation levels, varying from 1% to 100% in steps of 1%. (a) The contrast response functions to some narrow-band stimulus of six of the network units with similar orientation tuning but different spatial frequency tuning. (b) The dipper-shaped 75% correct threshold-versus-contrast function of the network. (c) Weight as a function of stimulus selectivity at the three different pedestal contrasts indicated by the arrows in panel b. (d) Same as panel c but averaged over 100 simulations.

such a weighting provides a crude approximation to optimal weighting. Figure 4a illustrates response rate as a function of contrast for a family of cortical neurons responding, in the absence of added external noise, to a narrow-band stimulus at their preferred orientation (Albrecht et al., 2002; Geisler & Albrecht, 1997).

In Figure 4a, the units with high saturation levels are those best tuned to the spatial frequency of the stimulus while progressively lower saturation levels represent units that respond less and less to the stimulus because their optimal stimuli have spatial frequencies that are further and further away from the spatial frequency of the stimulus. The close relation between saturation level and selectivity for spatial frequency allows us to use the relative saturation level of each unit as an indication of its selectivity for the sinusoidal signal and pedestal (see Equation 1). The saturation level of the 100 units in this simulation ranged from 100% to 1% of the maximum saturation level—a range that would represent a range of peak sensitivities of roughly 5 octaves, uniformly sampled on a logarithmic axis, around the frequency of the signal (we assume geometrically symmetrical tuning functions

with a bandwidth of approximately 1.5 octaves). These selectivities were mainly chosen for convenience. (Note that in behavioral studies, the contrast of narrowband stimuli rarely exceeds 50% and even this may be an order of magnitude greater than the contrasts that typically occur naturally (Frazor & Geisler, 2006). Thus, the behaviorally relevant region of Figure 4a lies below about 30% contrast.)

The behavior shown in Figure 4b captures the central features of contrast discrimination that appear in experiments without added external noise. Figure 4b illustrates the network's dipper-shaped threshold-versus-contrast function at 75% correct in a 2AFC task; the "dipper" shape constitutes the pedestal effect. Almost no unit in the simulation shows the pedestal effect by itself–rather the effect emerges from the pooled responses as will be discussed below.

Figures 4c and 4d both show the fraction of the total weight assigned to different units as a function of their selectivity. Results for three different pedestal contrasts– indicated by the colored arrows in Figure 4b–are shown: zero contrast in black, a just-visible pedestal contrast in

blue, and a highly visible pedestal contrast in red. Figure 4c shows the distribution of weights as a function of unit selectivity for one simulation; Figure 4d shows the weights averaged over 100 simulations. For each of these simulations, the set of cell selectivities was held constant, but all other parameters of our network were randomly selected from the distributions described in the Methods section.

When a blank stimulus—a uniform field of the same mean luminance as the sinusoidal grating—is presented (black curves), the units by definition respond at their (usually non-zero) spontaneous discharge rate. Consequently, the jagged black line reflects the probability distribution of spontaneous activity across cells. The distribution was chosen to be exponential (Olshausen & Field, 2005) with an average of 1.5% of maximal firing rate (see Methods). This cell characteristic—base rate activity—is independent of stimulus selectivity and the distribution of weights is uniform with respect to selectivity (as can be seen in Figure 4d, where the black line is the average weight in 100 simulations and is approximately horizontal). As stimulus contrast increases from zero, the relative magnitude of a unit's response begins to reflect not only its spontaneous rate but also a stimulus-driven part—the latter determined by the unit's selectivity (see Equation 1). Hence, at non-zero stimulus contrasts, higher weights are attracted by the better tuned units (i.e., more responsive units) in a fashion that is roughly linearly related to selectivity (the blue and red lines in Figures 4c and 4d). The higher weighting occurs, of course, because the weighting of a unit increases as its response increases. However, both the slope and the variability of this linear relation depend on stimulus contrast.

At low, barely visible pedestal contrasts (blue curves), spontaneous activity contributes a high proportion of the total cell response. As a consequence of the exponentially distributed spontaneous discharge levels, a few units with high base-rate activity will attract much of the weight (the blue curve in Figure 4c). Note that such units would be called "irrelevant" in the context of uncertainty models (Pelli, 1985). As pedestal contrast increases, more and more cells are driven by the stimulus and spontaneous activity gradually looses its influence on network activity. Consequently, the variability across the distribution of weights decreases with increasing contrast (the red curve in Figure 4c) and the slope of the distribution of weights increases (the red curve in Figure 4d).

Thus, with contrast detection and discrimination based on the combined response of many elements instead of a single element, the class of units contributing most to the decision statistic and distinguished by their stimulus selectivity or spatial-frequency tuning varies a lot. With weighting determined by responsiveness and especially at low stimulus contrasts, units tuned to spatial frequencies and orientations remote from the signal are often weighted heavily.

We now investigate in more detail the effects of weighted pooling on contrast detection and discrimination.

In the following simulations, we consider the simplified situation where the pool consists of 12 uncorrelated units. One of these units is optimally tuned to the signal (i.e., selectivity = 1), and the selectivity of the 11 other units, as indexed by their saturation level, is randomly drawn from a Gaussian distribution centered at 0.50, with a standard deviation of 0.17 and clipped at 0 and 1. For the 4-cycle/deg signal used in our simulations, this selectivity distribution is not inconsistent with spatial frequency tuning properties of cortical neurons (Geisler & Albrecht, 1997); i.e., many neurons are somewhat sensitive to the signal, while few are optimally tuned to the signal, or completely insensitive to the signal.

Contrast discrimination: Effects of response pooling

Figure 5a illustrates response rate as a function of contrast for all 12 units in one particular pool driven, in the absence of added external noise, by a narrow-band stimulus at their preferred orientation. The contrast response function of the unit that is optimally tuned to the signal is shown in purple. In comparison to monitoring only the most sensitive unit, pooling the responses from neurons of different sensitivity weighted by their responsiveness, i.e., weighted pooling across units, improves detectability at all stimulus contrasts except in some cases where the contrast is so low that few units are stimulus driven. Moreover, because of the proportionality rule, the improvement in sensitivity from weighted pooling relative to the most selective single unit increases as contrast increases.

To see this consider Figure 5b, which illustrates how the ratio of the mean pooled response to the standard deviation of the pooled response depends on the contrast of a narrowband stimulus. This ratio is closely related to the detectability of the stimulus and is labeled d' (Green & Swets, 1966). The ratio of the mean to the standard deviation for the pooled group is indicated by the green line, the same ratio for the most selective unit, by the purple line. The higher this ratio, the better the system is able to discriminate a low-contrast signal grating from a uniform field. At the very lowest stimulus contrasts, where the non-optimally tuned units mainly contribute noise, the purple line lies above the green line, indicating that performance of the most selective unit is slightly better than that of the pooled response. From a certain stimulus contrast on, however, the green line lies above the purple line, indicating that the pooled response outperforms the most selective unit. Moreover, as contrast increases, the difference between the two functions grows as a consequence of the proportionality rule and the changing weighting profile (Figure 4d). The pooled detectability function is thus more sharply accelerated and it has been suggested that this particular non-linearity underlies the pedestal effect (Nachmias, 1981; Nachmias & Sansbury, 1974; Smithson, Henning, MacLeod, & Stockman, 2009).
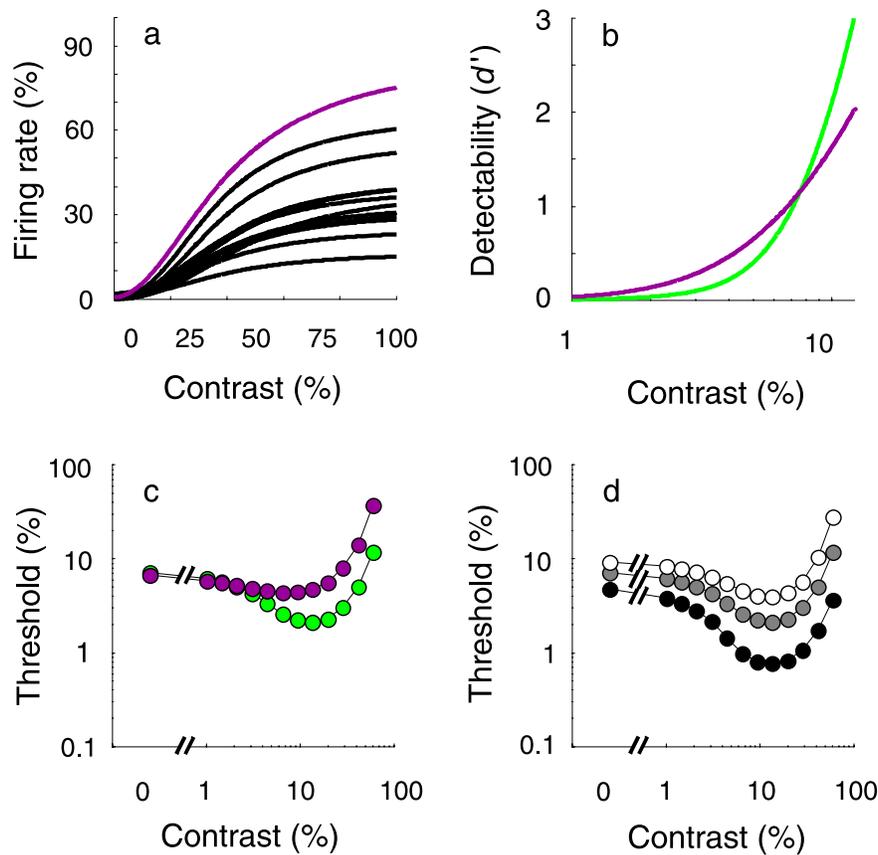
Figure 5. Illustration of the effects of pooling on contrast discrimination. (a) The response functions of the 12 units in one particular pool to signals of increasing contrast. The contrast response function of the unit that is optimally tuned to the signal is shown in purple. (b) Detectability, d', expressed as the ratio of the mean response to the standard deviation, as a function of contrast for the most sensitive unit (purple line) and the pooled decision statistic (green line) on semi-logarithmic coordinates. (c) Contrast-discrimination thresholds at 75% correct as a function of pedestal contrast on double logarithmic coordinates for the most sensitive unit (purple) and the pooled decision statistic (green). (d) Contrast-discrimination thresholds as a function of pedestal contrast on double logarithmic coordinates for the pooled decision statistic at three performance levels in 2AFC (black: 60%, gray: 75%, and white: 90%).

Figure 5c shows the 75% correct "thresholds" for detecting a narrowband stimulus as a function of the contrast of the pedestal, again in the absence of noise; the threshold for the pooled group is indicated by the green symbols, for the most selective unit, by the purple. In both cases, the thresholds are determined by calculating the difference between the mean response to the pedestal alone and the mean response to the signal-plus pedestal and dividing this difference by the square root of the sum of their variances (see Equation 6). It is clear in Figure 5c that the pooling of information in the network leads not only to better performance from a certain pedestal contrast on but also to a bigger pedestal effect. The pedestal effect of the pooled network is based on broadband-weighted pooling, i.e., on the use of information carried by units tuned to frequencies other than the signal frequency—i.e., a form of off-frequency looking. The finding that pooling responses weighted by their responsiveness increases the size of the pedestal effect, in addition, removes the puzzling discrepancy between the mild dip observed in

single cells in striate cortex (Geisler & Albrecht, 1997) and the deep dip observed psychophysically (Bird et al., 2002; Nachmias & Sansbury, 1974). However, note that pooling responses in a Bayesian optimal fashion does not produce a strong pedestal effect (see Chirimuuta & Tolhurst, 2005 and Figure 10b in Geisler & Albrecht, 1997).

In human vision, the strength of the pedestal effect is known to depend on the performance level taken to define the "threshold" (Bird et al., 2002; Goris et al., 2008a; Henning & Wichmann, 2007; Nachmias & Sansbury, 1974; Wichmann, 1999). At low performance levels, the maximal pedestal-induced threshold reduction is considerably larger than at high performance levels. This particular property of contrast discrimination has largely been ignored because only one performance contour is usually determined (Foley, 1994; Legge & Foley, 1980) but has already proven useful in model selection (Goris et al., 2008a; Wichmann, 1999). In Figure 5d, discrimination thresholds for the pooled network of weighted unit

Goris, Wichmann, & Henning

activation are plotted at 60% (black), 75% (gray), and 90% (white) correct response levels. The network clearly gives rise to the observed performance-level-dependent pedestal effect: Consistent with human data, the dip of the threshold-versus-pedestal contrast function is deeper at lower performance levels.

Another salient feature of contrast-discrimination studies is that, from a certain pedestal contrast level on, discrimination thresholds as a function of contrast plotted on double logarithmic coordinates rise in a nearly linear fashion (e.g., Bird et al., 2002). This, of course, is Weber's law. It can be seen in Figure 5c that the rise in discrimination thresholds of the pooled population response is also approximately linear on double logarithmic coordinates. This is a consequence of the roughly constant ratio of the variance and mean of the pooled response in the high-contrast region. Thus, because the variance of the pooled response is effectively proportional to the mean pooled response at suprathreshold contrasts, the ratio of the just-noticeable contrast increment and pedestal contrast is approximately constant, consistent with Weber's Law and the linear slope on double logarithmic coordinates.

The slope of this part of the dipper curve on the coordinates of Figure 5c is steeper than psychophysical estimates that are typically around or slightly below one (Bird et al., 2002; Wichmann, 1999). We speculate that two neurophysiological observations that are not included in our model may account for this deviation. First, if one or more of the network units saturated more slowly at higher contrasts, responding more linearly over the entire contrast range, the slope would be lower. Indeed, there is a great deal of heterogeneity among cortical cells and about 5% demonstrate a nearly linear relation between contrast and response magnitude (Albrecht et al., 2002). The effect of reaching saturation more slowly, i.e., extended linearity may be twofold: because their responses continue to increase after other units saturate, the more linear units will attract an increasing proportion of the weighting, provided, of course, that such units ultimately exhibit high saturation levels. Thus, although they may constitute a small proportion of the population, their contribution to the pool will be disproportionately large at high contrasts. Consequently, because of the large contribution of the linear units at high contrasts, the pool will be effectively linearized. The importance of linearization at higher levels of the visual system is a major issue (Eliasmith & Anderson, 2003).

Second, extracellular recordings in cat V1 suggest that firing-rate variance saturates at high firing rates and thus deviates from the proportionality rule used here (Carandini, 2004). Effectively, the signal-to-noise ratio of cortical cell responses may thus be better at high contrasts, which would lower the slope of the rising part of the dipper function.

In summary, pooling the responses of a limited number of units that resemble the contrast response functions and statistical properties found in primary visual cortex by simple weighted summation produces contrast discrimination predictions that resemble several features of human vision. Further, the pedestal effect in this model is based on the use of information carried by units tuned to frequencies other than the signal frequency. This is, of course, a form of off-frequency looking.

We now specifically investigate whether this simple implementation of off-frequency looking is able to capture the contrast discrimination in noise results reported by Henning and Wichmann (2007).

Contrast discrimination in noise: Effects of response pooling

Thus far, we have considered contrast processing in the absence of added external noise. Addition of spectrally flat noise to a narrow-band stimulus increases cell activity at low response rates but inhibits cell activity at high response rates (Carandini et al., 1997; see Appendix B). Presumably, the activation is caused by rectification and the inhibition by broadband normalization. These effects are thus in line with the normalization model and—as explained in detail in Appendix B—can be captured by extending the descriptive equation of the contrast response function (Equation B3). Finally, because noise elicits variable excitatory and inhibitory activation, the acceleration of the contrast response function is reduced or softened by noise (Miller & Troyer, 2002).

To the best of our knowledge, no physiological data on effects of notched noise are available. We thus base our predictions on the logic of the normalization model. For a unit tuned to the frequency in the center of the 1.5-octave wide notch, notched noise can only provide non-specific suppression. This is because the noise will have a negligible effect on spatial-frequency tuned excitatory frequencies of the unit. Given that broadband noise increases cell activity at low response rates, the inhibitory effect of the notched noise is likely to be stronger than the net inhibition provided by broadband noise. In our simulations, notched noise is assumed to cause twice as much inhibition as white noise for a unit tuned to the signal frequency and centred in the notch (see Equation B4 in Appendix B). For a unit tuned to a frequency outside the notch, however, the effects of the noise are assumed to be approximately similar to effects of broadband noise. As is the case for broadband noise, the acceleration of the contrast response function is softened by notched noise. Importantly, simulations with the normalization model showed that stronger noise suppression produces a more linear contrast response function (Figure 3). Thus, the linearizing effects of notched noise are most pronounced for units tuned to the signal frequency (i.e., Sel = 1) and decrease with selectivity.

In summary, the effects of broadband noise and notched noise on the contrast response function were derived from model simulations and described by Equations B3 and B4.

Our model has only one free parameter, i.e., the effective noise power. At a noise power which produces on average a four-fold elevation of the maintained discharge, we find that the model robustly produces contrast-discrimination data taken in noise, that are consistent with the observations of Henning and Wichmann (2007): the pedestal effect is slightly diminished in broadband noise but disappears almost completely in notched noise. (This noise power is not unreasonably large: Carandini et al. (1997)—making use of binary noise at several noise powers—report an average three-fold elevation in their experiments; Henning and Wichmann (2007) used the maximum noise-power density (1D noise) their display system allowed them to generate.)

Figure 6 illustrates the 12 noiseless contrast response functions for one particular pool with no external noise (Figure 6a), as well as the changes these functions undergo in spectrally flat (Figure 6b) and notched noise (Figure 6c). The contrast-discrimination functions for each of these conditions are shown in Figure 6d. The contrast-discrimination threshold functions shown in Figure 6d closely mimic the results of Henning and Wichmann (2007): the pedestal effect in the absence of noise (green) is slightly diminished and shifted on these

coordinates to higher contrasts in broadband noise (blue), but disappears almost completely in notched noise (red). Further, at higher pedestal contrast levels, discrimination thresholds of all noise conditions almost coincide, consistent with human data (Henning & Wichmann, 2007). We thus conclude that our implementation of noise effects and broadband-weighted response pooling is able to produce plausible results for contrast discrimination in noise.

Contrast detection: Effects of response pooling

Pooling the responses of many units with wide heterogeneity in spatial-frequency tuning produces a detectability function that is largely shifted to lower contrasts and more sharply accelerated relative to the detectability function for a single unit (see Figure 5b). The model presented in this paper thus shows that response pooling is a sensible strategy for contrast detection as well. Consequently, experimental manipulations that modify the contrast response functions of cortical cells will affect detection performance. Although not discussed by Henning and Wichmann (2007), the detection data they also gathered in
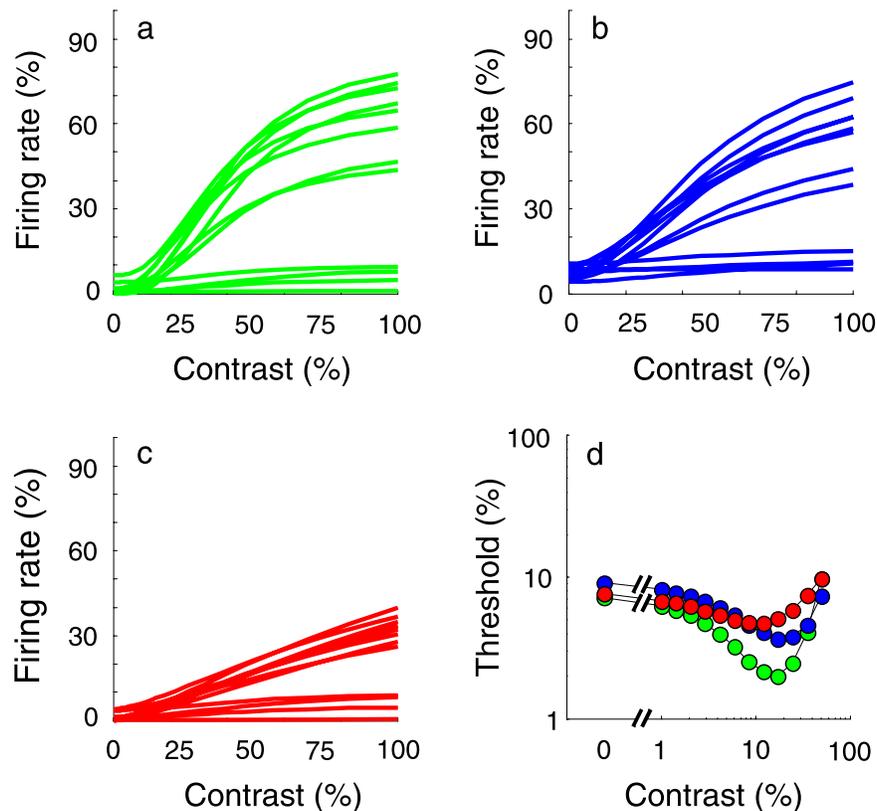

Figure 6. Contrast discrimination in noise. (a) The contrast response functions of the 12 units in one particular pool to the signal with no added external noise. (b) The contrast response functions for the signal embedded in broadband external noise. (c) The contrast response functions for the signal in the presence of notched noise. (d) Predicted contrast-discrimination thresholds for this pool as a function of pedestal contrast with no external noise (green symbols), in broadband noise (blue symbols), and in notched noise (red symbols).

their series of contrast-discrimination experiments provide a test of this hypothesis.

In their experiments, detection of a 4.0-cycle/deg sinusoidal grating was measured without noise and with notched noise added—the notch being 1.5 octaves wide and geometrically centered on the signal frequency. Figure 7a shows how signal detectability increases as a function of contrast without noise (green), in broadband noise (blue) and in notched noise (red) according to the

model. It will be noted that these functions are closely related to the strength of the pedestal effect, i.e., more rapidly accelerating detectability functions correspond to a stronger pedestal effect. Figure 7b plots the predicted psychometric functions, i.e., percentage correct in a 2AFC task as a function of signal contrast on semi-logarithmic coordinates, derived from the detectability functions shown in Figure 7a. First consider detection without noise and in notched noise.
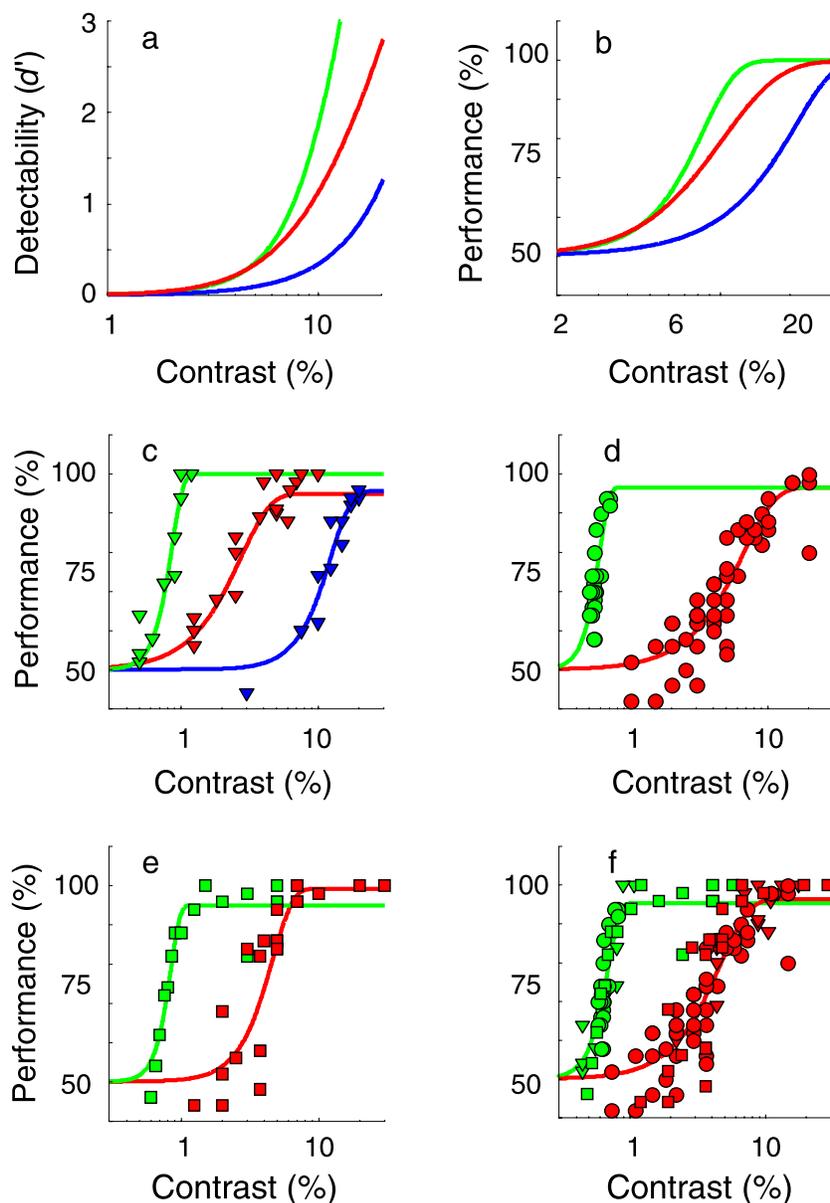

Figure 7. Illustration of the effect of channel pooling on contrast detection. (a) Signal detectability as a function of contrast for the pooled decision statistic on semi-logarithmic coordinates. Performance without noise is shown in green, in notched noise in red, and in broadband noise in blue. (b) The psychometric functions relating percentage correct in a 2AFC detection task to signal contrast (corresponding to the detectability functions shown in panel a). (c) Detection performance in a 2AFC task as a function of signal contrast for observer GBH. Green symbols refer to the no-noise condition, red symbols to the notched-noise condition, and blue symbols to the broadband-noise condition. (d) Detection performance for observer NAL. (e) Detection performance for observer TCC. (f) Detection performance as a function of (rescaled) signal contrast for all observers in the no-noise (green) and notched-noise (red) conditions.

According to the model, detection performance without noise should be better at most stimulus contrasts (see Figure 7b). Furthermore, the psychometric function relating performance (% correct) to signal contrast should be steeper without notched noise. Figures 7c–7e shows the percentage of correct responses as a function of signal contrast on semi-logarithmic coordinates obtained separately from three observers with no noise (green) and with (red) notched noise (Henning & Wichmann, 2007). For observer GBH, detection data with broadband noise (blue) are plotted as well. Fitted psychometric functions are also shown. As can be seen, the addition of notched noise not only hurts detection performance but also produces a more shallow psychometric function. Figure 7f shows the psychometric functions fitted to the combined data of all three observers, where the different symbols refer to different observers. (For each of the relevant conditions, the data for each observer have been scaled to equate the observers' 75% correct level within a given condition, thus increasing the statistical power of the test designed to compare the slopes of the psychometric functions.) It is clear that the data are consistent with the hypotheses formulated above: The signal contrast necessary to achieve 75% correct is lower in the no-noise condition (for each observer, $p < 10^{-6}$), and analysis of the β-parameter—the parameter controlling the steepness of the cumulative Weibull function fitted to the data—reveals that the psychometric function is also steeper than for the notched noise condition (for each observer, $p < 0.05$, for the combined data, $p < 10^{-6}$; Wichmann & Hill, 2001). These data are thus consistent with the notion that even detection of a sinusoidal grating may be based on pooled responses rather than on the most responsive channel.

For the model, the detection psychometric function in broadband noise is shallower than without noise but steeper than in notched noise if the noise-power density in the pass bands is kept constant. The 75% correct threshold, on the other hand, is expected to be higher in white noise than in both other conditions. Henning and Wichmann (2007) also measured detection in white noise, but noise power in the pass bands was only kept constant for observer GBH. Nevertheless, it is interesting to note that his data suggest that the detection psychometric function in white noise is indeed shallower than without noise ($p < 0.01$) but steeper than in notched noise ($p < 0.05$). Furthermore, as can be seen in Figure 7c, the signal contrast necessary to achieve 75% correct is higher in the white-noise condition than in both the no-noise condition ($p < 10^{-6}$) and the notched-noise condition ($p < 10^{-6}$). These data are thus consistent with the notion that noise may modify the neurons' contrast response functions, without, however, altering the pooling rules.

In summary, the population code model that successfully simulates contrast discrimination in noise was shown to predict effects of broadband-weighted pooling in detection as well. These predictions were tested and confirmed, thus suggesting that even the detection of a low-contrast sinusoidal grating may be based on the responses of many elements with a wide heterogeneity in spatial frequency tuning.

Larger pools with random spatial frequency tuning and correlated noise

A final consideration is the issue whether this model is able to produce realistic results in a less artificial situation. In the following simulation, we therefore consider a pool consisting of 250 correlated units tuned to a broad range of spatial frequencies.

In the type of network discussed in this paper, where the main noise source is signal-dependent, allowing more units to contribute to the decision statistic leads to better detection performance and a stronger pedestal effect. At first sight, larger pools may thus be expected to produce stronger pedestal effects. However, in this regard, it is important to note that single cell recordings have demonstrated that the responses of different cortical neurons in discrimination tasks are typically weakly correlated (Golledge, Panzeri, Zheng, Pola, & Scannell, 2003; Montani, Kohn, Smith, & Schultz, 2007; Panzeri, Golledge, Zheng, Tovée, & Young, 2001; Shadlen et al., 1996; Zohary, Shadlen, & Newsome, 1994). It is thus likely that channel responses are correlated to some extent. Indeed, at the psychophysical level, experimental evidence suggests that spatial-frequency channels may share some of their internal noise in contrast discrimination (Henning, Bird, & Wichmann, 2002).

Correlated noise has two fundamentally different effects (Averbeck, Latham, & Pouget, 2006): First, if unit activity is pooled according to a weighted average rule and, as proposed here, all weights are positive, correlated noise will substantially decrease the encoding capacity of the pool. Further, when the responses of units are correlated, pooling improvement with increasing numbers of units contributing to the pool approaches a limit once the number of units exceeds some critical number (Zohary et al., 1994). Second, correlated units might influence the computational decoding strategies appropriate for networks of neurons (Abbott & Dayan, 1999; Chen, Geisler, & Seidemann, 2006). If much additional information can be gained by taking into account the fact that neural noise is correlated, the decoding strategies used in the brain may be affected, i.e., it may be appropriate to introduce a whitening stage. For primary visual cortex (Golledge et al., 2003; Panzeri et al., 2001), correlations are estimated to be on the order of 0.1 to 0.15, but, Golledge et al. (2003), making use of information theory techniques to quantify the role of such small correlations, argue they would contribute less than 10% extra Shannon information in encoding visual information and it has been argued that such small correlations are unlikely to be taken into account in the decoding computations (Averbeck et al., 2006). Consequently, we did not apply whitening for our network.

Figure 8a illustrates the effects of correlated noise in our network. Detectability for a low-contrast grating (no pedestal) is plotted in Figure 8a as a function of the number of pooled units. Detectability is measured as the ratio of the mean response to its standard deviation, expressed as d'. Exactly one unit in the pool was optimally tuned to the signal. The selectivities of all other units were again randomly chosen from a Gaussian distribution

centered at 0.50 with a standard deviation of 0.17 and clipped at 0 and 1. The average inter-unit correlation is coded by the color indicated by the color bar on the right of the figure. If the noise is uncorrelated [the highest (red) curve], addition of more units progressively improves detectability and, in the limit, would yield an errorless observer. On the other hand, weakly correlated noise (the other curves) shows that the addition of more
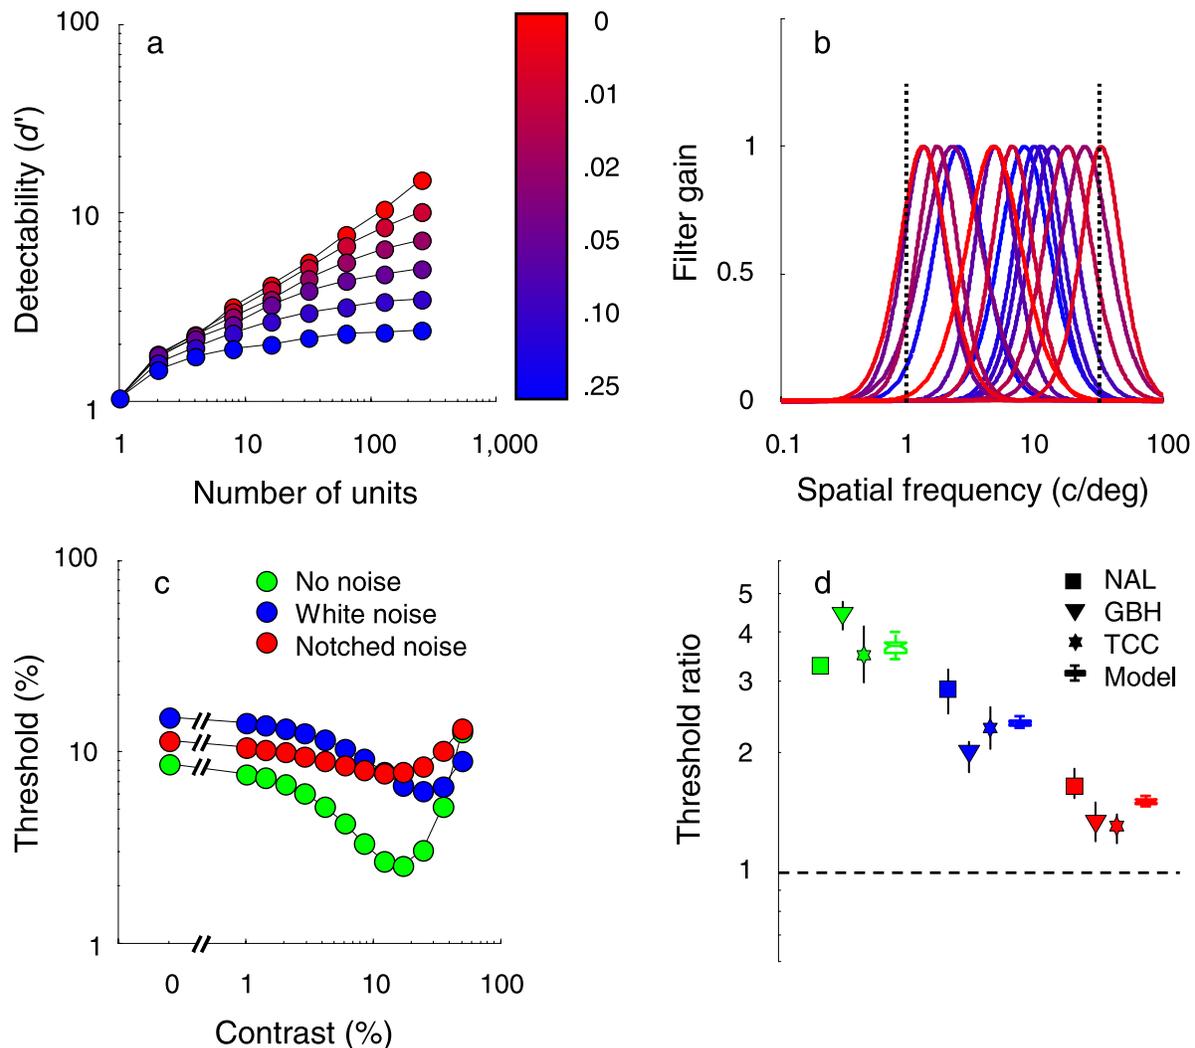

Figure 8. Contrast discrimination in larger pools with random spatial-frequency tuning and correlated noise. (a) The effect of correlated noise on network performance. Detectability of a low contrast sinusoidal grating is plotted as a function of pool size; color indicates the mean inter-unit correlation ranging from 0 (the topmost curve in red) to 0.15 (the bottom curve in blue). The standard deviation of this correlation in the simulations was chosen to be one tenth of the average correlation. Each estimate is based on 1,000 pools. (b) Spatial frequency tuning functions for 15 units of one particular pool of 250 units (for clarity, tuning functions of the other 235 units have been omitted). The dotted black lines indicate the boundaries of the geometrically uniform distribution from which the unit peak sensitivities were randomly sampled. (c) Simulated contrast-discrimination thresholds as a function of pedestal contrast with no external noise (green symbols), in broadband noise (blue symbols), and in notched noise (red symbols) for one particular pool of 250 units. (d) 75% correct threshold elevation ratios for the three observers that participated in the Henning and Wichmann (2007) experiments are shown for all noise conditions, indicated by the colors used in panel c. Error bars show the bootstrap-based, non-parameteric estimate of the 68% confidence intervals. Box plots summarize the same factors for 100 pools of 250 units. In all these box plots, the central horizontal line indicates the second quartile (i.e., the median threshold elevation ratio across pools); the other horizontal lines, where visible, indicate the first and third quartile. Whiskers indicate one and a half times the interquartile range. There were no outliers.

units has very little effect on detectability once a certain critical number of units is reached. As the average inter-unit correlation increases (downward or more blue in Figure 8a), that critical number of units drops. For inter-neuron correlations between 0.1 and 0.2 (typical estimates from single cell recordings), the improvement with increasing numbers approaches its asymptotic level between 50 and 100 units (Zohary et al., 1994). Pool size is thus expected to be rather limited.

In the following simulation, we consider pools of 250 units—thus well above the critical number—with an average inter-neuron correlation of 0.15 (the standard deviation of this correlation equals 0.05). We further postulate a spatial-frequency tuning function for each unit. To approximate tuning functions of cortical neurons (see e.g., Geisler & Albrecht, 1997), we opted to characterize spatial-frequency sensitivity with log-Gaussian-shaped tuning functions. Peak-sensitivity was randomly sampled

from an exponential distribution—and thus uniformly distributed on logarithmic coordinates—ranging between 1 and 33 cycles/deg. The average unit bandwidth equalled 1.5 octaves at half height, with a standard deviation of 0.2 octaves. Figure 8b illustrates the resulting tuning functions for 15 units from a pool (tuning functions for the other 235 units of the pool have been omitted, for the sake of clarity). The pools in this simulation are thus tuned to a broad range of spatial frequencies. Consequently, only few neurons are optimally sensitive for the 4-cycle/deg signal, while many are either somewhat sensitive or completely insensitive to the signal.

To test whether our model is able to produce realistic results with a broadly tuned pool consisting of units with correlated noise, 100 pools of 250 units were generated. The spatial-frequency tuning functions, contrast–response functions, and inter-unit correlations of these units were determined by randomly selecting parameter values from
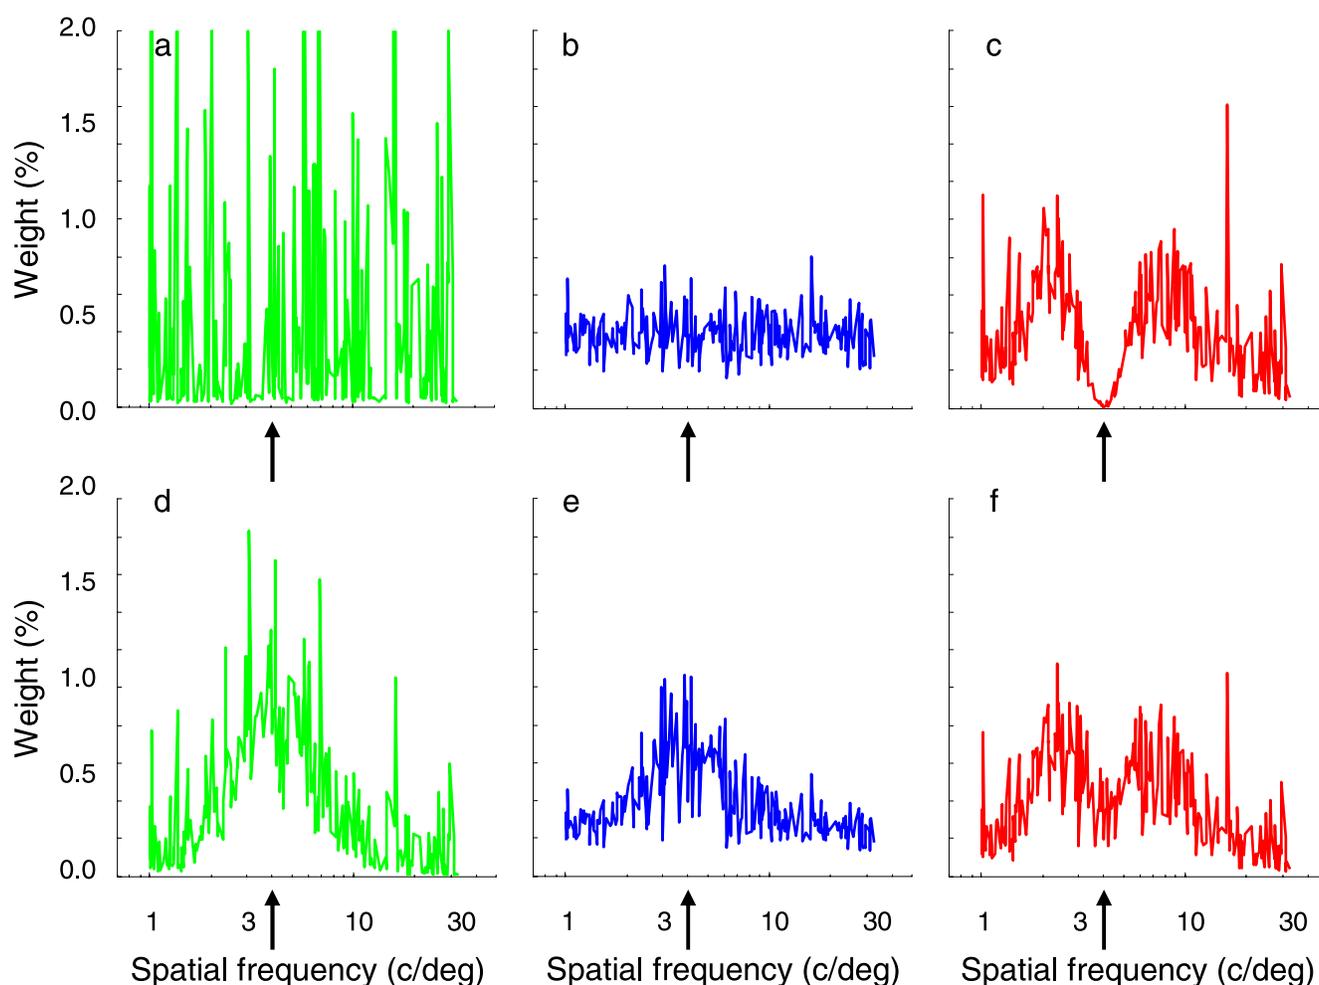

Figure 9. The effect of signal contrast and external noise on the weighting profile. (a–c) The distribution of weight at 0% signal contrast as a function of peak sensitivity for one particular pool of 250 units. Signal frequency is indicated by the arrow. (a) The weighting profile with no external noise added. (b) The weighting profile with broadband external noise added to the signal. (c) The weighting profile with notched noise added to the signal. (d–f) The distribution of weight at a signal contrast corresponding to the bottom of the dipper (i.e., approximately 15%) for the same noise conditions.

the appropriate distributions. The effects of noise were approximated as explained in the Contrast discrimination in noise: Effects of response pooling section; the effective noise power (i.e., the only free parameter in the model) was identical to the noise power used in the simulation shown in Figure 6. Contrast-discrimination thresholds in noise are shown for one of these pools in Figure 8c. These simulated results are very similar to the findings of Henning and Wichmann (2007). This can be seen in Figure 8d where the threshold elevation—expressing the strength of the pedestal effect—is plotted for all observers in the three noise conditions. Threshold elevation is the ratio of the 75% correct detection threshold (i.e., no pedestal) to the lowest contrast-discrimination threshold (bottom of the dipper) is plotted. Box plots summarize the same ratios in the same noise conditions for the 100 simulated pools. It is clear that our model produces contrast-discrimination data that closely mimic the psychophysical data. It can further be derived from this plot that the model behaves very robustly for pools of 250 units, as all box plots are remarkably small, indicating that the different pools produced very similar threshold elevations. We thus conclude that our model robustly produces plausible results for large pools with correlated noise and random spatial frequency tuning.

As explained in the section on weight assignment, response-based weighting produces a dynamical weighting profile, the exact shape of which depends on aspects such as signal contrast, unit selectivity, and other factors that affect units' responsiveness. Noise and the spectral characteristics of noise are such factors. It is thus interesting to "open" the model and to compare the weighting profiles in the different noise conditions. Figure 9 shows the distribution of weights as a function of the randomly sampled peak sensitivities for one particular pool of 250 in the absence of a signal (panels a–c) and at a contrast where the signal is just detectable (i.e., the signal contrast corresponds to the bottom of the dipper, panels d–f).

In the absence of a signal and external noise (Figure 9a), the weighting profile reflects nothing but the units' maintained discharge. With broadband noise added, all units respond to the noise in a similar fashion and thus produce a more regular, uniform weighting profile (Figure 9b). With notched noise added, response-based pooling produces a dip in the weighting profile (Figure 9c) because the responses of units tuned to frequencies within the notch are most strongly suppressed by the noise. Increasing the signal contrast to detection threshold strongly affects the weighting profiles (Figures 9d–9f). In the absence of external noise (Figure 9d), the weighting profile now peaks at the signal frequency (4 cycles/deg) and is approximately symmetrical on logarithmic coordinates. Units with peak-sensitivities more than one octave away from the signal frequency are usually not very selective for the signal and thus attract almost no weight. With broadband noise added (Figure 9e), the peak of the weighting profile is attenuated, while the tails are elevated

because units that are not tuned to the signal frequency are responding to the noise and thus attract some of the weight. With notched noise added (Figure 9f), response-based pooling produces a rippled weighting profile. The responses of units tuned to frequencies within the notch are most strongly suppressed by the noise, while units tuned to frequencies that are approximately one octave away from the signal frequency now attract most of the weight. In this regard, notched noise produces more rather than less off-frequency looking in our model—the biggest weights are off the signal frequency but within the notch.

Discussion

Current models of spatial vision cannot easily explain why the pedestal effect persists in broadband noise but disappears in notched noise. Indeed, predictions of the standard version of the psychophysical divisive inhibition model (Foley's model 3) are not in line with these findings. Further, there is a discrepancy between the deep dip attributed to single channels by these models and the mild dip observed in single cells of the striate cortex. If one assumes that notched noise prevents or reduces the use of information carried by channels tuned to frequencies other than the signal frequency, these findings suggest that the pedestal effect stems from off-frequency looking (Henning & Wichmann, 2007). However, the assumption that notched noise prevents off-frequency looking might be wrong.

At single cell level, the normalization model is consistent with many different observations on cortical cell behavior. Although the available data are limited, this model has also been shown to capture the effects of broadband noise reasonably well. Our implementation of noise effects is fully based on the logic of the normalization model. It is important to note that for notched noise, this implementation does not result in reduced information pooling across units. To the contrary, notched noise modifies units' contrast response functions, without preventing off-frequency looking. In this case, the dipper disappears in notched noise for two main reasons: First, because the notched noise, through the inhibitory effect of noise outside the notch on units tuned to frequencies within the notch, linearizes those units' response. Second, because notched noise prevents the (large) improvement in the signal-to-noise ratio of the responses of the units tuned to the signal frequency that occurs without noise and to a lesser degree in broadband noise. (Stated differently, in notched noise the pedestal does not lead to a stronger contribution of the "relevant" channels to the decision statistic.)

Here, we have shown that in a network consisting of units whose contrast response functions resemble those of cortical cells, weighted summation—with weights based simply on the magnitude of the response—produces

contrast-discrimination data that resemble many aspects of psychophysical observations. Similar to earlier neurophysiologically based models of vision, the model includes a spontaneous firing rate (a stimulus-independent base firing rate) and mimics the physiologically observed proportionality between mean and variance of the firing rate with signal-dependent noise. The pedestal effect in the model, however, arises because of information combination across units.

The model predicts not only the standard dipper effect but also how the dipper changes when spectrally flat and notch-filtered noise is added. Finally, the model is consistent with neurophysiological estimates of simple-cell contrast–response functions and thus, irrespective of the specific parameter settings of the model, resolves the discrepancy between single cell contrast–response functions, which display a weak or even no pedestal effect and the strong pedestal effect observed psychophysically.

In the model, we find that the detailed statistics of the components of the narrowband stimulus hardly matter: performance is principally determined by the signal-to-noise ratio of the decision statistic, which is based on the combined output of both sensitive and relatively insensitive units. Indeed, human contrast-discrimination performance has been reported to be largely independent of signal frequency (Bird et al., 2002).

Using information from non-optimally tuned cells is probably not a unique feature of contrast discrimination. For instance, Shadlen et al. (1996) considered the relation between behavioral and neurophysiological (MT) responses to visual motion and found that non-optimally tuned cells needed to be postulated and included in the neural pool in order to reconcile their behavioral and neurophysiological measures.

Employing a simple response-based weighting heuristic is a sensible strategy because under most realistic conditions, i.e., at sufficiently high contrasts, it resembles an "optimal" combination rule (for a maximum likelihood combination rule for detection, see Jazayeri & Movshon, 2006). Nevertheless, one may wonder how the visual system would manage to weight the responses of different units differently and even adjust these weights on a trial-to-trial basis. It may be helpful to notice that our particular decoding rule can also be thought of as the summed output of a neuronal layer in which the responses of (a reasonably large number of) randomly sampled V1 cells are squared and, perhaps, normalized by the (non-squared) responses of a similarly large random sample of V1 cells. The normalization does not alter model predictions, but the squaring is crucial to capture response-based weighting.

In this paper, we did not discuss effects of orientation tuning. There is, however, no reason to assume that pooling is limited to the spatial frequency dimension. Indeed, the selectivity parameter may be thought off as expressing effects of either or both spatial-frequency and orientation tuning. Consequently, our model predicts that performing a contrast-discrimination experiment in orientation-filtered noise will produce similar effects as the notched noise effects of Henning and Wichmann (2007).

Conclusion

Recent evidence has demonstrated that the pedestal effect in spatial vision is differently modified by spectrally flat and notch-filtered noise. Here, we have shown that a network consisting of units whose contrast response functions resemble those of the cortical cells believed to underlie human pattern vision can produce contrast-discrimination data consistent with psychophysical observations when the outputs of multiple units tuned to a range of spatial frequencies are combined by simple weighted summation. One implication of these findings is that even in processing low-contrast sinusoidal gratings, as in detection, the visual system may combine information across neurons tuned to different spatial frequencies and orientations.

Appendix A

To make the simulations tractable, we simplified the calculation of the weights by using mean responses (Equation 3). This simplification allowed us to estimate the variability of the pooled network response (Equation 5) and thus network discrimination performance (Equation 6) directly for any given combination of pedestal and signal stimulus, without having to simulate too many trials. However, in a real nervous system, weights would necessarily be based on responses alone. Compared with the simplified implementation, there is now trial-to-trial variation in the weights. This additional source of variation does not, however, change the important parts of the networks' behavior described in the paper. The crucial observation here is that in assigning weights, nothing need be known about the variance or indeed any other characteristic of a unit—the weights are determined solely by the strength of the unit's response.

This point is clarified by a simulation, the results of which are shown in Figure A1. Figure A1a illustrates how the detectability of a narrowband stimulus depends on stimulus contrast. Results for the most selective unit in a pool of 250 units—tuned to a broad range of spatial frequencies and with correlated noise included—are indicated by the purple symbols. Results for the pooled network response as approximated in the paper, i.e., by average-response-based weighting, are shown in green. Results for the pooled network response produced by simulating 2,500 trials for all signal contrasts and performing trial-by-trial response-based weighting are shown in black. It is clear that both pooled response functions outperform the most selective unit and are more sharply accelerated.
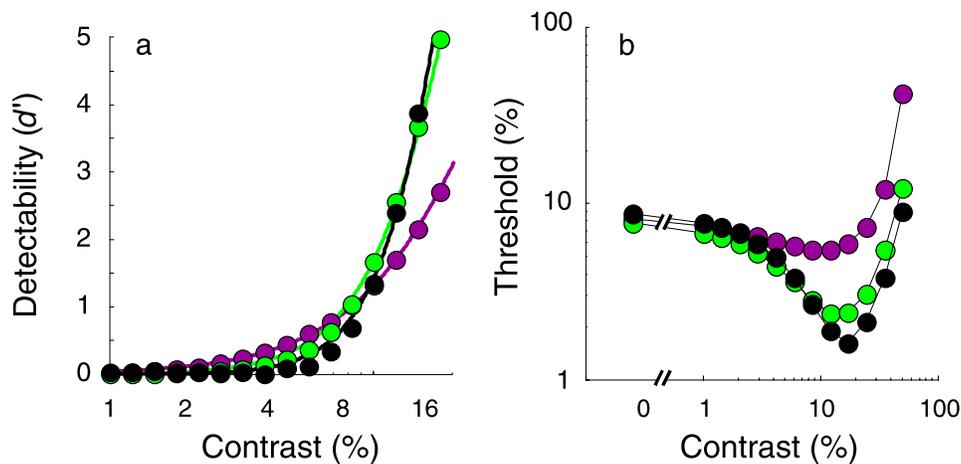
Figure A1. Average-response-based and response-based pooling. (a) Detectability, d', expressed as the ratio of the mean response to the standard deviation, as a function of contrast for the most sensitive unit (purple symbols) in a pool of 250 units, average-response-based pooling (green symbols), and trial-by-trial response-based pooling (black symbols) on semi-logarithmic coordinates. (b) Contrast-discrimination thresholds at 75% correct as a function of pedestal contrast on double logarithmic coordinates for the most sensitive unit (purple), average-response-based pooling (green), and trial-by-trial response-based pooling (black).

Consequently, the pooled response functions give rise to a stronger pedestal effect than the single most selective unit whether the effects of the pooling are derived from mean responses or, as would be the case in a real nervous system, from the trial by trial response strength. Estimated contrast-discrimination threshold functions are plotted in Figure A1b, making use of the same color conventions as in panel a. Discrimination thresholds were derived from a descriptive function—$d' = \frac{bC^p}{Z^q + C^q}$—fitted to the detectability results shown in Figure A1a (the fits are indicated by the colored lines in panel a). While the thresholds shown in green in panel b lie below the thresholds shown in black at low pedestal contrasts, it is clear that both pooled response functions produce a stronger pedestal effect than the most selective unit. In sum, the simplification in determining the weights is reasonable for the issues discussed in this paper.

Appendix B

To capture the effects of external noise on the contrast response function of units in our model, we used theory-based predictions combined with simulations. The basis of most models of cortical neurons is the concept of linear receptive field, followed by an instantaneous non-linear function. Broadband noise introduces stimulus variability at the preferred spatial frequency and phase of a linear filter and will thus increase the filter's response variance. Because neurons cannot give negative responses, this increased variance also increases the mean cell response at low response rates. At high response rates, this increased variance lowers the mean cell response slightly due to the non-linearity. To see all this, consider Figure B1, which shows how the contrast response function of a unit to a

preferred signal changes with different levels of external noise (panel a shows the mean response, panel b the variance of the response; lighter symbols refer to higher noise levels). To obtain these results, the simulated output of a linear filter was half-wave rectified and passed through the Naka–Rushton equation (see Equation B1), without any further rescaling:

$$R_{u,bn}(c) = \frac{\left(\max\left(0, R_f(c)\right)\right)^n}{c_{50}^n + \left(\max\left(0, R_f(c)\right)\right)^n},$$ (B1)

where $R_f(c)$ is the response of a linear filter as a function of signal contrast expressed as a fraction of the maximal response. As in Equation 2, the variance of the unit's response was proportional to its mean value.

For the unit shown in Figure B1, n equals 2.4 and c_{50} equals 0.38. We performed simulations for a wide range of parameter values and noise levels and found that the results could be captured by resetting the parameters of the unit's response function as given by Equation B2 (fits to the mean response are shown in Figure B1).

$$\overline{R_{u,bn}}(c) = r_{\text{Noise}} + \frac{c^{n-\Delta n}}{(c_{50} + \Delta c)^{n-\Delta n} + c^{n-\Delta n}},$$ (B2)

where r_{Noise} is the average noise evoked response, and Δn and Δc describe the change of the response exponent and semisaturation contrast in noise. Based on our simulations, these three parameters were estimated analytically for each network unit as a function of external noise level (this is a free parameter in the model), the response exponent n, and semisaturation contrast c_{50} (these are randomly selected parameters, as explained in Methods).
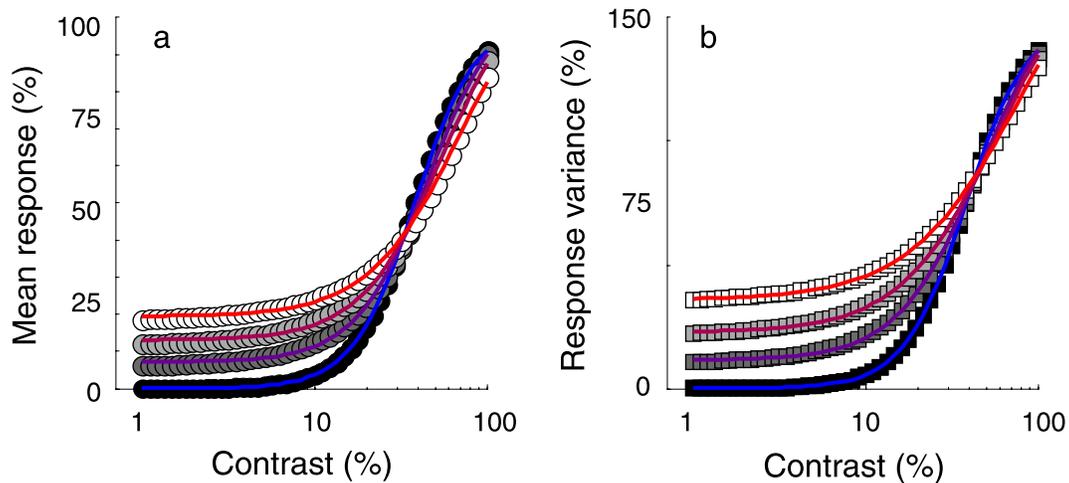
Figure B1. (a) The simulated mean contrast response functions for a given unit with different levels of external noise, making use of Equation B1 (lighter symbols refer to higher noise levels). Responses are expressed as a fraction of the maximal response. Colored lines indicate fits of the descriptive model discussed in the text (see Equation B2). (b) The simulated response variance for the same unit.

Compared with cortical cell data, the noise activation at low firing rates predicted by this simple model is too high. Similarly, the noise suppression at high firing rates is too low (Carandini et al., 1997). Both failures can be cured by an additional parameter α introduced by Carandini et al. (1997) to capture non-specific suppression effects of broadband noise, as given by Equation B3.

$$\overline{R_{u,bn}}(c) = \frac{r_{\text{Noise}}}{1 + \alpha NC}$$
$$+ \text{Sel}\left(\frac{c^{n-\Delta n}}{(c_{50} + \Delta c)^{n-\Delta n} + c^{n-\Delta n} + (\alpha NC)^{n-\Delta n}}\right).$$
$$(B3)$$

As in Equation 1, Sel expresses the unit's selectivity. Parameter values for α were drawn from an exponential distribution, appropriately scaled to approximate the estimates reported by Carandini et al. (1997). At the noise power chosen for the contrast-discrimination simulations in the paper, the response to noise alone was on average approximately four times higher than the spontaneous maintained discharge, r_0. This estimate is reasonably close to the roughly three-fold elevation found by Carandini et al. Further, at this noise level, $\Delta n = 0.45$ (standard deviation of 0.15 across units) and $\Delta c = 0.11$ (standard deviation of 0.03). Figure B2 illustrates the effects of broadband noise on the contrast response function for one unit, simulated with the normalization model—i.e., a
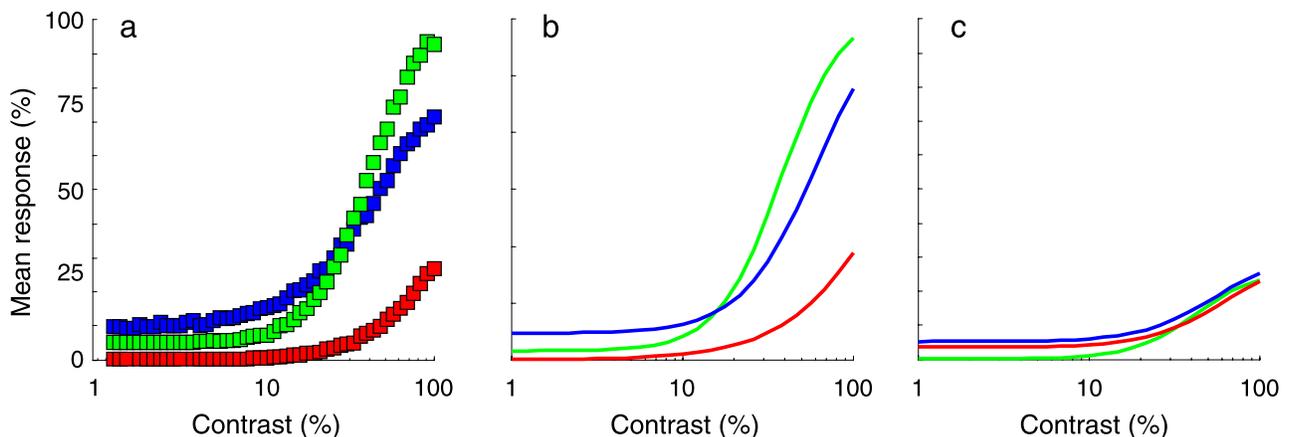

Figure B2. (a) The contrast response functions simulated with the normalization model for a unit tuned to the signal without noise (green), in broadband noise (blue), and notched noise (red). Without noise, the excitatory and inhibitory factors of the normalization model are constant at any given contrast; in broadband noise, both factors increase and show trial-to-trial variation; in notched noise, the response of the excitatory factor is constant, while the inhibitory factor increases and produces a variable response. (b) Illustration of the effective mean contrast response functions used in the paper for a given unit without noise (green), in broadband noise (blue), and notched noise (red). This unit is tuned to the signal frequency. (c) Same as in panel b for a unit that is not sensitive for the signal frequency. Note that the responses to notched noise and white noise are much more similar for such units.

narrowly tuned excitatory factor and a broadly tuned divisive inhibitory factor, both producing variable responses—and our approximation, making use of Equation B3.

The effect of notched noise can be inferred from the logic of the normalization model (Heeger, 1992a, 1992b). First, consider a unit tuned to the 1.5 octaves wide notch centered at the signal frequency. Notched noise contains no power at the unit's preferred spatial frequency and phase and will thus not elicit any response by itself. However, the noise has power at frequencies to which the broadband inhibitory gain-control pool is tuned. The normalization model thus predicts that notched noise will produce non-specific inhibition. Given that noise without notch elevates the mean cell firing rate, the inhibitory effect of notched noise is likely to be stronger than the inhibition provided by white noise. For the simulations in the paper, the inhibition provided by notched noise was chosen to be two times higher than the inhibition provided by white noise. (This factor may also partly capture the fact that Henning and Wichmann, 2007, increased the power in the pass-bands of the notched noise relative to the white noise for two of three observers.) Because the inhibitory response is largely noise driven and thus variable, notched noise also modifies the response exponent and semi-saturation contrast. As was the case for broadband noise, parameters Δn and Δc used to capture this modification were estimated for each unit based on the simulations of the simple model described above (see Equation B1). This is of course only an approximation but sufficient to capture the increase of Δn and Δc with stronger suppression. This can be seen in Figure B2, which illustrates the effects of notched noise on the contrast response function for one unit, simulated with the normalization model and our approximation, making use of Equation B4.

The effects of notched noise differ for a unit tuned to one of the pass-bands of the noise (see Figure B2, panel c). In short, as the selectivity of the unit for the signal decreases, the notched noise will elicit both excitatory and inhibitory activation and its effects will gradually approximate the effects of broadband noise. The effect of notched noise thus depends strongly on a unit's selectivity. Equation B4 shows how we implemented all these effects of notched noise on the contrast response function.

$$
\begin{aligned}
\overline{R_{u,nn}}\,(c) = &\ \frac{(1 - \mathrm{Sel})r_{\mathrm{Noise}}}{1 + \alpha NC} \\
&+ \mathrm{Sel}\left(\frac{c^{n-\Delta n}}{(c_{50} + \Delta c)^{n-\Delta n} + c^{n-\Delta n} + ([1+\mathrm{Sel}]\alpha NC)^{n-\Delta n}}\right).
\end{aligned}
$$
(B4)

At the noise level used for the contrast-discrimination simulations in the paper, $\Delta n = 0.43$ (standard deviation of 0.3 across units) and $\Delta c = 0.18$ (standard deviation of 0.2). The average modification of the response exponent in notched noise thus closely resembles the results in white

noise, while the noise-suppression is stronger. For both parameters, the standard deviation in notched noise is higher due to the effects of tuning. While this implementation captures the main effects of notched noise described above, it is a simplification and at best only an approximation. Nevertheless, this operationalization of notched noise effects proved to be sufficient to generate plausible contrast-discrimination data in noise.

For the sake of clarity, the variability in maximal response rate across units was ignored in this appendix. In the model simulations discussed in the paper, however, Equations B3 and B4 were multiplied by each unit's $r_{\max}$. Finally, to estimate a unit's response variance, we used the ratio of the variance to the mean simulated in the simple model (Equation B1) and multiplied this ratio with the mean responses deduced from Equations B3 and B4.

Corresponding author: Robbe L. T. Goris.
Email: Robbe.Goris@psy.kuleuven.be.
Address: Laboratory of Experimental Psychology, Tiensestraat 102, B-3000 Leuven, Belgium.

References

Abbott, L. F., & Dayan, P. (1999). The effect of correlated variability on the accuracy of a population code. *Neural Computation, 11,* 91–101. [PubMed]

Albrecht, D. G., Geisler, W. S., Frazor, R. A., & Crane, A. M. (2002). Visual cortex neurons of monkeys and cats: Temporal dynamics of the contrast response function. *Journal of Neurophysiology, 88,* 888–913. [PubMed] [Article]

Averbeck, B. B., Latham, P. E., & Pouget, A. (2006). Neural correlations, population coding and computation. *Nature Reviews, Neuroscience, 7,* 358–366. [PubMed]

Bird, C., Henning, G. B., & Wichmann, F. A. (2002). Contrast discrimination with sinusoidal gratings of

different spatial frequency. *Journal of the Optical Society of America A, Optics, Image Science, and Vision, 19,* 1267–1273. [PubMed]

Blakemore, C., & Campbell, F. W. (1969). On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images. *The Journal of Physiology, 203,* 237–260. [PubMed] [Article]

Campbell, F. W., & Robson, J. G. (1968). Application of Fourier analysis to the visibility of gratings. *The Journal of Physiology, 197,* 551–556. [PubMed] [Article]

Carandini, M. (2004). Amplification of trial-to-trial response variability by neurons in visual cortex. *PLoS Biology, 2,* e264. [PubMed] [Article]

Carandini, M., Heeger, D. J., & Movshon, J. A. (1997). Linearity and normalization in simple cells of the macaque primary visual cortex. *Journal of Neuroscience, 17,* 8621–8644. [PubMed]

Carandini, M., Heeger, D., & Senn, W. (2002). A synaptic explanation of suppression in visual cortex. *Journal of Neuroscience, 22,* 10053–10065. [PubMed] [Article]

Chen, Y., Geisler, W. S., & Seidemann, E. (2006). Optimal decoding of correlated neural population responses in the primate visual cortex. *Nature Neuroscience, 9,* 1412–1420. [PubMed] [Article]

Chirimuuta, M., & Tolhurst, D. J. (2005). Does a Bayesian model of V1 contrast coding offer a neurophysiological account of human contrast discrimination? *Vision Research, 45,* 2943–2959. [PubMed]

Deneve, S., Latham, P. E., & Pouget, A. (1999). Reading population codes: A neural implementation of ideal observers. *Nature Neuroscience, 2,* 740–745. [PubMed]

DeValois, R. L., & DeValois, K. K. (1988). *Spatial vision.* Oxford: Oxford University Press.

Eliasmith, C., & Anderson, C. H. (2003). *Neural Engineering: Computation, representation, and dynamics in neurobiological systems.* Cambridge: MIT Press.

Finn, I. M., Priebe, N. J., & Ferster, D. (2007). The emergence of contrast-invariant orientation tuning in simple cells of cat visual cortex. *Neuron, 54,* 137–152. [PubMed] [Article]

Foley, J. M. (1994). Human luminance pattern-vision mechanisms: Masking experiments require a new model. *Journal of the Optical Society of America A, Optics, Image Science, and Vision, 11,* 1710–1719. [PubMed]

Frazor, R. A., & Geisler, W. S. (2006). Local luminance and contrast in natural images. *Vision Research, 46,* 1585–1598. [PubMed]

Freeman, T. C. B., Durand, S., Kiper, D. C., & Carandini, M. (2002). Suppresion without inhibition in visual cortex. *Neuron, 35,* 759–771. [PubMed]

Geisler, W. S., & Albrecht, D. G. (1997). Visual cortex neurons in monkeys and cats: Detection, discrimination, and identification. *Visual Neuroscience, 14,* 897–919. [PubMed]

Golledge, H. D. R., Panzeri, S., Zheng, F., Pola, G., & Scannell, J. W. (2003). Correlations, feature-binding and population coding in primary visual cortex. *Neuroreport, 14,* 1045–1050. [PubMed]

Goris, R. L. T., Wagemans, J., & Wichmann, F. A. (2008a). Modelling contrast discrimination data suggest both the pedestal effect and stochastic resonance to be caused by the same mechanism. *Journal of Vision, 8*(15):17, 1–21, http://journalofvision.org/8/15/17/, doi:10.1167/8.15.17. [PubMed] [Article]

Goris, R. L. T., Zaenen, P., & Wagemans, J. (2008b). Some observations on contrast detection in noise. *Journal of Vision, 8*(9):4, 1–15, http://journalofvision.org/8/9/4/, doi:10.1167/8.9.4. [PubMed] [Article]

Graham, N. (1989). *Visual pattern analyzers.* Oxford: Oxford University Press.

Graham, N., & Nachmias, J. (1971). Detection of grating patterns containing two spatial frequencies: A comparison of single-channel and multiple-channels models. *Vision Research, 11,* 251–259. [PubMed]

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics.* New York: John Wiley & Sons, Inc.

Hammond, P., & MacKay, D. M. (1977). Differential responsiveness of simple and complex cells in cat striate cortex to visual texture. *Experimental Brain Research, 30,* 275–296. [PubMed]

Heeger, D. J. (1992a). Half-squaring in responses of cat striate cells. *Visual Neuroscience, 9,* 427–443. [PubMed]

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

Henning, G. B., Bird, C. M., & Wichmann, F. A. (2002). Contrast discrimination with pulse trains in pink noise. *Journal of the Optical Society of America A, Optics, Image Science, and Vision, 19,* 1259–1266. [PubMed]

Henning, G. B., & Wichmann, F. A. (2007). Some observations on the pedestal effect. *Journal of Vision, 7*(1):3, 1–15, http://journalofvision.org/7/1/3/, doi:10.1167/7.1.3. [PubMed] [Article]

Jazayeri, M., & Movshon, J. A. (2006). Optimal representation of sensory information by neural populations. *Nature Neuroscience, 9,* 690–696. [PubMed]

Legge, G. E., & Foley, J. M. (1980). Contrast masking in human vision. *Journal of the Optical Society of America, 70,* 1458–1471. [PubMed]

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

Miller, K. D., & Troyer, T. W. (2002). Neural noise can explain expansive, power-law nonlinearities in neural response functions. *Journal of Neurophysiology, 87,* 653–659. [PubMed] [Article]

Montani, F., Kohn, A., Smith, M. A., & Schultz, S. R. (2007). The role of correlations in direction and contrast coding in the primary visual cortex. *Journal of Neuroscience, 27,* 2338–2348. [PubMed]

Nachmias, J. (1981). On the psychometric function for contrast detection. *Vision Research, 21,* 215–223. [PubMed]

Nachmias, J., & Sansbury, R. V. (1974). Grating contrast: Discrimination may be better than detection. *Vision Research, 14,* 1039–1042. [PubMed]

Olshausen, B. A., & Field, D. J. (2005). How close are we to understanding V1? *Neural Computation, 17,* 1665–1699. [PubMed]

Panzeri, S., Golledge, H. D. R., Zheng, F., Tovée, M. J., & Young, M. P. (2001). Objective assessment of the functional role of spike train correlations using information measures. *Visual Cognition, 8,* 531–547.

Pelli, D. G. (1985). Uncertainty explains many aspects of visual contrast detection and discrimination. *Journal of the Optical Society of America A, Optics and Image Science, 2,* 1508–1532. [PubMed]

Pouget, A., Zemel, R. S., & Dayan, P. (2000). Information processing with population codes. *Nature Reviews. Neuroscience, 1,* 125–132. [PubMed]

Ringach, D. L., Hawken, M. J., & Shapley, R. (1997). Dynamics of orientation tuning in macaque primary visual cortex. *Nature, 387,* 281–284. [PubMed]

Shadlen, M. N., Britten, K. H., Newsome, W. T., & Movshon, J. A. (1996). A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *Journal of Neuroscience, 16,* 1486–1510. [PubMed] [Article]

Shadlen, M. N., & Newsome, W. T. (1998). The variable discharge of cortical neurons: Implications for connectivity, computation and information coding. *Journal of Neuroscience, 18,* 3870–3896. [PubMed] [Article]

Smithson, H. E., Henning, G. B., MacLeod, D. I. A., & Stockman, A. (2009). The effect of notched noise on flicker detection and discrimination. *Journal of Vision, 9*(5):21, 1–18, http://journalofvision.org/9/5/21/, doi:10.1167/9.5.21.

Vogels, R., Spileers, W., & Orban, G. (1989). The response variability of striate cortical neurons in the behaving monkey. *Experimental Brain Research, 77,* 432–436. [PubMed]

Wichmann, F. A. (1999). *Some aspects of modelling human spatial vision: Contrast discrimination.* Unpublished doctoral dissertation. Oxford, UK: The University of Oxford.

Wichmann, F. A., & Hill, N. J. (2001). The psychometric function: II. Bootstrap-based confidence intervals and sampling. *Perception & Psychophysics, 63,* 1314–1329. [PubMed] [Article]

Zohary, E., Shadlen, M. N., & Newsome, W. T. (1994). Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature, 370,* 140–143. [PubMed]