

Neural Noise Can Explain Expansive, Power-Law Nonlinearities in Neural Response Functions

KENNETH D. MILLER¹ AND TODD W. TROYER²

¹Departments of Physiology and Otolaryngology, W. M. Keck Center for Integrative Neuroscience, Sloan Center for Theoretical Neurobiology, University of California, San Francisco, California 94143-0444; and ²Department of Psychology, Neuroscience and Cognitive Science Program, University of Maryland, College Park, Maryland 20742

Received 22 May 2001; accepted in final form 19 October 2001

Miller, Kenneth D. and Todd W. Troyer. Neural noise can explain expansive, power-law nonlinearities in neural response functions. *J Neurophysiol* 87: 653–659, 2002; 10.1152/jn.00425.2001. Many phenomenological models of the responses of simple cells in primary visual cortex have concluded that a cell's firing rate should be given by its input raised to a power greater than one. This is known as an expansive power-law nonlinearity. However, intracellular recordings have shown that a different nonlinearity, a linear-threshold function, appears to give a good prediction of firing rate from a cell's low-pass-filtered voltage response. Using a model based on a linear-threshold function, Anderson et al. showed that voltage noise was critical to converting voltage responses with contrast-invariant orientation tuning into spiking responses with contrast-invariant tuning. We present two separate results clarifying the connection between noise-smoothed linear-threshold functions and power-law nonlinearities. First, we prove analytically that a power-law nonlinearity is the only input-output function that converts contrast-invariant input tuning into contrast-invariant spike tuning. Second, we examine simulations of a simple model that assumes instantaneous spike rate is given by a linear-threshold function of voltage and voltage responses include significant noise. We show that the resulting average spike rate is well described by an expansive power law of the average voltage (averaged over multiple trials), provided that average voltage remains less than about 1.5 SDs of the noise above threshold. Finally, we use this model to show that the noise levels recorded by Anderson et al. are consistent with the degree to which the orientation tuning of spiking responses is more sharply tuned relative to the orientation tuning of voltage responses. Thus neuronal noise can robustly generate power-law input-output functions of the form frequently postulated for simple cells.

INTRODUCTION

Responses of visual cortical simple cells are commonly described by simple phenomenological models in which a linear filtering of the stimulus is followed by an expansive power-law nonlinearity to determine an instantaneous firing rate (Albrecht and Geisler 1991; Albrecht and Hamilton 1982; Anzai et al. 1999; Carandini et al. 1997, 1999; Emeryson et al. 1989; Gardner et al. 1999; Heeger 1992, 1993; Murthy et al. 1998; Sclar et al. 1990). By a power-law nonlinearity, we mean that a cell's firing rate r depends on its input or voltage V as $r = k([V]^+)^n$ for constants k and n , where the 0 voltage is set equal to the mean voltage at rest

and $[V]^+ = \max(V, 0)$.¹ By an expansive nonlinearity, we mean that $n > 1$. The linear filtering aspect of this model receives support from a number of studies showing that voltage responses of simple cells are remarkably linear functions of the visual stimulus (Anderson et al. 2000; Jagadeesh et al. 1993, 1997; Lampl et al. 2001) (but see DISCUSSION).

Despite the success of using an expansive power law nonlinearity in phenomenological models, direct experimental investigations have shown that the transformation from instantaneous voltage to instantaneous spike rate is well approximated by a linear-threshold function $r = k([V - T]^+)$, where V is the voltage after removal of spikes and low-pass filtering, r is the low-pass-filtered spike train, and T is an effective spike threshold (Anderson et al. 2000; Carandini and Ferster 2000). In this paper, we show a simple and surprisingly robust connection between linear-threshold models and expansive power-law nonlinearities: if the voltage trace includes significant stimulus-independent noise and if the conversion from instantaneous voltage to instantaneous firing rate is a linear threshold function, then the conversion from trial-averaged voltage to trial-averaged firing rate will be well described by an expansive power law (cf. Suarez and Koch 1989).

This work is inspired by the recent results of Anderson et al. (2000). They studied the intracellular basis for the observation that orientation tuning of visual cortical neurons is contrast invariant, i.e., changing stimulus contrast simply scales the magnitude of a neuron's response, without changing the shape of its orientation tuning curve (Sclar and Freeman 1982; Skottun et al. 1987). Anderson et al. (2000)'s results can be separated into three main findings. First, they found that a cell's trial-averaged voltage response \bar{V} showed contrast-invariant orientation tuning (both the size of the mean voltage response and the amplitude of voltage modulation had this property). Second, they found that spiking responses also showed contrast-invariant tuning (as expected) and that a given neuron's spiking response was more narrowly tuned for orientation than its intracellular voltage response. Finally, they used computer

¹ We use the rectified voltage $[V]^+$ in our definition of a power law to ensure that responses are an increasing function of voltage, i.e., we assume that hyperpolarization cannot increase response.

The costs of publication of this article were defrayed in part by the payment of page charges. The article must therefore be hereby marked "advertisement" in accordance with 18 U.S.C. Section 1734 solely to indicate this fact.

Address for reprint requests: K. D. Miller, Dept. of Physiology, UCSF, San Francisco, CA 94143-0444 (E-mail: ken@phy.ucsf.edu).

simulations to demonstrate that a linear threshold model for converting instantaneous voltage to instantaneous spike rate could account for this data, but that the noise in the actual voltage traces was critical for this result: if a linear threshold model was applied to the actual noisy voltage traces, or to the trial-averaged voltage traces with the addition of random noise, then contrast-invariant spiking tuning was attained; but if a linear threshold model was simply applied to the trial-averaged voltage traces without additional noise, the spiking tuning became contrast dependent.

Here we present two basic results that serve to clarify the connection between the work of Anderson et al. (2000) and the many successful phenomenological descriptions of simple cell responses that assume linear input plus an expansive power law nonlinearity. First, while others have noted that a power law nonlinearity converts contrast-invariant input into contrast-invariant output (e.g., Carandini et al. 1997; Heeger 1992; Heeger et al. 1996), here we prove that a power law is the *only* function that achieves this. This result, in combination with the results of Anderson et al. (2000), implies that adding noise to a linear threshold function *must* yield power law behavior to a good approximation. Second, we quantify the degree to which a noise-smoothed linear threshold function can indeed be approximated by a power law, finding that the approximation holds over a wide range of parameters. We also show that the exponent in the best-fit power law decreases with increasing noise level and that the sizes of signal and noise measured by Anderson et al. (2000) predict an exponent that accounts well for the observed sharpening of spiking orientation tuning relative to voltage tuning.

RESULTS

Contrast-invariant tuning

A response function is contrast invariant if changes in stimulus contrast simply scale responses without affecting the tuning to other parameters. For example, let $V(c, \theta)$ describe the voltage response of a given neuron as a function of contrast level c and orientation θ . Let $g_V(\theta)$ be the orientation tuning function at maximal contrast. If the cell displays contrast-invariant orientation tuning, then changing to a different contrast c simply scales the response, i.e., $V(c, \theta) = f_V(c)g_V(\theta)$, where $f_V(c)$ is the cell's contrast response function (Fig. 1). Therefore saying that orientation tuning is contrast invariant is equivalent to saying that the contrast response function is orientation invariant.

We will work in units in which $V = 0$ represents the mean voltage at rest, in the absence of a stimulus, so that V represents the stimulus-induced voltage. It is easy to see that if the voltage response is contrast invariant, and the stimulus-induced spike rate $R(V)$ (i.e., the spike rate after subtracting off the background rate) is equal to the stimulus-induced voltage raised to a power n , then this spike rate is contrast invariant:

$$R(V) = k([V]^+)^n = k([f_V(c)g_V(\theta)]^+)^n = kf_V(c)^n([g_V(\theta)]^+)^n \quad (1)$$

(In the last step, we used the fact that the contrast response function is nonnegative.) In APPENDIX A, we prove the converse, i.e., a power law is the only function that transforms contrast-invariant inputs into contrast-invariant spiking responses.

Only a pure power law yields contrast invariant responses.

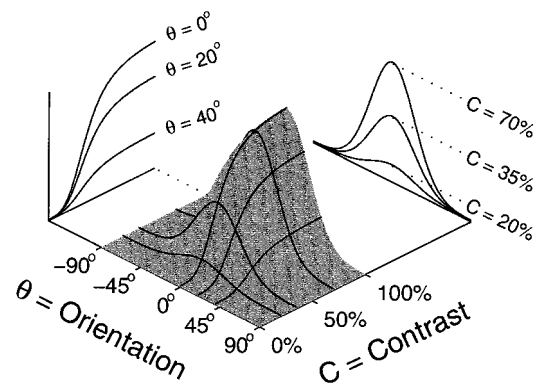


FIG. 1. Illustration of a response function with contrast-invariant orientation tuning. Vertical axis shows response $R(c, \theta)$ to contrast c and orientation θ . The response separates into the product of a function of contrast times a function of orientation, $R(c, \theta) = f(c)g(\theta)$. This condition ensures that orientation tuning is simply scaled by changes in contrast as illustrated by the projections along the orientation axis for different contrasts. This also implies that contrast tuning is simply scaled by changes in orientation as illustrated by the projections along the contrast axis for different orientations.

Power law functions with nonzero threshold, i.e., $r(V) = k([V - T]^+)^n$ for $T > 0$, do not yield contrast-invariant responses, instead they typically lead to an “iceberg” effect—tuning widens with increasing contrast as more orientations receive suprathreshold input.

Accuracy of power-law approximation

Now we turn to the question of whether noise-smoothed threshold linear functions can be approximated by power-law nonlinearities. We let the instantaneous spike rate r be a threshold linear function of voltage V : $r_T(V) = k[V - T]^+$, where $T > 0$ is a threshold and again $V = 0$ represents the mean voltage at rest. We assume that the voltage can be written as a sum of a trial-averaged voltage \bar{V} and zero-mean Gaussian noise v with SD σ : $V = \bar{V} + v$, where (letting an overbar represent an average over trials) $\bar{v} = 0$ and $\overline{v^2} = \sigma^2$. One can then derive an equation for $\bar{r}_T(\bar{V})$, the average response (averaged over the stochastic noise) as a function of the trial-averaged voltage \bar{V} (see Fig. 2A; APPENDIX B). Experimental results usually report the stimulus-induced response with background response subtracted off. Thus we will study the quantity

$$\bar{R}_T(\bar{V}) = \bar{r}_T(\bar{V}) - \bar{r}_T(0) \quad (2)$$

We work in units of the noise, taking $\sigma \equiv 1$. In these units, both \bar{V} and T are measured as number of SDs of the noise above rest e.g., $T = 3$ means that threshold is 3 noise SDs above rest. We take the gain k to be 1, which simply sets the units of response. With these choices, the form of the function $\bar{R}_T(\bar{V})$ is determined by the single parameter T .

Figure 2B, top, shows \bar{R}_T versus \bar{V} as continuous lines for a range of values of threshold T . To determine how well this function might be approximated by a power law, for each T , we found the best-fit power law $k\bar{V}^n$ (least-mean-squares fit; shown as dashed lines) over the range $0 \leq \bar{V} \leq T + \bar{V}_{hi}$ (upper limit shown as vertical dotted line segments). We illustrate, and initially consider, the case $\bar{V}_{hi} = 1.5$. The power law gives an excellent fit to \bar{V} over the fitted range for all values of T . Fig. 2B, bottom, shows the same fits on a log-log plot. This shows

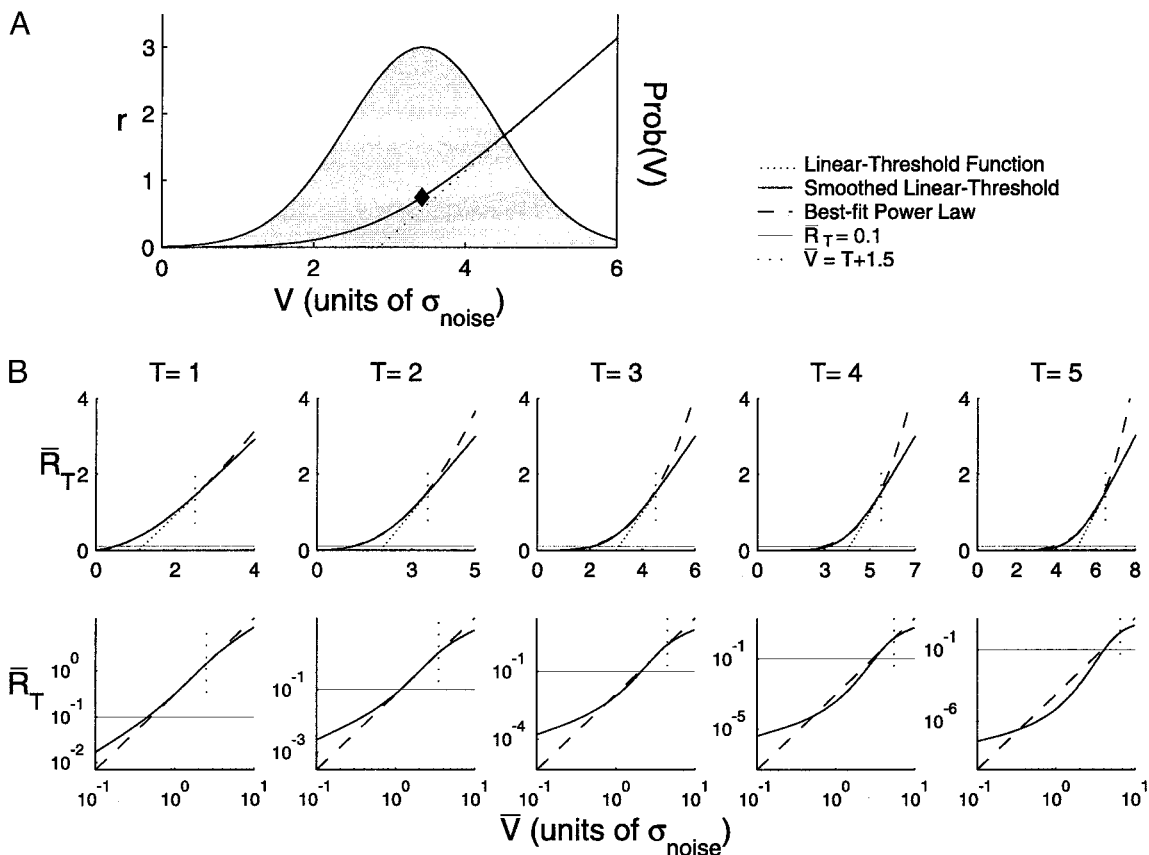


FIG. 2. A: how noise converts a linear threshold function (dashed line) into a smooth function (solid line). The response r is given by a linear threshold function of the voltage, $r(V) = k[V - T]^+$ (dotted line). Assuming $V = \bar{V} + v$, where v is 0-mean Gaussian noise (shown in gray), the average response \bar{r} vs. average voltage \bar{V} follows the smooth solid line. Diamond shows (\bar{V}, \bar{r}) for the Gaussian shown. Values illustrated were taken from those observed experimentally ($\sigma_{\text{noise}} = 3.5$ mV, $\bar{V} = 12$ mV = $3.43\sigma_{\text{noise}}$, $T = 10$ mV = $2.86\sigma_{\text{noise}}$; see discussion in text section *Comparison to experimental data*). B: the smooth function that arises from the combination of noise and a linear-threshold response function closely approximates a power law. Mean output \bar{R}_T (vertical axes) is shown vs. mean voltage \bar{V} (horizontal axes), on linear axes (top) or as log-log plots (bottom) for varying values of threshold T . Continuous curves: exact noise-smoothed function \bar{R}_T (Eqs. 2 and B1). Dashed lines: best-fitting power law $k\bar{V}^n$. Both \bar{V} and T are expressed as number of SDs of the noise above rest. Dotted vertical lines: $\bar{V} = T + 1.5$, the upper boundary of the region over which the power law was fit. Thin solid horizontal lines: $\bar{R}_T = 0.1$. Power law provides a good fit to \bar{R}_T wherever $\bar{V} \leq T + 1.5$ and $\bar{R}_T \geq 0.1$. x axes extend from 0 to $T + 3$ (top plots) or from 0.1 to 10 (bottom plots). Power law was found as fit giving least mean-square error over $0 \leq \bar{V} \leq T + 1.5$, using Matlab 6.0 function `fminsearch`.

that the power law fails for small inputs. However, these small inputs correspond to very low output rates ($\bar{R}_T = 0.1$ shown as thin horizontal lines) that occur only for values of \bar{V} well below threshold and result in negligible absolute differences between actual and fitted functions.

The best-fit exponent [n in the power law $\bar{R}_T \approx k([\bar{V}]^+)^n$] is always greater than one and increases with increasing threshold T (Fig. 3). Intuitively, a larger exponent means that the response \bar{V}^n remains small for larger values of \bar{V} , consistent with a higher threshold (cf. Carandini et al. 1997). Note that because the threshold T is expressed in units of the noise, increasing T for fixed noise is equivalent to decreasing the noise for fixed T . Thus the exponent is expected to be a decreasing function of noise level, i.e., higher noise leads to lower exponents (Anderson et al. 2000).

We quantified the robustness of the power law approximation to $\bar{R}_T(\bar{V})$ in two ways. First, for each value of the threshold T , we calculated the size of the error at a given voltage \bar{V} relative to the size of the response $\bar{R}_T(\bar{V})$, i.e., we plot $|\bar{R}_T(\bar{V}) - k\bar{V}^n|/\bar{R}_T(\bar{V})$ for a range of values of \bar{V} and T (Fig. 4). (Recall that

k and n are determined from optimizing the fit for $0 \leq \bar{V} \leq T + 1.5$). Not surprisingly, the power law breaks down for large \bar{V} . In this range, all values of $\bar{V} + v$ are above threshold, and the input-output function becomes the underlying linear threshold function. For low thresholds, the best fit power-law is more nearly linear (exponents near 1.0, see Fig. 3) and a good fit extends well beyond $T + 1.5$ ($T + 1.5$ is indicated by the upper dashed line in Fig. 4). For $T > 2$, however, accurate power-law fits do not extend much beyond this upper bound of the fitting range. The power law also breaks down for small values of \bar{V} . This indicates that the power-law fit does not capture the exact shape of the transfer function as it bends away from $\bar{R}_T = 0$ (see Fig. 2B). However, errors at small \bar{V} correspond to very low firing rates ($\bar{R}_T < 0.1$ shown as lower dashed line in Fig. 4). At such low rates, these large relative errors reflect small absolute differences between $\bar{R}_T(\bar{V})$ and the best-fit power law (see Fig. 2B, top).

To determine the range of voltages over which the power-law can give a good fit, we varied V_{hi} , the upper voltage cutoff of the range of fit $[0, T + V_{\text{hi}}]$. For each V_{hi} , we calculated the

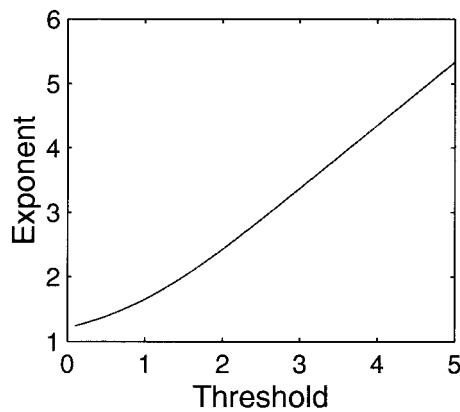


FIG. 3. Dependence of power-law exponent $n(T)$ on threshold T , for best-fit power law $k(T)\bar{V}^{n(T)}$. Fits were done as described in legend of Fig. 2.

average absolute error of the approximation (Fig. 5A), and the average error relative to the response (Fig. 5B), for integral values of threshold T from 1 to 5. As T increases beyond 2 SDs of the noise, good power-law fits are only obtained for ranges that extend about 1.5 SDs above threshold.

In summary, both methods of assessing the accuracy of power law fits reveal that, across a wide range of thresholds, a power law gives a good fit in the range $[0, T + 1.5]$.

Comparison to experimental data

We can compare these results to data as follows. The noise in the recordings of Anderson et al. (2000) was generally $\sigma = 3\text{--}4$ mV (rms). Their thresholds were roughly 10 mV from rest, yielding $T = 2.5\text{--}3.3$ (expressed in units of the noise). Stimulus-induced voltage changes (DC + F1) at the highest contrast studied at the preferred orientation were in the range of 8–12 mV, yielding $\bar{V} = 2\text{--}4$ (again in units of the noise) or no more than 1.5 above T . (Only contrasts $\leq 64\%$ were studied, but responses of most cat V1 cells are nearly or entirely

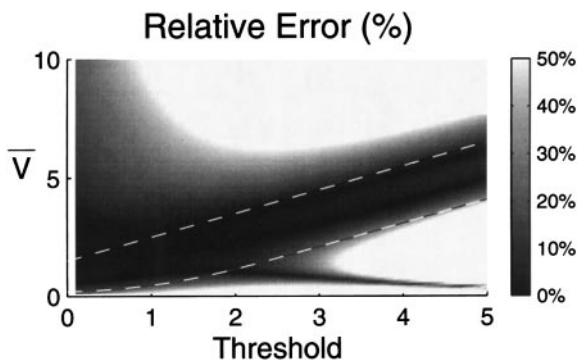


FIG. 4. Error of the best-fit power-law approximation $k\bar{V}^n$, expressed relative to the trial-averaged response \bar{R}_T . Relative error is plotted versus threshold T (horizontal axis) and input \bar{V} (vertical axis). Relative error is defined as $|\bar{R}_T(\bar{V}) - k\bar{V}^n| / \bar{R}_T(\bar{V})$ (expressed as a percentage). This error is indicated as grayscale, linear from black = 0 to white = 50%; all values $\geq 50\%$ are set to white. Upper dashed line indicates $\bar{V} = T + 1.5$, the upper boundary of the region over which the power law was fit. Lower dashed line indicates contour along which $\bar{R} = 0.1$. As in Fig. 2, $k(T)$ and $n(T)$ in power law $k\bar{V}^n$ were fit separately for each T over the range $0 \leq \bar{V} \leq T + 1.5$. Regions of large relative error below the $\bar{R} = 0.1$ contour show little absolute error: if absolute error $|\bar{R}_T(\bar{V}) - k\bar{V}^n|$ rather than relative error is plotted on the same scale (i.e., black = 0, white = 0.5), these regions become black (not shown; see also Fig. 2).

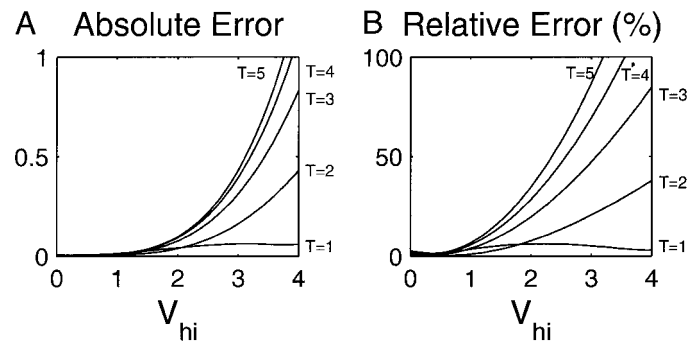


FIG. 5. Average absolute error (A) and average relative error (B) for variations of the upper bound V_{hi} of the range over which a power law was fit. Power law was fit over range $0 \leq \bar{V} \leq T + V_{hi}$. Error is plotted vs. V_{hi} . Average absolute error is average over $V \in [0, T + V_{hi}]$ of $|\bar{R}_T(\bar{V}) - k\bar{V}^n|$, while average relative error is average over the same range of $|\bar{R}_T(\bar{V}) - k\bar{V}^n| / \bar{R}_T(\bar{V})$. The five lines in each plot correspond to $T = 1, 2, 3, 4, 5$, as labeled. For $T > 1$, error rapidly grows for $V_{hi} > 1.5$; previous figures showed results for $V_{hi} = 1.5$.

saturated at that contrast, e.g., Albrecht 1995.) Thus they found that even the strongest visual cortical voltage responses remain in the range in which a linear threshold function yields a power law as the averaged input-output function.

A power law with exponent n is expected to sharpen tuning by a factor of \sqrt{n} . This is based on consideration of a Gaussian tuning curve: if such a curve has SD σ , raising it to the power n produces a curve with SD σ/\sqrt{n} . Thus we can compare independent estimates of the exponent n , one obtained from measures of sharpening of tuning, and the other obtained by using the relationship of noise levels to exponents shown in Fig. 3. The range $T = 2.5\text{--}3.3$ obtained from the recordings of Anderson et al. (2000) yields exponents $n = 2.9\text{--}3.7$ (Fig. 3), corresponding to a sharpening of tuning by factors of $\sqrt{n} = 1.7\text{--}1.9$. Carandini and Ferster (2000) found that voltage orientation tuning had a half-width-half-height (HWHH) of $38 \pm 15^\circ$ (mean \pm SD, averaged over cells), while spiking orientation tuning had a HWHH of about $23 \pm 8^\circ$ [see also Volgushev et al. (2000), who also found spike tuning to be sharper than tuning of intracellular potentials]; however, their spiking tuning estimate was almost certainly overly broad, because their methods did not allow resolution of spiking HWHHs $< 20^\circ$. These mean values represent a sharpening by a factor of 1.65, the square of which suggests an exponent $n = 2.72$, which is attained in our model when $T = 2.3$. Given that this is almost certainly an underestimate of the true sharpening, this agrees well with the estimate $n = 2.9\text{--}3.7$. Gardner et al. (1999) examined the same issue using extracellular recording by comparing the tuning predicted from a cell's noise-mapped linear receptive field to that observed in response to gratings; they found sharpening corresponding to power-law exponents that had a geometric mean across cells of 3.15. Under the assumption that the linear receptive field approximates the transformation of stimuli into membrane voltage, this degree of sharpening agrees well with the noise-based estimate.

DISCUSSION

Here we have shown two things. First, a power law is the only input-output function that converts contrast-invariant voltage tuning into contrast-invariant spiking tuning. Second, given a linear-threshold function relating instantaneous voltage

to instantaneous firing rate, addition of Gaussian noise to the voltage yields a relationship between trial-averaged voltage and trial-averaged firing rate that is well approximated by an expansive power law. This approximation seems quite good provided that the trial-averaged voltage does not exceed the threshold by more than about 1.5 SDs of the noise. We have gone on to compare these results to existing data. The voltage responses reported by Anderson et al. (2000) remained within the range in which a power law gives a good approximation, even at high contrasts. The exponents predicted from their reported noise and threshold levels predict a sharpening of spiking tuning, relative to voltage tuning, that agrees well with published data on the degree of such sharpening (Carandini and Ferster 2000; Gardner et al. 1999).

Mechanisms yielding contrast-invariant tuning

Anderson et al. (2000) used numerical simulations to demonstrate that neural noise and a threshold-linear transfer function could transform contrast-invariant voltage responses into contrast-invariant spike responses. Our results, both theoretical (APPENDIX A) and computational (Figs. 2 and 4), indicate that the invariance of spike tuning was due to the fact that a noise-smoothed threshold-linear function is well approximated by a power law. Thus the approach of Anderson et al. (2000) to contrast-invariant orientation tuning, based on noise-smoothed linear threshold models, resembles phenomenological descriptions in which a linear filtering of the stimulus is followed by an expansive power law nonlinearity (Carandini et al. 1997; Heeger 1992; Heeger et al. 1996). It would be interesting to see if similar mechanisms explain contrast-invariance of other response properties, such as spatial frequency tuning (Albrecht and Hamilton 1982).

However, a noise-induced power-law nonlinearity only explains half the problem of contrast-invariant tuning of V1 simple cells, namely how voltage tuning that scales with contrast is converted into spiking tuning that scales with contrast. The mechanisms by which the input to simple cells from neurons in the lateral geniculate nucleus (LGN) is converted into voltage tuning that scales with contrast remain to be elucidated. The voltage responses do not result from a simple linear filtering of the input as postulated by the phenomenological models. If the voltage resulted from a linear filtering of the drifting sinusoidal grating stimulus, then the mean voltage would be independent of orientation and contrast because changes in these parameters do not change mean luminance. The experiments of Anderson et al. (2000) found, instead, that the mean voltage responses to drifting sinusoidal grating stimuli were well-tuned for orientation and grew with contrast. From a more neural point of view, one might expect the voltage response of a simple cell to arise from a linear filtering of the firing rates of the LGN neurons that are presynaptic to that cell. Because LGN firing rates have a rectification nonlinearity—their firing rates can greatly increase but can decrease only to zero—their mean rate of firing increases with contrast. However, the output of such a filter should again be untuned for orientation because the mean responses of LGN cells are largely untuned for orientation, and the mean voltage response under a linear filtering of the LGN would be obtained by a weighted sum of the mean responses of each LGN input (Ferster and Miller 2000; Troyer et al. 1998). Thus the ob-

served voltage responses do not arise from a linear filtering either of the stimulus or of the LGN firing rates.

This raises the question as to how this mean LGN input, which is untuned for orientation and grows with contrast, is converted into a tuned mean voltage response. We have suggested that feedforward inhibition (inhibition from interneurons driven by LGN input) can suppress the untuned mean input from the LGN and hence explain the lack of contrast-dependent responses to stimuli oriented perpendicular to the preferred orientation (Ferster and Miller 2000; Miller et al. 2001; Troyer et al. 1998; see also McLaughlin et al. 2000; Wielaard et al. 2001, who propose a similar feedforward inhibitory mechanism but in the context of a somewhat different circuit). The increases in mean voltage for preferred stimuli may arise through the interaction between cellular or circuit nonlinearities and the large voltage modulations experienced at these orientations: both the reversal potential nonlinearity and intracortical excitation from other cells with a threshold nonlinearity will cause stimuli that yield larger voltage modulations to be accompanied by larger voltage mean responses. That is, the tuning of the voltage mean may largely be inherited from the tuning of the voltage modulation. More generally, the results of Anderson et al. (2000) point away from the nonlinearities involved in converting intracellular voltages to spikes (e.g., the threshold nonlinearity) as being the key issue related to contrast-invariant tuning, and toward an investigation of mechanisms that contribute to the tuning properties of the intracellular voltage.

Comparison to previous theoretical work

It is widely known that noise can smooth and in certain respects linearize a threshold nonlinearity (e.g., Knight 1972; Spekreijse 1969; Stemmler 1996), making otherwise subthreshold inputs become “visible.” However, the present work is showing something far more specific, namely that the specific smooth function that results closely approximates a power law. Furthermore it is key that it is the voltage deviations from background that are raised to a power [i.e., $\bar{R}_T(\bar{V}) = k([\bar{V}]^+)^n$]. Had the form of the output function instead been, say, $k([\bar{V} - T]^+)^n$ for $T > 0$, this would not yield contrast-invariant tuning—more and more of the input would be suprathreshold at higher contrasts. Because the exact equation for \bar{R}_T (Eqs. 2 and B1) depends on \bar{V} only through its dependence on $\bar{V} - T$, it is surprising that this is well approximated by $k\bar{V}^n$ over a significant range—we know of no simple analytic reason why this empirical finding should be true.

We are aware of two other works that relate threshold-linear functions and power laws. First, Carandini et al. (1997) showed graphically that a power law can roughly approximate a threshold-linear function with higher thresholds corresponding to larger exponents. This suggests that for some response properties, models based on linear threshold and power law nonlinearities may yield similar predictions. However, contrast-invariant orientation tuning requires that where responses are significantly larger than zero, the *ratio* of responses between various orientations must remain constant at different levels of contrast. Even though absolute differences between a power law and an unsmoothed linear threshold function might be moderate, the relative error between the functions is very large near threshold, and as a result linear threshold models do not

yield contrast-invariant tuning (Anderson et al. 2000). The key point we are making is not simply that a linear threshold function resembles a power law but rather that noise *converts* a linear threshold function into a different function that approximates a power law sufficiently closely to achieve contrast-invariant tuning. Second, in a finding close in spirit to the present work, Suarez and Koch (1989) showed that, given a linear threshold model, adding noise to the input that is uniformly distributed over some range (or, having a population of cells receiving identical input but with a uniform distribution of thresholds) acts like taking the integral of the linear threshold function, yielding a quadratic input-output function. However, the argument is not robust. It only yields a power law of the form $k([\bar{V}]^+)^n$ [rather than $k([\bar{V} - T]^+)^n$] when the upper bound of the noise (at rest) is equal to spike threshold and can only yield power law exponents exactly equal to 2.

Conclusions

Neural noise can convert an instantaneous linear-threshold input-output function into a power-law relationship between mean input and mean output. Given reports suggesting that spontaneous voltage fluctuations (“noise”) in neocortex in vivo are large and of comparable size to stimulus-induced voltage modulations (Arieli et al. 1996; Azouz and Gray 1999; Ho and Destexhe 2000; Paré et al. 1998; Tsodyks et al. 1999), it will be of great interest to determine if the response properties of cells in other regions of the neocortex are best modeled by an expansive power law nonlinearity. In particular, it will be interesting to see if such a power law might be related more generally to tuning for stimulus form that is invariant to changes in stimulus magnitude.

APPENDIX A

Here we show that a rectified power law $\bar{R}(\bar{V}) = k([\bar{V}]^+)^n$ is the only static input-output function that converts contrast-invariant input tuning into contrast-invariant output tuning, assuming that $\bar{R}(\bar{V})$ is nondecreasing (an increase in voltage cannot give a decrease in response) and nonnegative. The more general expression without the latter assumptions is also a form of power law.

We let c be contrast, θ be the other parameters (such as orientation) that show contrast-invariant tuning, \bar{V} be the input and $\bar{R}(\bar{V})$ be the output. Contrast-invariant input tuning implies that

$$\bar{V}(c, \theta) = f(c)g(\theta) \tag{A1}$$

for some continuous functions f, g , and we assume contrast scaling is nonnegative: $f(c) \geq 0$. Contrast-invariant output tuning implies that

$$\bar{R}(f(c)g(\theta)) = F(f(c))G(g(\theta)) \tag{A2}$$

for some continuous functions F, G , and again nonnegativity of contrast scaling implies $F \geq 0$. We assume that $F(x)$ and $G(x)$ are differentiable, at least for $x \neq 0$.

Differentiating both sides of Eq. A2 with respect to f and g yields

$$\frac{d\bar{R}(fg)}{df} = g\bar{R}'(V) = F'(f)G(g) = \frac{d(F(f)G(g))}{df} \tag{A3}$$

$$\frac{d\bar{R}(fg)}{dg} = f\bar{R}'(V) = F(f)G'(g) = \frac{d(F(f)G(g))}{dg} \tag{A4}$$

We begin by assuming that $f, F(f), g$, and $G(g)$ are all nonzero. Then

$$\bar{R}'(V) = \frac{F'(f)G(g)}{g} = \frac{F(f)G'(g)}{f} \tag{A5}$$

which yields

$$\frac{F'(f)f}{F(f)} = \frac{G'(g)g}{G(g)} = n \tag{A6}$$

Here, n is an arbitrary constant; because the quantity on the left is a function only of c , and that in the middle is a function only of θ , the only way these two quantities can be equal to one another is if they are both equal to a constant (something that depends neither on c or θ), which we call n . Focusing on the function G , we obtain

$$\frac{G'(g)}{G(g)} = \frac{n}{g} \tag{A7}$$

Integrating both sides with respect to g yields

$$\int dg \frac{G'(g)}{G(g)} = \int dg \frac{n}{g} \tag{A8}$$

$$\ln G(g) = n \ln g + k_1 \tag{A9}$$

where k_1 is a constant of integration. Exponentiating yields

$$G(g) = k_2 g^n \tag{A10}$$

where $k_2 = e^{k_1}$. Identical reasoning shows that $F(f) = k_3 f^n$ for some constant k_3 . Therefore

$$\bar{R}(\bar{V}) = \bar{R}(fg) = F(f)G(g) = k_2 k_3 f^n g^n = k\bar{V}^n \tag{A11}$$

where $k = k_2 k_3$.

We obtained these results on the assumption that f, F, g , and G were all nonzero. Combining these results with the continuity of G and F , however, we can conclude that 1) $G(0) = 0$ and $F(0) = 0$ and 2) F and G are either strictly zero or strictly nonzero (and equal to a power law) on each open half-infinite interval $(-\infty, 0)$ and $(0, \infty)$. Finally, since we have assumed that f is nonnegative and $\bar{R}(\bar{V})$ is nondecreasing and nonnegative, $G(g) = 0$ for $g < 0$. Thus, over the full range of \bar{V}

$$\bar{R}(\bar{V}) = k([\bar{V}]^+)^n \tag{A12}$$

APPENDIX B

Let the voltage $V = \bar{V} + v$ where overbar represents an average, $\bar{v} = 0$ and $\bar{v}^2 = \sigma^2$. We assume the instantaneous spike rate is given by $r_T(V) = k[V - T]^+$. Then the trial-averaged spike rate is

$$\begin{aligned} \bar{R}_T(\bar{V}) &= k[\bar{V} + v - T]^+ \\ &= \frac{k}{\sqrt{2\pi\sigma^2}} \int_{T-\bar{V}}^{\infty} dv \exp\left(-\frac{v^2}{2\sigma^2}\right) (\bar{V} - T + v) \\ &= k\sigma \left[\frac{\bar{V} - T}{2\sigma} \left(1 + \operatorname{erf}\left(\frac{\bar{V} - T}{\sqrt{2}\sigma}\right) \right) + \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(\bar{V} - T)^2}{2\sigma^2}\right) \right] \end{aligned} \tag{B1}$$

where erf is the error function, $\operatorname{erf}(x) = (2/\sqrt{\pi}) \int_0^x dy \exp(-y^2)$. To work in units of the noise, we set $\sigma = 1$ in Eq. B1.

We thank B. Bialek, J. Anderson, and D. Ferster for useful conversations. We gratefully acknowledge support by National Eye Institute Grant EY-11001 (K. D. Miller).

After this work was completed, we became aware that the independent work of D. Hansel and C. van Vreeswijk reached similar conclusions.

REFERENCES

ALBRECHT DG. Visual cortex neurons in monkey and cat: effect of contrast on the spatial and temporal phase transfer functions. *Vis Neurosci* 12: 1191–1210, 1995.
 ALBRECHT DG AND GEISLER WS. Motion selectivity and the contrast-response function of simple cells in the visual cortex. *Vis Neurosci* 7: 531–546, 1991.

- ALBRECHT DG AND HAMILTON DB. Striate cortex of monkey and cat: contrast response function. *J Neurophysiol* 48: 217–237, 1982.
- ANDERSON JS, LAMPL I, GILLESPIE D, AND FERSTER D. The contribution of noise to contrast invariance of orientation tuning in cat visual cortex. *Science* 290: 1968–1972, 2000.
- ANZAI A, OHZAWA I, AND FREEMAN RD. Neural mechanisms for processing binocular information. I. Simple cells. *J Neurophysiol* 82: 891–908, 1999.
- ARIELI A, STERKIN A, GRINVALD A, AND AERTSEN A. Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* 273: 1868–1871, 1996.
- AZOUZ R AND GRAY CM. Cellular mechanisms contributing to response variability of cortical neurons. *J Neurosci* 19: 2209–2223, 1999.
- CARANDINI M AND FERSTER D. Membrane potential and firing rate in cat primary visual cortex. *J Neurosci* 20: 470–484, 2000.
- CARANDINI M, HEEGER DJ, AND MOVSHON JA. Linearity and normalization in simple cells of the macaque primary visual cortex. *J Neurosci* 17: 8621–8644, 1997.
- CARANDINI M, HEEGER DJ, AND MOVSHON JA. Linearity and gain control in V1 simple cells. In: *Cerebral Cortex. Models of Cortical Circuits*, edited by Uliński PS, Jones EG, and Peters A. New York: Kluwer Academic/Plenum, 1999, vol. 13, p. 401–443.
- EMERSON RC, KORENBERG MJ, AND CITRON MC. Identification of intensive nonlinearities in cascade models of visual cortex and its relation to cell classification. In: *Advanced Methods of Physiological System Modeling*, edited by Marmarelis VZ. New York: Plenum, 1989, p. 97–111.
- FERSTER D AND MILLER D. Neural mechanisms of orientation selectivity in the visual cortex. *Annu Rev Neurosci* 23: 441–471, 2000.
- GARDNER JL, ANZAI A, OHZAWA I, AND FREEMAN RD. Linear and nonlinear contributions to orientation tuning of simple cells in the cat's striate cortex. *Vis Neurosci* 16: 1115–1121, 1999.
- HEEGER DJ. Half-squaring in responses of cat striate cells. *Vis Neurosci* 9: 427–443, 1992.
- HEEGER DJ. Modeling simple-cell direction selectivity with normalized, half-squared, linear operators. *J Neurophysiol* 70: 1885–1898, 1993.
- HEEGER DJ, SIMONCELLI EP, AND MOVSHON JA. Computational models of cortical visual processing. *Proc Natl Acad Sci USA* 93: 623–627, 1996.
- HO N AND DESTEXHE A. Synaptic background activity enhances the responsiveness of neocortical pyramidal neurons. *J Neurophysiol* 84: 1488–1496, 2000.
- JAGADEESH B, WHEAT HS, AND FERSTER D. Linearity of summation of synaptic potentials underlying direction selectivity in simple cells of the cat visual cortex. *Science* 262: 1901–1904, 1993.
- JAGADEESH B, WHEAT HS, KONTSEVICH LL, TYLER CW, AND FERSTER D. Direction selectivity of synaptic potentials in simple cells of the cat visual cortex. *J Neurophysiol* 78: 2772–2789, 1997.
- KNIGHT BW. Dynamics of encoding in a population of neurons. *J Gen Physiol* 59: 734–766, 1972.
- LAMPL I, ANDERSON J, GILLESPIE D, AND FERSTER D. Prediction of orientation selectivity from receptive field architecture in simple cells of cat visual cortex. *Neuron* 30: 263–274, 2001.
- MCLAUGHLIN D, SHAPLEY R, SHELLEY M, AND WIELAARD DJ. A neuronal network model of macaque primary visual cortex (V1): orientation selectivity and dynamics in the input layer 4 α . *Proc Natl Acad Sci USA* 97: 8087–8092, 2000.
- MILLER KD, SIMONS DJ, AND PINTO DJ. Processing in layer 4 of the neocortical circuit: new insights from visual and somatosensory cortex. *Curr Opin Neurobiol* 11: 488–497, 2001.
- MURTHY A, HUMPHREY AL, SAUL AB, AND FEIDLER JC. Laminar differences in the spatiotemporal structure of simple cell receptive fields in cat area 17. *Vis Neurosci* 15: 239–256, 1998.
- PARÉ D, SHINK E, GAUDREAU H, DESTEXHE A, AND LANG EJ. Impact of spontaneous synaptic activity on the resting properties of cat neocortical pyramidal neurons in vivo. *J Neurophysiol* 79: 1450–1460, 1998.
- SCLAR G AND FREEMAN RD. Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. *Exp Brain Res* 46: 457–461, 1982.
- SCLAR G, MAUNSELL JH, AND LENNIE P. Coding of image contrast in central visual pathways of the macaque monkey. *Vision Res* 30: 1–10, 1990.
- SKOTTUN BC, BRADLEY A, SCLAR G, OHZAWA I, AND FREEMAN RD. The effects of contrast on visual orientation and spatial frequency discrimination: a comparison of single cells and behavior. *J Neurophysiol* 57: 773–786, 1987.
- SPEKREIJSE H. Rectification in the goldfish retina: analysis by sinusoidal and auxiliary stimulation. *Vision Res* 9: 1461–1472, 1969.
- STEMMLER M. A single spike suffices: the simplest form of stochastic resonance in model neurons. *Network* 7: 687–716, 1996.
- SUAREZ H AND KOCH C. Linking linear threshold units with quadratic models of motion perception. *Neural Comput* 1: 318–320, 1989.
- TROYER TW, KRUKOWSKI A, PRIEBE NJ, AND MILLER KD. Contrast-invariant orientation tuning in cat visual cortex: feedforward tuning and correlation-based intracortical connectivity. *J Neurosci* 18: 5908–5927, 1998.
- TSODYKS MV, KENET T, GRINVALD A, AND ARIELI A. Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science* 286: 1943–1946, 1999.
- VOLGUSHEV M, PERNBERG J, AND EYSEL UT. Comparison of the selectivity of postsynaptic potentials and spike responses in cat visual cortex. *Eur J Neurosci* 12: 257–263, 2000.
- WIELAARD DJ, SHELLEY M, MCLAUGHLIN D, AND SHAPLEY R. How simple cells are made in a nonlinear network model of the visual cortex. *J Neurosci* 21: 5203–5211, 2001.