

Blue-Yellow Signals Are Enhanced by Spatiotemporal Luminance Contrast in Macaque V1

Gregory D. Horwitz,^{1,2} E. J. Chichilnisky,² and Thomas D. Albright^{1,2}

¹Howard Hughes Medical Institute, ²Systems Neurobiology Laboratories, The Salk Institute for Biological Studies, La Jolla, California

Submitted 21 July 2004; accepted in final form 16 October 2004

Horwitz, Gregory, E. J. Chichilnisky, and Thomas D. Albright. Blue-yellow signals are enhanced by spatiotemporal luminance contrast in macaque V1. *J Neurophysiol* 93: 2263–2278, 2005; doi:10.1152/jn.00743.2004. We measured the color tuning of a population of S-cone-driven V1 neurons in awake, fixating monkeys. Analysis of randomly chosen color stimuli that were effective in evoking action potentials showed that these neurons received opposite sign input from the S cones and a combination of L and M cones. Surprisingly, these cells also responded to LM cone contrast irrespective of polarity, a nonlinear sensitivity that was masked by conventional linear analysis methods. Taken together, these observations can be summarized in a nonlinear model that combines nonopponent and opponent signals such that luminance contrast enhances color processing. These findings indicate that important aspects of the cortical representation of color cannot be described by classical linear analysis, and reveal a possible neural correlate of perceptual color-luminance interactions.

INTRODUCTION

Color perception results from a complex neuronal computation. The early steps of this computation are well understood: the sensitivity of the cone photoreceptors to lights of various wavelength has been characterized thoroughly (Baylor et al. 1987; Wandell 1995) as has the subsequent synergistic and antagonistic combination of cone signals in the retina and lateral geniculate nucleus (Derrington et al. 1984; DeValois 1965; Gouras 1968). How color information is processed in cortex, however, is less clear. The goal of the current experiments was to extend our understanding of how neurons in the primary visual cortex (V1) combine signals that originate in the cones.

One possibility is that V1 neurons combine cone signals linearly. This assumption has been made, either implicitly or explicitly, in studies that employ cone-isolating stimuli to estimate the weights with which cone inputs are integrated (Conway 2001; Johnson et al. 2001; Landisman and Ts'o 2002). While the linearity assumption is justified for many V1 neurons, it is clearly inappropriate for others (Conway 2001; Hanazawa et al. 2000; Hubel and Wiesel 1968; Lennie et al. 1990; Vautin and Dow 1985). For example, a V1 neuron that responds to S cone stimulation, but only when L- and M-cone excitations are appropriately balanced, does not integrate cone inputs linearly and thus cannot be characterized with cone-isolating stimuli (Hanazawa et al. 2000).

Recently developed data-analysis tools have provided new ways to characterize neurons that combine inputs nonlinearly (de Ruyter van Steveninck and Bialek 1988; Paninski 2003;

Rust et al. 2004; Schwartz et al. 2001; Sharpee et al. 2004; Simoncelli et al. 2004; Touryan et al. 2002). Here we use one of these techniques to reveal a surprising nonlinear computation performed by blue-yellow neurons in V1.

We excited V1 neurons in awake monkeys with a dynamic, randomly colored stimulus and analyzed the stimulus sequences that preceded spikes. Analysis proceeded in two steps. First, we computed the average stimulus that preceded a spike. This analysis identified a group of S-cone-dominated, color-opponent neurons. Had these neurons combined cone signals linearly, this analysis would have characterized their color tuning completely. The second step in our analysis, however, demonstrated that approximately half of these neurons were fundamentally nonlinear. An analysis of the *variability* of the stimuli that preceded spikes revealed that these cells received a rectified nonopponent signal from L and M cones, which combined synergistically with the opponent signal. This interaction between opponent and nonopponent signals emphasizes color information in regions of a visual scene with high luminance contrast and may be related to enhancements of color vision that occur with the addition of luminance contrast.

METHODS

Subjects

Four alert rhesus monkeys (*Macaca mulatta*), weighing between 8 and 10 kg, served as subjects in these experiments. Animals had normal color vision and no significant refractive error. Experimental protocols were approved by the Salk Institute Animal Care and Use Committee and conform to U.S. Department of Agriculture regulations and to the National Institutes of Health guidelines for the humane care and use of laboratory animals.

Animal preparation

Procedures for surgical preparation, behavioral training, and electrophysiological recording were routine and similar to those described previously (Dobkins and Albright 1994). Briefly, each monkey was implanted with a stainless steel head post and monocular scleral search coil. During neural recording, head movements were prevented by securing the head post to a mating piece on the monkeys' chair. Three monkeys had craniotomies over the occipital operculum from which single units were isolated. The fourth monkey had a craniotomy that allowed access to the thalamic lateral geniculate nucleus (LGN).

Visual stimuli

Stimuli were generated using an 8-bit graphics card (ELSA GLoria Synergy) controlled by PC software (Cortex 5.94, NIMH), and were

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 and other correspondence: G. D. Horwitz, Vision Center Laboratory, The Salk Institute for Biological Studies, 10010 N. Torrey Pines Rd., La Jolla, CA 92037 (E-mail: horwitz@salk.edu).

displayed on a 19-in CRT monitor refreshed at 100 Hz (Sony F500). Phosphor emission spectra were measured in 4-nm bins using a spectroradiometer (PR-650, Photoresearch). The intensity of each phosphor was measured at 35 levels spanning its dynamic range, and all of our stimuli were constructed to compensate for the nonlinear voltage-intensity relationship. Monitor calibration routines were adapted from those included in the Matlab Psychophysics Toolbox (Brainard 1997). We verified that spectra produced by mixtures of the three phosphors were well predicted by linear combinations of the three component spectra, and we recalibrated the monitor periodically to compensate for slow changes in CRT characteristics that occurred over the course of data collection.

Monkeys were trained to maintain visual fixation while a randomly flickering 8×8 stimulus grid, shown in Fig. 1, was presented at the receptive field (RF) of one or more individually isolated V1 neurons. The color of each $0.22 \times 0.22^\circ$ square in the grid changed randomly and independently every 10 ms. Color changes were implemented by independently modulating the intensities of the phosphors according to Gaussian distributions. The space-time averaged intensity of each phosphor was equal to its contribution to the background, which was metameric with an equal-energy white at 65 cd/m^2 . The SD of each distribution was $\sim 9\%$ of the range physically achievable, corresponding to luminance contrasts of 3.9, 11.5, and 2.5% for the red, green, and blue phosphors, respectively.

We also stimulated some V1 neurons and all LGN neurons with a 3° diam uniformly colored disk centered on the RF. The color of this disk modulated with statistics identical to a single pixel in the 8×8 stimulus grid.

The SD of our stimulus in the L-, M-, and S-cone excitation directions was 0.012, 0.012, and 0.013, respectively, which yields 0.11, 0.13, and 0.26 units of cone contrast when each is divided by the corresponding mean cone excitation due to the background. These contrasts are comparable to those used in previous studies (Conway 2001; Johnson et al. 2001, 2004). Unlike stimuli used in these previous studies, however, our stimulus largely modulated the three cone types together. Thus cone contrasts in cone *isolating* directions (e.g., L cone direction *conditional* on M and S cone activations constrained to be at their mean levels) were smaller: (L: 0.03, M: 0.03, S: 0.22). These L and M cone contrasts are 3–10 times smaller than those used by Johnson et al. (2001, 2004) and Conway (2001), but the S-cone contrast is comparable.

We were concerned that eye movements made during the inter-trial interval would cause the RFs of recorded neurons to move off the monitor, thereby causing a profound drop in the light level at the RF and changing the adaptation state. To mitigate such changes, monkeys viewed the stimulus through an aluminum-foil-lined tunnel. The tunnel was rectangular in cross-section (of the same dimensions as the

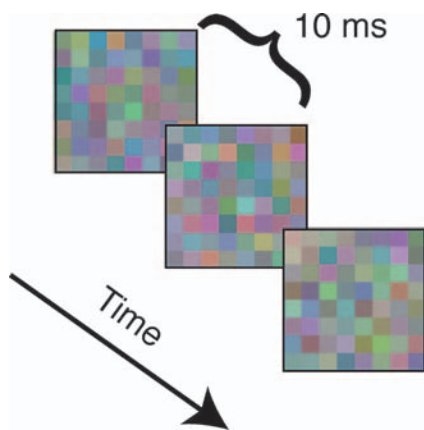


FIG. 1. Representative stimulus frames. Each $0.22 \times 0.22^\circ$ pixel in our stimulus changed color randomly and independently every 10 ms. Color changes were implemented by modulating the intensities of the red, green, and blue monitor phosphors according to independent Gaussian distributions.

CRT screen), spanned the distance from the CRT to the monkey's chair, and was closed on all four sides except for an opening in the top through which the experimenter monitored the monkey via a closed-circuit video camera. Stimulus presentations lasted until the monkey broke fixation (80% of trials) or until 10 s had elapsed (20% of trials). The mean stimulus duration was 4.97 ± 3.39 (SD) s. The mean number of trials per cell was 212.

Electrophysiological recording

For V1 recordings, glass-coated platinum-iridium electrodes (Frederick Haer, Brunswick, ME) were lowered through the dura mater via a hydraulic microdrive. For LGN recordings, parylene-coated tungsten electrodes (Frederick Haer) were inserted into transdural guide tubes and lowered by hydraulic microdrive. Electrical signals were amplified and filtered, and action potentials were isolated via template matching algorithms either on-line (Alpha Omega Engineering, Nazareth, Israel) or off-line (Plexon, Dallas, TX). Electrode impedances were 1–4 M Ω at 1 kHz.

On isolating a neuron, we mapped RF boundaries with oriented bars and spots of light. We then positioned the 8×8 stimulus grid on the estimated center of the RF and collected quantitative data. In early recording sessions, we recorded from each stable, well-isolated neuron we encountered that responded well to the stimulus. Later, we hunted specifically for blue-yellow neurons.

Data analysis

Analyses were performed using custom software written in Matlab (MathWorks). Stimulus movies were reconstructed and represented numerically as phosphor intensities at each frame and pixel in the display. Recorded spike trains were temporally aligned with the stimulus reconstruction so that each spike served as a pointer to the frame on which it occurred. We extracted the 20 frames preceding each spike to obtain a collection of stimulus sequences or "spike-triggered stimuli." We averaged spike-triggered stimuli together to compute linear RFs. Each linear RF had 3,840 components ($64 \text{ pixels} \times 3 \text{ phosphors} \times 20 \text{ frames}$). To calculate nonlinear RFs, we used a principal components analysis, as described in the following text.

Color spaces

Colored stimuli can be described in terms of the intensities of the primary lights used to create the stimuli, or equivalently, as the photon absorptions in the L, M, and S cones. Because the distribution of our stimuli was rotationally symmetric in the space of monitor primaries (phosphors), we used this space to calculate linear and nonlinear RFs (Chichilnisky 2001). We then transformed linear and nonlinear RFs to a two-dimensional color space with axes representing S cone weight and $3L+2M$ cone weight ("LM") (coefficients on L and M were derived from a least-squares fit to the human photopic luminosity function, V_λ). Only two color directions are reported here because the third is poorly constrained by the data (see following text).

R, G, B triplets were transformed to S, LM excitations by multiplication with a 2×3 matrix, which was derived from the phosphor emission spectra of our monitor and a set of cone fundamentals (Stockman et al. 1993). To construct this matrix, we computed the pairwise dot products of the emission spectra of the three phosphors (R, G, and B) and the two fundamentals used in this study: S and LM. This matrix can be represented as

$$A = \begin{bmatrix} R \cdot LM & G \cdot LM & B \cdot LM \\ R \cdot S & G \cdot S & B \cdot S \end{bmatrix} = \begin{bmatrix} 0.0533 & 0.1658 & 0.0453 \\ 0.0018 & 0.0089 & 0.1451 \end{bmatrix}$$

This matrix transforms lights expressed in R, G, B space to excitations in S, LM space. The R, G, B values derived from linear and nonlinear RFs do not represent *lights*, however, but visual *mechanisms* (Kno-

blanch and D'Zmura 2001). A light that is orthogonal to a mechanism in one space must be orthogonal to that mechanism in every space (e.g., isoluminant lights are orthogonal to the luminance mechanism irrespective of the color space in which they are represented). This would not be true in general if lights and mechanisms were transformed identically. The appropriate transformation for mechanisms is the inverse transpose of the transformation matrix for lights: A^{-T} (Knoblauch and D'Zmura 2001).

The color direction (3M-2L) that is orthogonal to the two used in this report (S, and 3L+2M) was largely absent from our display: stimulus modulation was 7.0 times greater in the 3L+2M direction, and 5.7 times greater in the S direction, than in the 3M-2L direction. Tuning estimates in this direction were thus poorly constrained by our data. The main conclusions of this study: that many blue-yellow neurons had significant nonlinear RFs, that these nonlinear RFs revealed sensitivity to LM modulation, and that this signal combined roughly multiplicatively with an S-cone-dominated cone-opponent signal did not change when this color direction was included in the analysis.

Principal-components analysis

The collection of spike-triggered stimuli contains a great deal of information about the stimulus patterns that caused a recorded neuron to fire. The linear RF (the average of these stimuli), which is often the only aspect of these data reported in scientific papers, is a gross distillation of the raw data and may conceal important stimulus preferences. Just as a probability distribution is not fully characterized by the mean, the distribution of spike-triggered stimuli is not fully characterized by the linear RF. Here we consider another aspect of the distribution of spike-triggered stimuli: its longest axis, or first principal component.

Principal-components analysis (PCA) allows high-dimensional data (such as the collection of spike-triggered stimuli) to be represented in a lower dimensional subspace in such a way that much of the structure in the original data is preserved. Principal components are ordered: the first principal component provides the best one-dimensional description of the data after the average has been subtracted. The second principal component provides the best one-dimensional description of the data once the contribution of the first has been taken into account, and so on.

Intuitively, PCA on the collection of spike-triggered stimuli can reveal sensitivity to complementary stimulus patterns that mask each other when averaged. For example, an ON cell responds to increments in light intensity, so the collection of spike-triggered stimuli tend to be brighter than the background. The average of these stimuli, the linear RF, is bright and reveals the stimulus preference. In contrast, an ON-OFF cell responds to both increments and decrements in light intensity at the same spatial location so both bright and dark stimuli appear in the collection of spike-triggered stimuli. These stimuli cancel when averaged, leading to an uninformative linear RF. The first principal component of the spike-triggered stimuli, the nonlinear RF, can reveal the stimulus preferences of such a cell.

Previous studies have shown that PCA can reveal stimulus features that excite complex cells (Rust et al. 2004; Touryan et al. 2002). To validate our implementation of the technique, we used PCA to reveal the orientation tuning of a complex cell. Figure 2A shows the orientation tuning curve of the cell measured with an achromatic grating. In a separate block of trials, we stimulated this neuron with the random noise stimulus (Fig. 1). As shown in Fig. 2B, the nonlinear RF accurately predicted the preferred orientation, but the linear RF did not.

Formally, the principal components are the eigenvectors of the spike-triggered covariance matrix. The spike-triggered covariance matrix can be written

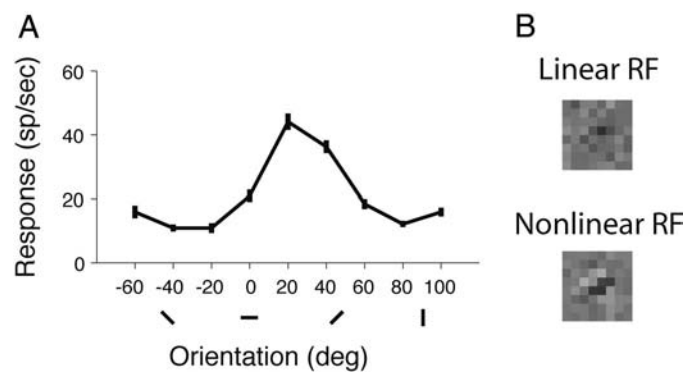


FIG. 2. Responses of a V1 complex cell. *A*: response as a function of the orientation of a 2 cycle/° achromatic grating drifting at 5 Hz. *B*: linear and nonlinear receptive fields (RFs) of the same neuron measured in a separate block of trials. Although the linear RF was largely unstructured, the nonlinear RF was structured and suggested a preferred orientation of ~22.5° (halfway between 0 and 45°), in close agreement with the preferred orientation measured with drifting gratings.

$$C = \frac{1}{n-1} \sum_i^n (s_i - a) \cdot (s_i - a)^T \quad (1)$$

where n is the number of spikes in an experiment, s_i is the i th spike-triggered stimulus, (represented as a vector whose entries indicate the intensity of each phosphor at each location and each time relative to the occurrence of a spike), and a is the spike-triggered average (that is, $\frac{1}{n} \sum_i^n s_i$). The first principal component is the eigen-

vector of this matrix associated with the largest eigenvalue, the second principal component is the eigenvector associated with the second largest eigenvalue, and so on (Strang 1988). It can be shown that the first principal component is the direction in stimulus space along which the spike-triggered stimuli have the greatest variance, the second is the direction along which the spike-triggered stimuli have the greatest variance in the subspace orthogonal to the first, and so on.

By construction, the raw stimuli have equal variance in every direction. If stimuli in certain regions of the stimulus space increase or decrease the firing probability, the collection of spike-triggered stimuli will not have equal variance in every direction. In this case, directions along which the collection of spike-triggered stimuli have particularly high or low variance can reveal the classes of stimuli that drive the cell. In what follows, we examine the stimulus direction that exhibited the highest variance, the first principal component.

How large is the first principal component relative to the spike-triggered average? A natural measure of the former is the square root of the largest eigenvalue of the spike-triggered covariance matrix. A natural measure of the latter is its vector norm. Comparison between these two measures is not straightforward, and a rigorous treatment of the issue is beyond the scope of this paper. Fortunately, these magnitudes are irrelevant under the class of models that we consider here, as will be described in the following text in *Models* (see also Chichilnisky 2001).

Projecting out the spike-triggered average

In an experiment such as ours, the first principal component of the spike-triggered stimuli can be redundant with the spike-triggered average (Schwartz et al. 2001). We were interested specifically in neurons for which this was not the case. We therefore subtracted from each spike-triggered stimulus whatever projection it had onto the linear RF (see following text), prior to performing PCA. All spike-triggered stimuli were therefore orthogonal to the linear RF and would therefore be invisible to a linear cell.

Formally, the projection, p , of a stimulus, s , onto the linear RF, a , is defined as

$$p = \frac{s^T \cdot a}{|a|^2} a \quad (2)$$

The projection of the stimulus *orthogonal* to the linear RF can be written

$$p_{\perp} = s - p \quad (3)$$

The covariance matrix of these projections, C_{ortho} , can be found by substituting p_{\perp} for $(s_i - a)$ in Eq. 1. An equivalent but faster approach is to calculate the spike-triggered covariance matrix by Eq. 1 and perform the equivalent matrix projection prior to the eigenvector decomposition

$$P = I - a(a^T \cdot a)^{-1} a^T \quad (4)$$

$$C_{\text{ortho}} = P \cdot C \cdot P^T \quad (5)$$

The matrix P in Eq. 4 is a projection matrix into the subspace orthogonal to a (Draper and Smith 1998). The transformation of the covariance matrix in Eq. 5 is equivalent to computing the covariance of the spike-triggered stimuli, after having multiplied each stimulus vector by the matrix, P (Leon-Garcia 1994).

Forcing orthogonality between linear and nonlinear RFs guarantees that they will not be redundant, so we impose this constraint throughout this report. For the 27 blue-yellow V1 neurons that are the focus of this report, however, this constraint was inconsequential. Nonlinear RFs (computed from the raw spike-triggered stimuli) were nearly orthogonal to linear RFs (median correlation coefficient = 0.09), even without this projection. As a result, nonlinear RFs computed from the raw spike-triggered stimuli were nearly identical to nonlinear RFs computed from stimulus projections orthogonal to the linear RF (median correlation coefficient = 0.99).

Significance testing

Statistical significance of the nonlinear RF was assessed with a nonparametric randomization test on the largest eigenvalue of the spike-triggered covariance matrix. To estimate the distribution of this statistic under the null hypothesis of no relationship between stimuli and spikes, we randomly shifted spike trains in time relative to the reconstructed stimulus movie, recalculated the spike-triggered covariance matrix, and obtained the largest eigenvalue. This procedure was repeated 2000 times. If the largest eigenvalue from the unrandomized data exceeded 95% of the largest eigenvalues from the randomized data sets, we concluded that the nonlinear RF was significant at the 0.05 level. Because we project the spike-triggered stimuli orthogonal to the linear RF prior to performing PCA, linear cells should not have significant nonlinear RFs, irrespective of any static nonlinearity that follows linear integration of cone inputs.

Pixel and frame selection

Many pixels in our stimulus fell outside the spatial extent of a given cell's RF, and many of the frames we analyzed lay outside a given cell's temporal integration window. These "nuisance dimensions" add appreciable noise to the nonlinear RFs. To focus our analyses on the pixels and frames that drove each cell most strongly, we used the following procedure. For each cell, in addition to calculating the spike-triggered average, we calculated the spike-triggered variance. Like the average, the variance has 3,840 components (64 pixels \times 3 phosphors \times 20 frames): it is the variance of the spike-triggered stimuli, or equivalently, the elements on the main diagonal of the spike-triggered covariance matrix. We converted means and variances to z scores by subtracting each from their theoretical means and then dividing by their theoretical SDs, which were calculated under the null

hypothesis of independence between the stimulus and Poisson-generated spikes. The formulas for these theoretical means and variances were

$$\mu_{\text{mean}} = \mu$$

$$\sigma_{\text{mean}} = \frac{\sigma}{\sqrt{n}}$$

$$\mu_{\text{var}} = \sigma^2$$

$$\sigma_{\text{var}} = \sigma^2 \sqrt{\frac{2}{n-1}}$$

Where μ is the mean phosphor intensity, σ is the SD of the phosphor intensities, and n is the number of spikes.

We then added the absolute value of the z -scored mean vector to the z -scored variance vector, and summed across phosphors. This provided a 1,280-component vector (64 pixels \times 20 frames) that reflects both first- and second-order deviations from chance. The frame and pixel in which this deviation was greatest was selected for analysis.

Models

Under a linear model, a neuron's response, r , is equal to the dot product of a stimulus vector, s , onto the cell's linear kernel, w

$$r = w \cdot s \quad (6)$$

An obvious failure of this model is the existence of spiking thresholds. This can be formalized as a static nonlinearity that maps a linear "generator signal" to a nonnegative firing rate (Chichilnisky 2001). The inclusion of such a static nonlinearity provides a simple nonlinear model that subsumes the linear model as a special case

$$r = f(w \cdot s) \quad (7)$$

The weights with which such a cell combines cone inputs (which may vary as a function of time and space) are given by w , the cell's linear kernel. Conveniently, under the conditions of our experiment, the spike-triggered average converges to w under either model (Chichilnisky 2001).

A natural extension of the simple nonlinear model in Eq. 3 includes multiple $w_i s$

$$r = f(w_1 \cdot s, w_2 \cdot s, w_3 \cdot s, \dots, w_n \cdot s) \quad (8)$$

Under this more general model, the neuron's response to a stimulus, s , is a nonlinear function of the stimulus projection onto n different vectors. If the dimensionality of the subspace spanned by the vectors $\{w_1 \dots w_n\}$ is smaller than the dimensionality of the stimulus space, this model provides a parsimonious description of the neuron's response properties. Note that any choice of $\{w_1 \dots w_n\}$ that spans the same subspace can lead to an equivalent description since the static nonlinearity, $f()$, is unconstrained. Similarly, note that changes in the relative magnitude of the w_i vectors can be compensated for exactly by a change in the domain of the nonlinearity, $f()$. The magnitudes of the w_i vectors are therefore arbitrary without specification of $f()$.

In this report, we considered the two-dimensional version of the model in Eq. 8

$$r = f(w_1 \cdot s, w_2 \cdot s) \quad (9)$$

For each neuron studied, we took the spike-triggered average stimulus as an estimate of w_1 and the first principal component of the spike-triggered stimuli as an estimate of w_2 . We have no theoretical reason to expect that this choice (which is arbitrary apart from its simplicity) should be interpretable in terms of the underlying biology. That this was the case for blue-yellow color opponent neurons was fortuitous (as elaborated in the DISCUSSION).

The two-dimensional response surface, $f()$, was estimated by representing each stimulus as a vector, projecting these vectors onto the estimates of w_1 and w_2 , binning projection magnitudes, and calculating firing rates for stimuli within each bin. To avoid bias, different portions of the data were used to estimate the w_i and the response surface. This procedure was repeated 100 times with different dataset partitions, and the resultant response surfaces were averaged together.

RESULTS

Linear cells

We studied 158 V1 neurons in three monkeys. For each neuron, the collection of stimuli preceding spikes was averaged to derive the spike-triggered average stimulus. For neurons that integrate cone inputs linearly, the spike-triggered average reveals the spatiotemporal pattern of weights with which cone signals are combined and thus comprises the neuron's "linear receptive field (RF)." The *left column* of Fig. 3 shows linear RFs for six example V1 neurons. Linear RFs of example cells 1, 2, and 3 had distinct ON and OFF subunits with large LM cone weights that were in phase with weaker S cone weights. Example cells 4–6 had spatially opponent linear RFs and opponent S and LM weights. The presence of multiple RF subunits with complementary color preferences is consistent with previous reports (Conway et al. 2002; Gouras 1974; Hubel and Wiesel 1968; Johnson et al. 2001; Livingstone and Hubel 1984; Michael 1978; Thorell et al. 1984).

Linear RFs can conceal the stimulus preferences of nonlinear cells. We define the "nonlinear RF" as the first principal component of the spike-triggered stimuli. This nonlinear RF has been used to describe the nonlinear properties of complex cells (see Rust et al. 2004; Touryan et al. 2002) (see also METHODS). Panels on the right of Fig. 3, show nonlinear RFs of the six example neurons. Nonlinear RFs of these neurons were unstructured and not significant (randomization test, $P > 0.05$), consistent with the idea that these cells combine cone inputs linearly.

Figure 4 shows linear and nonlinear RFs in space and time. Spatial RF maps are replotted from Fig. 3, and a single pixel is highlighted in each one, selected according to the criteria in the METHODS. Next to each spatial map, we plot a "temporal RF" computed at the highlighted pixel. Temporal linear RFs for example cells 1–3 were LM-dominated, had a short-duration, and peaked ~50–60 ms before the spike. Temporal linear RFs for example cells 4–6 showed opposite signs of response to LM and S-cone excitation, had a longer duration, and peaked ~60–70 ms before the spike. None of the temporal nonlinear RFs were significant ($P > 0.05$).

Nonlinear cells

In contrast to the example cells in Fig. 3 and 4, many cells in our dataset had responses above and beyond those revealed by the linear RF. Linear RFs of the cells shown in Fig. 5 were similar to those shown in Fig. 3 (example cells 4–6), but nonlinear RFs were importantly different. Each nonlinear RF contained at least one pixel with a large positive LM weight adjacent to at least one pixel with a large negative LM weight. The presence of this pattern in the nonlinear RF implies that these cells responded either to the particular light/dark pattern displayed or to the reversed pattern, as explained in METHODS

and verified in *Linear and nonlinear RF interaction*. These nonlinear RFs were highly structured and significant ($P < 0.05$), indicating that these cells did not combine cone inputs linearly.

Figure 6 shows temporal linear and nonlinear RFs at the selected pixel. Linear RFs resembled those shown in Fig. 4. Nonlinear RFs were highly structured, exhibiting mono- or biphasic modulations of LM cone weights and smaller, usually in-phase, S cone weight modulations. This pattern is similar to the linear RFs of example cells 1–3 in Fig. 4. All of the temporal nonlinear RFs in Fig. 6 were statistically significant ($P < 0.05$).

We recorded from 44 blue-yellow neurons (31 dominantly blue-ON/yellow-OFF and 13 dominantly blue-OFF/yellow-ON cells). Receptive fields of these 44 neurons ranged in eccentricity from 2.3 to 7.7° (mean: 5.3°). Twenty-seven (60%) of these neurons had significant nonlinear RFs ($P < 0.05$). This is strong evidence against the hypothesis that this class of V1 cell combines cone inputs linearly. Importantly, this is true irrespective of any static nonlinearity (e.g., spike thresholds) that occurs after a linear combination of cone signals (see METHODS). Neither chromatic aberration nor eye movements can account for this result (see *Eye movements* and *Chromatic aberration*). Similar results were obtained when we omitted the first 2 s of each trial from the analysis, suggesting that the nonlinear RFs we observed were not due to transient processes associated with the onset of visual stimulation (data not shown).

Cone weights of linear and nonlinear RFs

Linear and nonlinear RFs in Fig. 6 had strikingly different spectral signatures: linear RFs were S-dominated whereas nonlinear RFs were LM-dominated. To quantify the color tuning of linear and nonlinear RFs, we considered the single pixel selected by the criteria provided in METHODS, thereby minimizing errors incurred by pooling across differently tuned subunits. Figure 7A illustrates the procedure we used to estimate cone weights. First, we calculated temporal linear and nonlinear RFs in the space of monitor phosphors. These temporal RFs were approximately separable in color and time, so we used a singular value decomposition to find the R, G, B triplet that best characterized the color tuning (Strang 1988). We then transformed this R, G, B triplet into S and LM cone weights as described in METHODS.

The distribution of estimated cone weights is plotted as an angular histogram in Fig. 7B. Linear RFs (gray wedges) had relatively large S cone weights that were either positive or negative and usually had an LM contribution of the opposite sign. Nonlinear RFs (black wedges), in contrast, had large LM weights with smaller, usually nonopponent S cone weights.

Linear and nonlinear RF interaction

Neither the linear nor the nonlinear RF, in isolation, characterizes completely the responses of the neurons we studied. Instead, each reveals a different aspect of the stimuli that excited the cell. To explore how neural responses depended on the stimulation of both RFs jointly, we reconstructed our stimulus movie and, on each frame, measured the degree of excitation to the linear and nonlinear RFs (see METHODS). We




FIG. 3. Linear and nonlinear spatial RFs of 6 example cells. Each row is data from a single cell. Linear and nonlinear RFs are shown as grayscale maps of S and LM cone weights and in color representing the superposition of these 2 maps. Cone weights were scaled by the SD of the cone weights computed at a noncausal delay. This scaling equates the noise level across different linear RF maps and is an arbitrary choice for the nonlinear RF maps. Magenta dots represent weights that were >3 SD away from 0. Magenta dots are not plotted on nonlinear RFs as dependencies across the pixels render this analysis uninterpretable. Prespike delays for calculation of linear and nonlinear RFs was 50, 60, 50, 70, 70, and 60 ms for the 6 example cells, respectively.

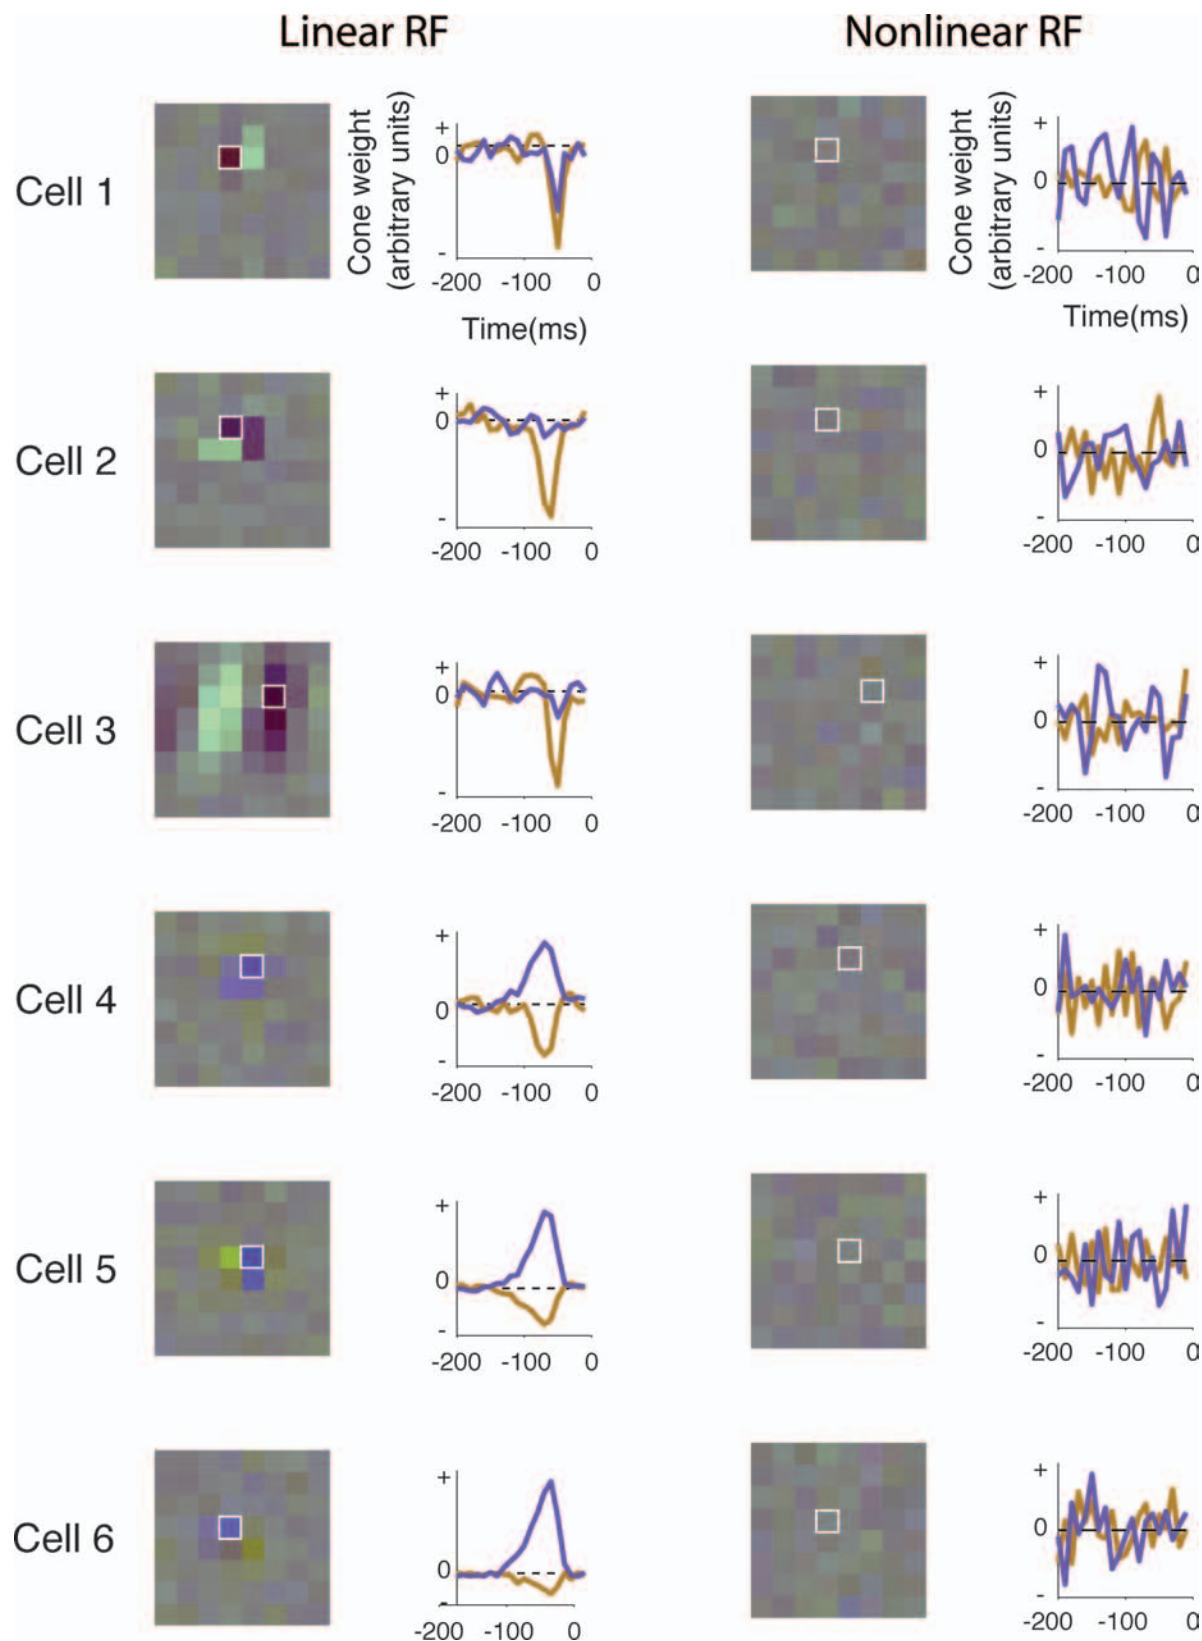


FIG. 4. Linear and nonlinear RFs of 6 example cells. Each row is data from a single cell. Results from 4 analyses are shown: spatial linear RF (1st column), temporal linear RF (2nd column), spatial nonlinear RF (3rd column), and temporal nonlinear RF (4th column). White boxes in spatial linear RFs indicate the pixel selected for temporal analysis by the criteria provided in METHODS. Blue traces in the temporal plots represent S cone weights and brown traces represent LM weights. Dashed lines in temporal plots indicate a cone weight of 0.




FIG. 5. Linear and nonlinear RFs of 6 example nonlinear blue-yellow cells. Conventions are as in Fig. 3. Prespike delays for calculation of linear and nonlinear RFs was 70 ms for all of the cells except *cell 1* for which the delay was 60 ms.

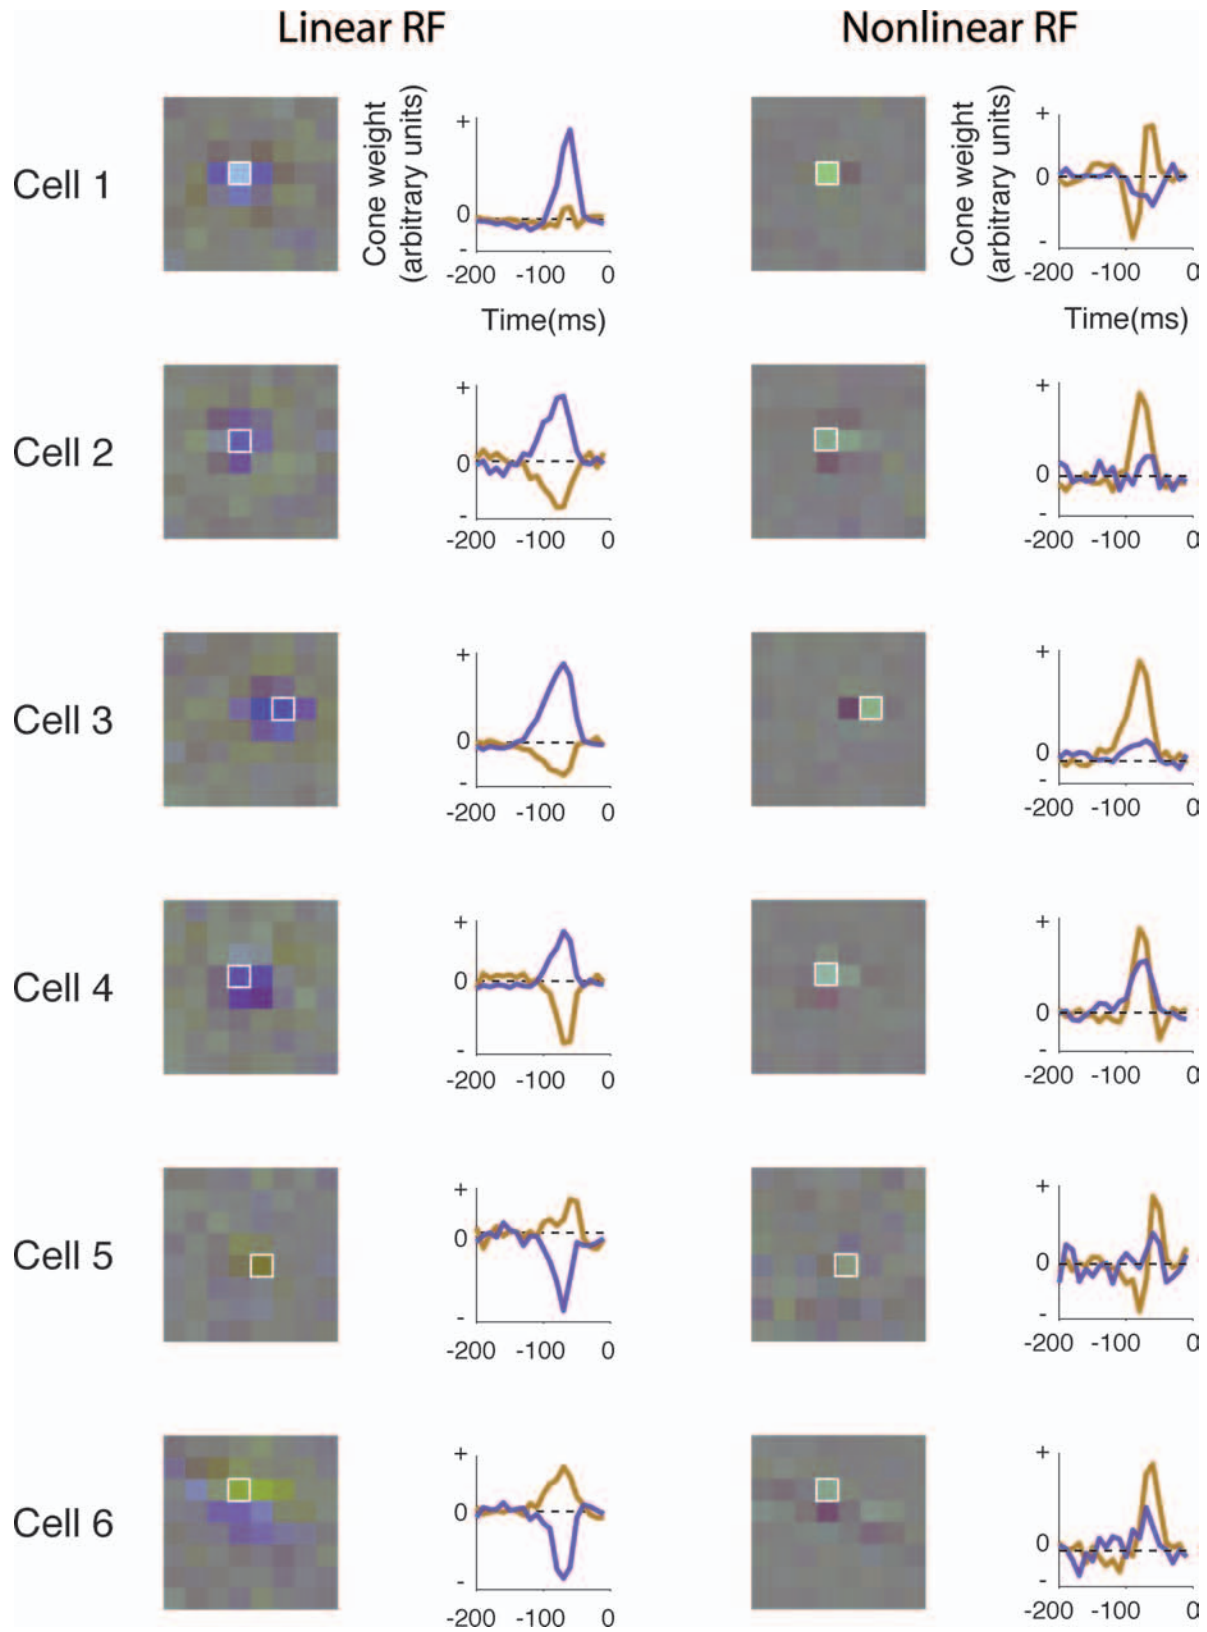


FIG. 6. Linear and nonlinear RFs of 6 example nonlinear blue-yellow cells. Conventions are as in Fig. 4.

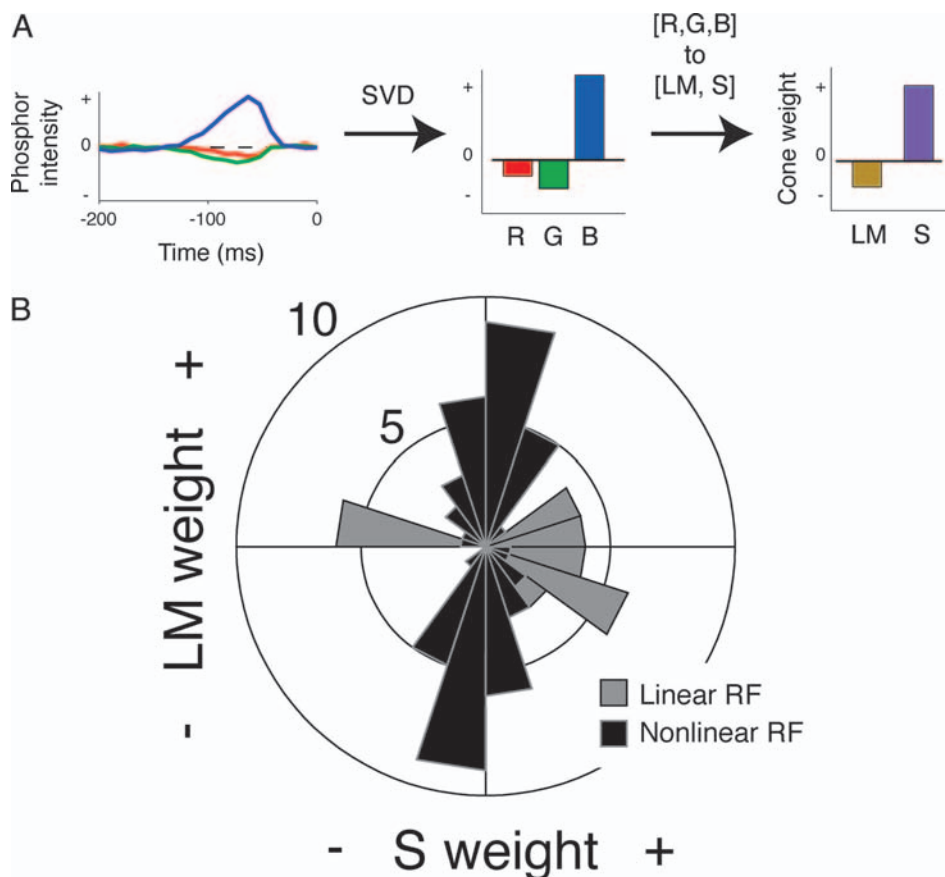


FIG. 7. *A*: procedure for estimating cone weights from temporal linear and nonlinear RFs (see text for details). SVD indicates the singular value decomposition. [R,G,B to LM,S] is a matrix multiplication (see METHODS). *B*: angular histogram of normalized cone weights of linear (gray) and nonlinear (black) RFs. S and LM cone weights were derived from each linear and nonlinear RF and normalized so that their squares summed to 1. Each pair of S and LM cone weights can thus be represented as a 2-element vector the angle of which reveals the relative S to LM cone weight. This plot shows the distribution of the angles of these vectors, or equivalently, the arctangent of the normalized S and LM cone weights. The radius of each wedge corresponds to the number of cells in each angular bin. The relative sign of S and LM cone weights is meaningful for both linear and nonlinear RFs, but the overall sign of the nonlinear RF is arbitrary, so nonlinear RFs are plotted on both sides of the origin.

then pooled stimuli according to how strongly they excited the linear and nonlinear RFs and calculated the responses these stimuli produced. Repeating this process for nonoverlapping groups of stimuli allowed us to construct a two-dimensional response surface for each neuron that showed how neural response depended on RF stimulation. The results of this analysis were qualitatively similar whether we considered spatial or temporal RFs (data shown are from the temporal analysis).

Figure 8 shows response surfaces for the six example neurons shown in Figs. 5 and 6 (*cells 1, 2, etc.* appear in *A, B, etc.*). Within each panel, firing rates increased from bottom to top, indicating that stimuli that activated the linear RF elicited more spikes than stimuli that did not. Importantly, however, for any given level of linear RF stimulation, the firing rate increased with either positive or negative stimulation of the nonlinear RF. The neurons thus responded well to stimuli that matched the linear RF but even more strongly to stimuli that matched the superposition of the linear RF and the nonlinear RF. These composite stimuli resemble blue (or yellow) patches superimposed on achromatic edges.

To quantify the influence of the nonlinear RF on firing rate, we considered responses to stimuli that excited the nonlinear RF weakly (Fig. 8*G, inset*, ○), and compared this response to that driven by stimuli that excited the nonlinear RF strongly (Fig. 8*G, inset*, ×). Critical to this analysis, the two sets of stimuli excited the linear RF identically. The percent change in the firing rate between these two conditions is shown as a histogram in Fig. 8*G*. On average, exciting the nonlinear RF

increased responses by 42%. The nonlinear RFs thus plays a substantial role in driving these 27 blue-yellow neurons.

In principle, response surfaces could be difficult to describe succinctly, but we found that each was determined almost perfectly from its one-dimensional (1-D) marginals. For each response surface, we fit a separable model by multiplying together the 1-D X and Y marginal functions, so that each row (column) in the fitted response surface was constrained to be a scaled version of every other row (column). Correlations between model fits and empirical response surfaces were quite high (mean *r*: 0.995, SD: 0.005), indicating that the responses of these neurons can be approximated as the product of two signals, each of which is the nonlinearly transformed output of a linear filter (the nonlinearity assigned to each filter is provided by the corresponding marginal function of the response surface). For comparison, we also considered a model in which the response of the neuron was given by the *sum* of two signals (i.e., the rows and columns of the fitted response surface were constrained to be shifted, as opposed to scaled, versions of each other). This model also fit the data quite well, but demonstrably less well than the separable model (data not shown).

Absence of nonlinear RFs in LGN neurons

Neurons in the lateral geniculate nucleus (LGN) have been shown to integrate cone inputs nearly linearly (Derrington et al. 1984), thus providing a natural control on our analysis method. We studied 78 multiunit sites from the LGN of a fourth monkey in search of blue-yellow neurons. Eight of these sites were dominated by one or two single units, yielded clear, stable

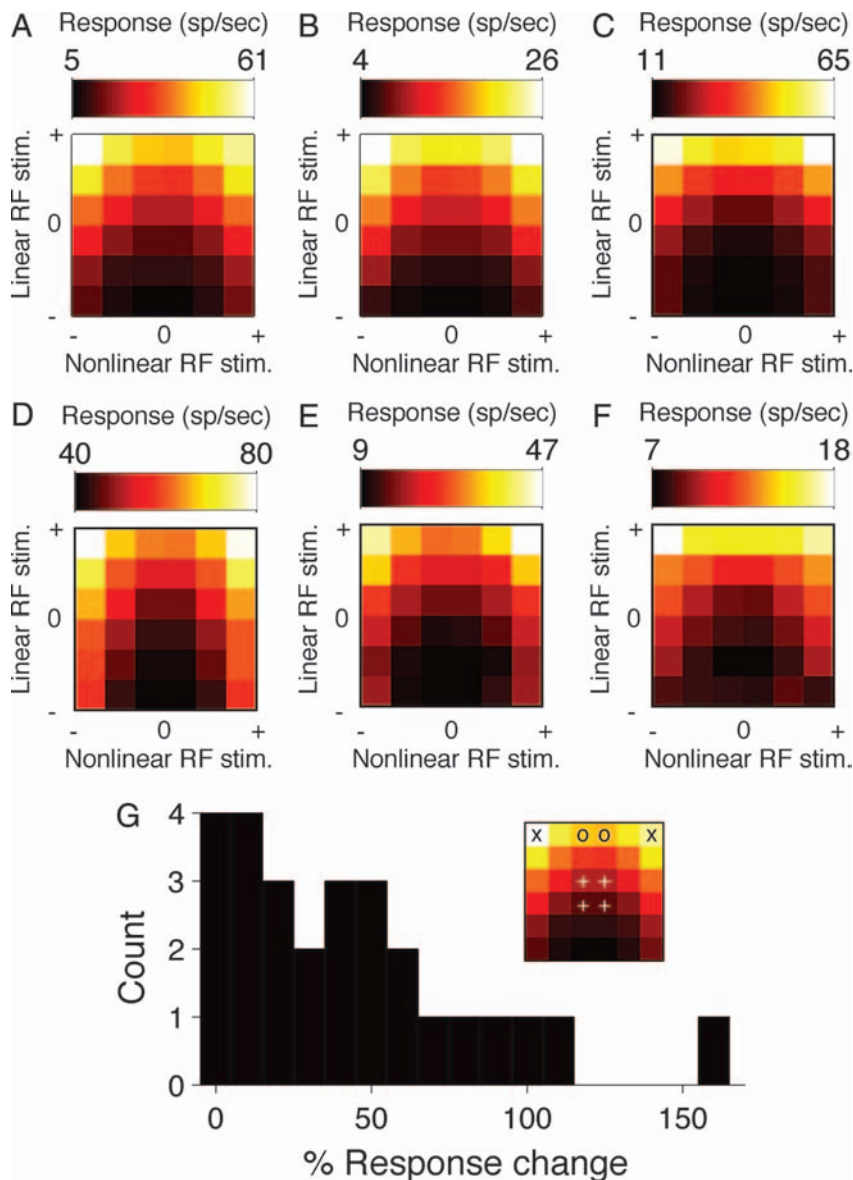


FIG. 8. *A–F*: firing rate of neurons in Fig. 3 as a function of linear and nonlinear RF excitation. These neurons respond most strongly to stimuli that excite the linear RF positively and the nonlinear RF either positively or negatively. Color bars at the top of each panel indicate firing rate ranges. *G*: histogram of firing rate increases that result from nonlinear RF stimulation. ○, responses to stimuli that did not excite the nonlinear RF and ×, those that did compared after subtraction of baseline firing rate (+). Each neuron had different linear and nonlinear RFs, but on average, ○ corresponded to a $46 \pm 8\%$ (mean \pm SD) change in S cone contrast and an $8 \pm 5\%$ change in LM contrast (of opposite sign), and × corresponded to a $38 \pm 13\%$ change in S cone contrast and a $14 \pm 6\%$ change in LM contrast.

linear RFs, and exhibited blue-yellow opponency. Here we consider these eight sites. Importantly, multiunit recording is conservative with respect to the hypothesis we are testing: pooling spikes from heterogeneously-tuned neurons would be expected to increase the significance of the nonlinear RF. RF eccentricities ranged from 4.2 to 12° (mean: 7.9°). The spatial substructure of parafoveal LGN RFs is too fine to be resolved by our display (Fig. 1), so we tested LGN cells with a 3° diam disk whose color modulated randomly with statistics identical to a single pixel in the 8×8 grid.

Figure 9, *A–C*, shows linear and nonlinear RFs from three representative LGN recordings. Every linear RF exhibited a consistent, short-latency blue-yellow signal, and every nonlinear RF was unstructured and not significant ($P > 0.05$). To control for the possibility that spatially uniform stimulation never reveals nonlinear RFs, we also stimulated 25 blue-yellow V1 neurons with the 3° diam disk. Twenty-one of these neurons had significant nonlinear RFs. Linear and nonlinear RFs obtained from V1 neurons under spatially uniform stimulation often resembled those obtained with spatially-varying

stimulation (Fig. 9, *D–F*). The absence of nonlinear RFs in LGN cells was therefore not a trivial consequence of uniform spatial stimulation. We conclude that our analysis method does not create nonlinear RFs where none exist, and that the interaction between opponent and nonopponent signals that we have documented in V1 is unlikely to reflect signal integration occurring at the level of the LGN or retina.

Eye movements

Monkeys in this study were trained to maintain stable fixation, and trials were terminated whenever the eye left a $1 \times 1^\circ$ electronically defined window surrounding the fixation point. Small eye movements within this window, however, caused the stimulus to shift on the retina.

Critically for this study, fixational eye movements would not be expected to give rise to the nonlinear RFs we observed. Eye movements can create nonlinear RFs by moving individual pixels among subunits, but the color signature of such nonlinear RFs will reflect the combination of the subunits from which

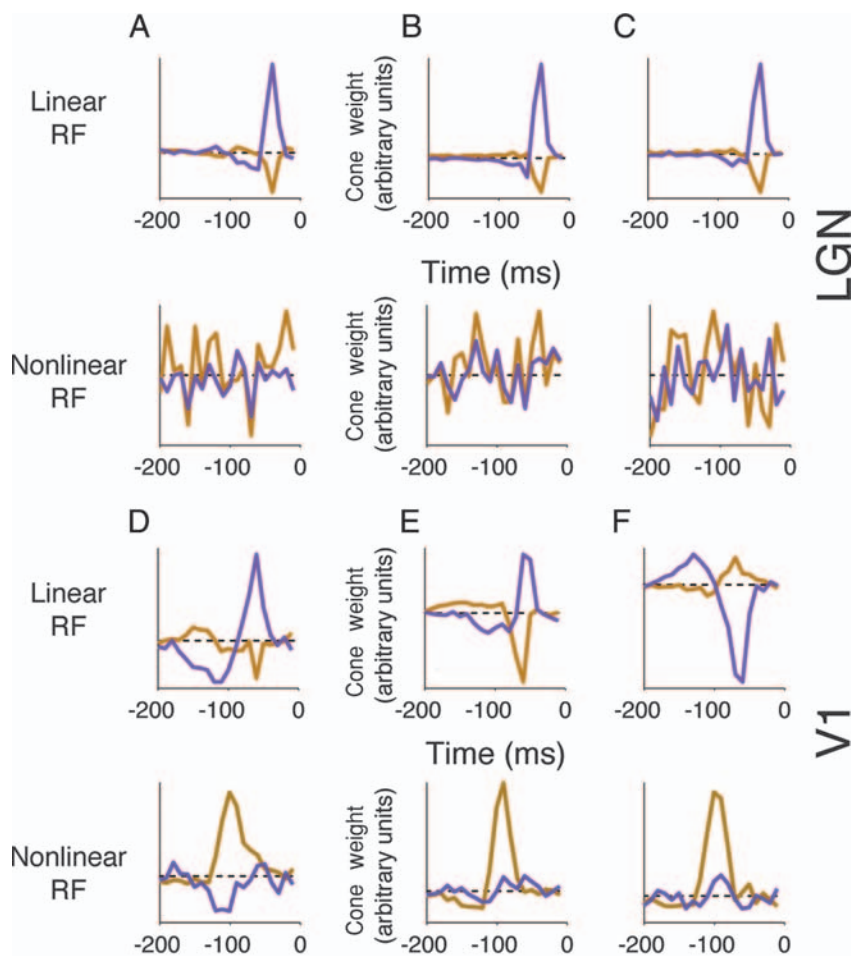


FIG. 9. Linear and nonlinear RFs from 3 multiunit sites in the LGN; (A–C) and from 3 single units in V1 (D–F). ---, cone weights of 0. Nonlinear RFs for all LGN recordings were unstructured and not significant ($P > 0.05$). Nonlinear RFs for all V1 recordings shown are significant ($P < 0.05$).

it arose. For example, sequential activation of blue and yellow subunits by a single pixel could give rise to a blue-yellow nonlinear RF because some spike-triggered stimuli will be blue and others will be yellow. The nonlinear RFs we observed had a color signature that was absent in the linear RFs and thus did not arise from shifts of the stimulus on the retina. Moreover, significant nonlinear RFs were observed with spatially uniform stimulation of V1 cells (e.g., Fig. 9, D–F). Because this stimulus lacked spatial contrast within the spatial extent of the RF, eye movements would not be expected to alter visual responses whether reflected in the linear nor nonlinear RF.

The location of the stimulus on the screen and the position of the eye in the head can be used to compute the position of the stimulus on the retina. Our measurements thus allowed us, in principle, to compensate for changes in fixation position. The quality of our eye position records, however, did not permit such compensation consistently. Our eye position estimates, although accurate over short time scales (noise $< 0.01^\circ$), were inaccurate over long time scales during which slow drifts changed estimated eye position over as much as several tenths of a degree (> 1 pixel width) (Read and Cumming 2003; Tsao et al. 2003). Some of the slow changes in estimated eye position were measurement artifacts related to small changes in room temperature.

Within individual trials, fixation could be measured accurately and was reasonably stable. Seventy percent of eye-position samples in individual trials were confined to an area the size of a single stimulus pixel ($0.22 \times 0.22^\circ$). The SD of

eye position (averaged across horizontal and vertical channels) within trials was 0.07° . As expected, this SD depended on trial duration ($r = 0.28$, $P < 0.0001$) because small displacements in eye position accumulate over time. However, eye position was only slightly more variable (SD: 0.08°) for trials that lasted the maximal duration (10 s). Seventy-eight percent of saccades within the fixation window were smaller than the width of a single stimulus pixel (mean saccade amplitude: 0.16° , SD: 0.10°). Eye movements within individual trials thus affected our spike-triggered analyses negligibly.

Estimating changes in eye position between trials is more difficult than estimating changes within trials because of slow drifts in the eye position signal. To estimate between-trial fixation variability, we measured differences in median eye position on consecutive trials, between which relatively little drift presumably occurred (Read and Cumming 2003). Fifty-percent of the median differences lay within the single pixel boundary. The SD of the median differences was 0.14° . We conclude that measured fixation positions were reasonably consistent across consecutive trials, although slow drifts in estimated eye position undoubtedly contributed to the measured variability. The fact that many of the linear RFs that we studied (e.g., Figure 3) had distinguishable, adjacent ON and OFF subunits attests further to the stability of fixation.

Although eye movements would not be expected to lead artifactually to the nonlinear RFs we studied, the combination of eye movement and the spatial coarseness of our stimulus undoubtedly obscured high-frequency structure in the spatial

RFs we studied. One specific possibility is that a spatially phase-specific LM signal may coexist with the nonlinear LM signal we report here, but that this signal may have been of sufficiently high spatial frequency to be masked in our linear RFs. Experiments in anesthetized animals could test this hypothesis by providing higher resolution spatial maps of linear and nonlinear RFs.

Chromatic aberration

Any stimulus that varies spatially and spectrally is subject to chromatic aberration that shifts (transverse chromatic aberration) and defocuses (axial chromatic aberration) light in a wavelength dependent fashion. Several lines of evidence indicate that our results did not arise trivially from chromatic aberration. First, both axial and transverse chromatic aberrations over the spatial extent of our stimulus can be modeled as linear transformations (Marimont and Wandell 1994; Thibos et al. 1990). Linear transformation of the stimulus might distort linear RF estimates but would not manifest as a significant nonlinear RF in a linear neuron. Second, blue-yellow V1 cells stimulated with a uniformly-colored 3° diam disk yielded significant, LM-dominated nonlinear RFs (Fig. 9, *D–F*). Chromatic aberration should not occur within the spatial extent of the RF with this stimulus because spatial shifts caused by transverse chromatic aberration would be expected to be only ~0.04° (Thibos et al. 1990), which is small relative to the distance from the stimulus edge to the RF. Third, the power spectral density of our stimulus was dominated by low frequency components. Approximately 60% of the power in our stimulus that was in the optically relevant range (<45 cycles/°) was <2 cycles/° and an additional 20% lay between 2 and 5 cycles/°. Most of the power in our stimulus was thus affected minimally by axial chromatic aberration (Marimont and Wandell 1994). Fourth, although chromatic aberration was present to the same degree in all experiments, we found many cells without significant nonlinear RFs indicating that one is not an obligatory consequence of the other. Finally, although significant nonlinear RFs were common among blue-yellow cells, some lacked it (Fig. 3, *cells 4–6*) indicating that a significant nonlinear RF is not a necessary consequence of sensitivity to blue-yellow modulation.

DISCUSSION

We studied a class of neurons in macaque V1 with linear RFs that exhibited S and LM opponency. Principal components analysis on spike-triggered stimuli revealed that, in addition to being color-opponent, many of these cells responded to spatiotemporal LM contrast irrespective of polarity. Joint consideration of both response properties revealed that firing rates could be modeled as a product of opponent and nonopponent signals. Control experiments showed that this interaction between opponent and nonopponent signals was not inherited from the LGN, and, along with arguments provided in the RESULTS, argue against an artifactual origin from fixational eye movements or chromatic aberration. In this DISCUSSION, we relate our results to previous findings, consider the neural substrate underlying the responses we observed, describe some promising extensions of the analysis technique used in this study, and speculate on the perceptual significance of the V1 responses we observed.

Gain control in V1

The interaction we observed between opponent and nonopponent signals is consistent with gain modulation. Gain control in V1 has been studied intensively and can account for a host of nonlinear response properties (Albrecht and Geisler 1991; Heeger 1992). The phenomenon we studied, however, differs from the classic form of gain control that is thought to implement contrast normalization. First, classical contrast gain control is suppressive (Albrecht et al. 1984; Bonds 1991; Schwartz et al. 2001), whereas the phenomenon we studied was facilitatory. Second, suppressive contrast gain control is a nearly ubiquitous property of V1 neurons (Bonds 1991), but we found evidence for excitatory nonlinear interactions in only a small subset of neurons.

Similarly, the nonlinear RFs we observed are not an expected consequence of nonlinear, LM-tuned surrounds. First, unlike the phenomenon we studied, nonlinear surrounds are usually suppressive (Solomon et al. 2004; Ts'o and Gilbert 1988). Second, the energy in the nonlinear RF was usually spatially coincident with the center of the linear RF, not the surrounding area (see Fig. 5). Third, the spatial profile of most nonlinear RFs was an edge, which is inconsistent with any model in which the LM mechanism gives only one sign of response (e.g., ON or OFF).

Relationship to previous studies

Nonlinear color tuning has been documented in previous studies of V1 (Cottaris and DeValois 1998; DeValois et al. 2000; Hanazawa et al. 2000; Lennie et al. 1990; Solomon et al. 2004; Wachtler et al. 2000). Modeling attempts, however, have focused on a restricted class of nonlinear models, that is: linear summation of cone inputs followed by a static nonlinearity (*Eq. 7* in METHODS). Although this class of models is a significant improvement over the strictly linear model, it qualitatively fails to explain the color tuning of the blue-yellow neurons we studied.

Johnson et al. (2001) found that color-luminance cells that receive strong S cone input tend to prefer high spatial frequency luminance gratings but relatively low spatial frequency S cone isolating gratings (see also Thorell et al. 1984, Fig. 12*B*). These neurons may overlap with the population we studied, in which case this counterintuitive result can be understood in terms of the spatial frequency tuning of the linear and nonlinear RFs. As shown in Fig. 5, nonlinear RFs tended to be tuned for higher spatial frequencies than linear RFs. Thus an S cone isolating grating, which should stimulate primarily linear RFs, would drive the cells maximally at a relatively low spatial frequency, and a luminance grating, which should stimulate primarily nonlinear RFs, would drive the cells maximally at a relatively high spatial frequency. We would predict that these cells would respond to a drifting S cone isolating grating with a modulated response and a luminance grating with a relatively unmodulated response.

Stimulation of V1 neurons with large, uniformly colored stimuli leads to color-tuning dynamics in some neurons (Cottaris and DeValois 1998). We observed qualitatively similar color-tuning dynamics in our data (e.g., Fig. 9*D*). These dynamics might be expected if opponent and nonopponent signals act on the cell with different latencies. Indeed, under spatially

uniform stimulation, nonlinear RFs (which were dominantly nonopponent) tended to peak before linear RFs (which were dominantly opponent). This explanation, however, cannot account for the findings of Cottaris and DeValois First, the stimulus used by Cottaris and DeValois did not vary in photometric luminance, so their data may have been affected little by the nonopponent signal we studied. Second, the temporal pattern of responses observed by Cottaris and DeValois was opposite what we would expect from the current results. Cottaris and DeValois found that many cells carried long-latency blue-yellow opponent signals, whereas we would have predicted a short-latency blue-yellow signal, preceded by the projection of the nonopponent signal onto the isoluminant plane. The color-tuning dynamics observed by Cottaris and DeValois is therefore unlikely to arise from the nonlinear interaction we studied.

Neural substrate of color interaction

We do not know the pathways by which the opponent and nonopponent signals that we studied reach V1, but previous physiological and anatomical studies suggest possibilities. For instance, the magnocellular pathway could provide an LM-dominated signal (via spiny stellate cells in layer 4C α) that modulates the gain of blue-yellow signals carried by the koniocellular (and possibly parvocellular) afferents to layers 4A and 2/3 of V1, as schematized in Fig. 10 (Allison et al. 2000; Chatterjee and Callaway 2003; Lachica et al. 1993; Nealey and Maunsell 1994; Schiller and Malpeli 1978; Yabuta et al. 2001). In this scenario, we would expect that ablation of magnocellular neurons would eliminate nonlinear RFs but spare linear RFs.

The interaction between cone-opponent and nonopponent signals we have documented may not be specific to a blue-yellow system. In our study, linear RFs showed blue-yellow opponency and nonlinear RFs showed nonopponent sensitivity. This dissociation suggests that the mechanisms underlying linear and nonlinear RFs may have been fortuitously close to orthogonal in the (device-specific) space in which color channels modulated independently (see METHODS). The possibility remains that other types of color-opponent V1 neurons are excited by LM contrast like blue-yellow neurons are, but that opponent and nonopponent mechanisms are sufficiently non-orthogonal in these neurons as to yield confounded linear and nonlinear RFs under the current experimental conditions.

Extensions of the approach

In this study, we ignored all but the first principal component of the spike-triggered stimulus distribution, but most of the neurons we studied had multiple significant principal components. These additional principal components did not reflect other gain modulating signals, but rather appeared to reflect a common underlying LM-dominated mechanism (data not shown).

Spatially, these principal components tended to resemble luminance edges of different orientations and spatial phases. Temporally, principal components tended to resemble phase-shifted versions of each other, consistent with the integration of contrast energy in a particular temporal frequency band or the presence of non-Poisson spiking statistics (Pillow and Simoncelli 2003). Neither explanation is consistent with the model on which our analysis is based (Eq. 9 in METHODS). An extension of the model—including a second linear filtering stage after the nonlinearity—might account for these multiple principal components and provide a concise, biologically plausible characterization of the neurons we studied (Victor 1988; Victor and Shapley 1979).

Linear and nonlinear RFs are not expected, in general, to reflect biologically meaningful mechanisms. Thus a general-purpose technique for recovering the precise characteristics of the mechanisms underlying linear and nonlinear RFs would require additional steps (e.g., combination of linear and nonlinear RFs). Support for this idea comes from the fact that linear and nonlinear RFs of some apparently dissimilar neurons in our dataset spanned a similar plane in stimulus space. These neurons may be fed by identical mechanisms that differ only in their relative contribution to the response. Furthermore, many response surfaces like those shown in Fig. 8 can be made more separable with a rotation of the coordinate axes. One area for future research is the development of techniques for finding biologically relevant bases that span the subspace of relevant stimuli.

PCA has the advantage of having well-understood properties, and it is closely related to the second-order Wiener kernel, which has been used to characterize nonlinear neurons previously (Marmarelis and Marmarelis 1978). However, PCA considers only second-order stimulus interactions, and it may be misleading when applied to neurons whose responses depend on higher-order stimulus interactions. A generalization of the approach used in this study is to parameterize a set of stimuli,




FIG. 10. A model that accounts for the responses of the blue-yellow cells we studied. A blue-yellow color-opponent signal (perhaps derived from the koniocellular and/or parvocellular pathway) is passed through a monotonic nonlinearity, and this signal is multiplied with a nonopponent signal (perhaps derived from the magnocellular pathway) that has been passed through a nonmonotonic nonlinearity.

each of which can then be represented as a point in an abstract stimulus space, and to characterize the neuron's responses as the ratio of the response-triggered stimulus density (spike-triggered stimulus density, in our study), to the a priori stimulus density (de Ruyter van Steveninck and Bialek 1988; Sharpee et al. 2004). The typically high dimensionality of visual stimulus parameterizations makes a brute force approach impractical, but any of a number of dimensionality reduction techniques can render the problem tractable.

Perceptual significance

Monkeys in our study were rewarded for simply fixating, so we cannot draw firm conclusions about the relationship between the neuronal responses we recorded and color perception, but we can speculate on the relationship between our neurophysiological results and those from human psychophysics. Colored targets are more readily detected when superimposed on luminance pedestals (Eskew et al. 1994; Gowdy et al. 1999; Gur and Akri 1992; Hilz et al. 1974; Switkes et al. 1988). This perceptual phenomenon may be a consequence of LM-induced facilitation of color-opponent V1 neurons. Luminance patterns containing sharp edges improve chromatic detection profoundly (Gowdy et al. 1999), and such patterns would be expected to activate the nonlinear RFs we observed particularly strongly (Fig. 5). Chromatic disks are most easily detected ~20 ms after a spatially-aligned luminance flash (Eskew et al. 1994), and in our experiments with uniformly-colored disks, the LM signal in the nonlinear RF often led the blue-yellow signal in the linear RF by approximately this latency (Fig. 9, D-F).

A facilitatory interaction between color and luminance could provide an efficient means to estimate surface reflectance. Changes in surface reflectance occur preferentially at object boundaries, which are usually marked by discontinuities in both luminance and chromaticity (Fine et al. 2003; Kingdom 2003; Ruderman et al. 1998). Enhancement of chromatic signals at luminance edges may thus be an important step in the computation of object color.

ACKNOWLEDGMENTS

We thank J. Costanza, L. Abavare, D. Diep, and F. Davalian for technical assistance during the course of the study. B. Wandell, C. Barberini, E. Callaway, E. Simoncelli, O. Schwartz, and several anonymous reviewers provided valuable comments on the manuscript. E. Simoncelli also provided statistical assistance.

GRANTS

G. D. Horwitz was supported by the Helen Hay Whitney Foundation. E. J. Chichilnisky was supported by the McKnight Foundation and Sloan Foundation. T. D. Albright is an investigator of the Howard Hughes Medical Institute.

REFERENCES

Albrecht DG, Farrar SB, and Hamilton DB. Spatial contrast adaptation characteristics of neurones recorded in the cat's visual cortex. *J Physiol* 347: 713–739, 1984.

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.

Allison JD, Melzer P, Ding Y, Bonds AB, and Casagrande VA. Differential contributions of magnocellular and parvocellular pathways to the contrast response of neurons in bush baby primary visual cortex (V1). *Vis Neurosci* 17: 71–76, 2000.

Baylor DA, Nunn BJ, and Schnapf JL. Spectral sensitivity of cones in the monkey *Macaca fascicularis*. *J Physiol* 390: 145–160, 1987.

Bonds AB. Temporal dynamics of contrast gain in single cells of the cat striate cortex. *Vis Neurosci* 2: 239–255, 1991.

Brainard DH. The psychophysics toolbox. *Spat Vis* 10: 433–436, 1997.

Chatterjee S and Callaway EM. Parallel color-opponent pathways to primary visual cortex. *Nature* 426: 668–671, 2003.

Chichilnisky EJ. A simple white noise analysis of neuronal light responses. *Network* 12: 199–213, 2001.

Conway BR. Spatial structure of cone inputs to color cells in alert macaque primary visual cortex (V-1). *J Neurosci* 21: 2768–2783, 2001.

Conway BR, Hubel DH, and Livingstone MS. Color contrast in macaque V1. *Cereb Cortex* 12: 915–925, 2002.

Cottaris NP and DeValois RL. Temporal dynamics of chromatic tuning in macaque primary visual cortex. *Nature* 395: 896–900, 1998.

de Ruyter van Steveninck R and Bialek W. Coding and information transfer in short spike sequences. *Proc Soc Lond B Biol Sci* 234: 379–414, 1988.

Derrington AM, Krauskopf J, and Lennie P. Chromatic mechanisms in the lateral geniculate nucleus of macaques. *J Physiol* 357: 241–265, 1984.

DeValois RL. Analysis and coding of color vision in the primate visual system. *Cold Spring Harbor Symp Quant Biol* 30: 567–579, 1965.

DeValois RL, Cottaris NP, Elfar SD, Mahon LE, and Wilson JA. Some transformations of color information from lateral geniculate nucleus to striate cortex. *Proc Natl Acad Sci USA* 97: 4997–5002, 2000.

Dobkins KR and Albright TD. What happens if it changes color when it moves?: the nature of chromatic input to macaque visual area MT. *J Neurosci* 14: 4854–4870, 1994.

Draper NR and Smith H. *Applied Regression Analysis*. New York: Wiley, 1998.

Eskew RT Jr, Stromeyer CF 3rd, and Kronauer RE. The time-course of chromatic facilitation by luminance contours. *Vision Res* 34: 3139–3144, 1994.

Fine I, MacLeod DI, and Boynton GM. Surface segmentation based on the luminance and color statistics of natural scenes. *J Opt Soc Am A Opt Image Sci Vis* 20: 1283–1291, 2003.

Gouras P. Identification of cone mechanisms in monkey ganglion cells. *J Physiol* 199: 533–547, 1968.

Gouras P. Opponent-color cells in different layers of foveal striate cortex. *J Physiol* 238: 583–602, 1974.

Gowdy PD, Stromeyer CF 3rd, and Kronauer RE. Facilitation between the luminance and red-green detection mechanisms: enhancing contrast differences across edges. *Vision Res* 39: 4098–4112, 1999.

Gur M and Akri V. Isoluminant stimuli may not expose the full contribution of color to visual functioning: spatial contrast sensitivity measurements indicate interaction between color and luminance processing. *Vision Res* 32: 1253–1262, 1992.

Hanazawa A, Komatsu H, and Murakami I. Neural selectivity for hue and saturation of color in the primary visual cortex of the monkey. *Eur J Neurosci* 12: 1753–1763, 2000.

Heeger DJ. Normalization of cell responses in cat striate cortex. *Vis Neurosci* 9: 181–198, 1992.

Hilz RL, Huppmann G, and Cavonius CR. Influence of luminance contrast on hue discrimination. *J Opt Soc Am* 64: 763–766, 1974.

Hubel DH and Wiesel TN. Receptive fields and functional architecture of monkey striate cortex. *J Physiol* 195: 215–243, 1968.

Johnson EN, Hawken MJ, and Shapley R. The spatial transformation of color in the primary visual cortex of the macaque monkey. *Nat Neurosci* 4: 409–416, 2001.

Johnson EN, Hawken MJ, and Shapley R. Cone inputs in macaque primary visual cortex. *J Neurophysiol* 91: 2501–2514, 2004.

Kingdom FA. Color brings relief to human vision. *Nat Neurosci* 6: 641–644, 2003.

Knoblauch K and D'Zmura M. Reply to letter to editor by M. J. Sankeralli and K. T. Mullen (published in *Vision Res* 41: 53–55, Lights and neural responses do not depend on choice of color space). *Vision Res* 41: 1683–1684, 2001.

Lachica EA, Beck PD, and Casagrande VA. Intrinsic connections of layer III of striate cortex in squirrel monkey and bush baby: correlations with patterns of cytochrome oxidase. *J Comp Neurol* 329: 163–187, 1993.

Landisman CE and Ts'o DY. Color processing in macaque striate cortex: electrophysiological properties. *J Neurophysiol* 87: 3138–3151, 2002.

Lennie P, Krauskopf J, and Sclar G. Chromatic mechanisms in striate cortex of macaque. *J Neurosci* 10: 649–669, 1990.

Leon-Garcia A. *Probability and Random Processes for Electrical Engineering*. Reading, MA: Addison-Wesley, 1994.



- Livingstone MS and Hubel DH.** Anatomy and physiology of a color system in the primate visual cortex. *J Neurosci* 4: 309–356, 1984.
- Marimont DH and Wandell BA.** Matching color images: the impact of axial chromatic aberration. *J Opt Soc Am A* 111: 3113–3122, 1994.
- Marmarelis P and Marmarelis V.** *Analysis of Physiological Systems: The White Noise Approach*. New York: Plenum, 1978.
- Michael CR.** Color vision mechanisms in monkey striate cortex: dual-opponent cells with concentric receptive fields. *J Neurophysiol* 41: 572–588, 1978.
- Nealey TA and Maunsell JH.** Magnocellular and parvocellular contributions to the responses of neurons in macaque striate cortex. *J Neurosci* 14: 2069–2079, 1994.
- Paninski L.** Convergence properties of three spike-triggered analysis techniques. *Network Comput Neural Syst* 14: 437–464, 2003.
- Pillow J and Simoncelli E.** Biases in white noise analysis due to non-Poisson spike generation. *Neurocomputing* 52–54: 109–115, 2003.
- Read JC and Cumming BG.** Measuring V1 receptive fields despite eye movements in awake monkeys. *J Neurophysiol* 90: 946–960, 2003.
- Ruderman D, Cronin T, and Chiao C.** Statistics of cone responses to natural images: implications for visual coding. *J Opt Soc Am A* 15: 2036–2045, 1998.
- Rust N, Schwartz O, Movshon J, and Simoncelli E.** Spike-triggered characterization of excitatory and suppressive stimulus dimensions in monkey V1. *Neurocomputing* 58–60: 793–799, 2004.
- Schiller PH and Malpeli JG.** Functional specificity of lateral geniculate nucleus laminae of the rhesus monkey. *J Neurophysiol* 41: 788–797, 1978.
- Schwartz O, Chichilnisky E, and Simoncelli E.** Characterizing neuronal gain control using spike-triggered covariance. In: *Advances in Neural Information Processing Systems*, edited by Dietterich T, Becker S, and Ghahramani Z. Cambridge, MA: MIT Univ. Press, 2002, p. 269–276.
- Sharpee T, Rust N, and Bialek W.** Analyzing neural responses to natural signals: maximally informative dimensions. *Neural Comput* 16, 2004.
- Simoncelli E, Paninski L, Pillow J, and Schwartz O.** Characterization of neural responses with stochastic stimuli. In: *The New Cognitive Neurosciences* (3rd ed.), edited by Gazzaniga M. Boston, MA: MIT Univ. Press, 2004, p. 327–338.
- Solomon SG, Peirce JW, and Lennie P.** The impact of suppressive surrounds on chromatic properties of cortical neurons. *J Neurosci* 24: 148–160, 2004.
- Stockman A, MacLeod DI, and Johnson NE.** Spectral sensitivities of the human cones. *J Opt Soc Am A* 10: 2491–2521, 1993.
- Strang G.** *Linear Algebra and its Applications*. San Diego: Harcourt, Brace, Jovanovich, 1988.
- Switkes E, Bradley A, and De Valois KK.** Contrast dependence and mechanisms of masking interactions among chromatic and luminance gratings. *J Opt Soc Am A* 5: 1149–1162, 1988.
- Thibos LN, Bradley A, Still DL, Zhang X, and Howarth PA.** Theory and measurement of ocular chromatic aberration. *Vision Res* 30: 33–49, 1990.
- Thorell LG, De Valois RL, and Albrecht DG.** Spatial mapping of monkey V1 cells with pure color and luminance stimuli. *Vision Res* 24: 751–769, 1984.
- Touryan J, Lau B, and Dan Y.** Isolation of relevant visual features from random stimuli for cortical complex cells. *J Neurosci* 22: 10811–10818, 2002.
- Tsao DY, Conway BR, and Livingstone MS.** Receptive fields of disparity-tuned simple cells in macaque V1. *Neuron* 38: 103–114, 2003.
- Ts'o DY and Gilbert CD.** The organization of chromatic and spatial interactions in the primate striate cortex. *J Neurosci* 8: 1712–1727, 1988.
- Vautin RG and Dow BM.** Color cell groups in foveal striate cortex of the behaving macaque. *J Neurophysiol* 54: 273–292, 1985.
- Victor JD.** The dynamics of the cat retinal Y cell subunit. *J Physiol* 405: 289–320, 1988.
- Victor JD and Shapley RM.** The nonlinear pathway of Y ganglion cells in the cat retina. *J Gen Physiol* 74: 671–689, 1979.
- Wachtler T, Sejnowski TJ, and Albright TD.** Responses of cells in macaque V1 to chromatic stimuli are compatible with human color constancy. *Soc Neurosci Abstr* 7.8: 2000.
- Wandell BA.** *Foundations of Vision*. Sunderland, MA: Sinauer, 1995.
- Yabuta NH, Sawatari A, and Callaway EM.** Two functional channels from primary visual cortex to dorsal visual cortical areas. *Science* 292: 297–300, 2001.