

Spatiotemporal Elements of Macaque V1 Receptive Fields

Nicole C. Rust,^{1,*} Odelia Schwartz,³
J. Anthony Movshon,^{1,4} and Eero P. Simoncelli^{1,2,4}

¹Center for Neural Science and

²Howard Hughes Medical Institute

New York University

New York, New York 10003

³Salk Institute and

Howard Hughes Medical Institute

La Jolla, California 92037

Summary

Neurons in primary visual cortex (V1) are commonly classified as simple or complex based upon their sensitivity to the sign of stimulus contrast. The responses of both cell types can be described by a general model in which the outputs of a set of linear filters are nonlinearly combined. We estimated the model for a population of V1 neurons by analyzing the mean and covariance of the spatiotemporal distribution of random bar stimuli that were associated with spikes. This analysis reveals an unsuspected richness of neuronal computation within V1. Specifically, simple and complex cell responses are best described using more linear filters than the one or two found in standard models. Many filters revealed by the model contribute suppressive signals that appear to have a predominantly divisive influence on neuronal firing. Suppressive signals are especially potent in direction-selective cells, where they reduce responses to stimuli moving in the nonpreferred direction.

Introduction

Neurons in primary visual cortex are classically divided into two groups. Simple cells respond precisely to the location and contrast polarity of features in the visual scene, while complex cells measure the magnitude of local contrast without regard to the polarity or precise position of stimulus features (Hubel and Wiesel, 1962; Hubel and Wiesel, 1968). Simple receptive fields can be described by a single linear spatiotemporal filter whose output is half-wave rectified and squared (Figure 1A) (Movshon et al., 1978b; Heeger, 1992a). Complex receptive fields are economically described by two linear spatiotemporal filters whose outputs are squared and summed (the “energy model,” Figure 1B) (Movshon et al., 1978a; Adelson and Bergen, 1985; Spitzer and Hochstein, 1985). The two cell types are characteristically found in different cortical layers and are thought to receive different patterns of input, suggesting that they form distinct groups (Hubel and Wiesel, 1962; Hubel and Wiesel, 1968). Recent theory and experiments have shown, however, that many simple and complex

response properties lie on a continuum (Chance et al., 1999; Mechler and Ringach, 2002; Priebe et al., 2004). It is therefore of interest to note that the standard models of these cells can both be captured by a single general framework in which the outputs of a set of linear spatiotemporal filters are nonlinearly combined (Figure 1C).

To use such a model, it is necessary to estimate its components from physiological data. For each cell, one would like to extract a set of relevant filters and the nonlinear rule by which their outputs are combined. Simple cells can be analyzed using the well-known method of reverse correlation (Jones and Palmer, 1987; DeAngelis et al., 1993). Specifically, an unbiased estimate of the linear filter can be recovered by taking the spatiotemporal average of the stimuli preceding spikes—the *spike-triggered average* (STA). However, the nonlinearity of complex cells prevents analysis by reverse correlation, since these cells are driven equally by a given stimulus and an otherwise identical stimulus of opposite contrast polarity. These stimuli cancel when averaged, and the STA is therefore flat. But an analysis of the *spike-triggered covariance* (STC) can resolve mechanisms that have this type of symmetric nonlinear influence on response (de Ruyter van Steveninck and Bialek, 1988; Brenner et al., 2000; Simoncelli et al., 2004). STC analysis has successfully revealed some receptive field elements of cat V1 complex cells (Touryan et al., 2002) and macaque retinal ganglion cells (Schwartz et al., 2002).

Here we use a combination of STA and STC analysis to reveal unexpected structure within the receptive fields of both simple and complex cells in macaque V1. In particular, the responses of both cell types are best described using more linear filters than the standard models predict. Moreover, some of the filters contribute suppressive signals that appear to have a predominantly divisive influence on neuronal firing. These additional filters have a significant impact on responses, even to simple stimuli. These findings extend and enrich our understanding of the computations performed by V1 cells and place additional constraints on biologically based models. A brief report of some of this work has appeared elsewhere (Rust et al., 2004).

Results

Recovering the Linear Filters

We stimulated each neuron with a binary random bar stimulus whose bars covered a region slightly larger than the classical receptive field and were aligned with the preferred orientation (Figure 1D, left). We assume an “LNP” functional model consisting of a set of linear filters (L), an instantaneous nonlinearity (N) that combines filter outputs to yield a rate, and a Poisson spike generator (P) that converts the rate signal into spikes (Figure 1C). We assume that the filters operate within a 16 frame (160 ms) interval (Figure 1D, middle). The ensemble of these space-time stimulus intervals pre-

*Correspondence: rust@cns.nyu.edu

⁴These authors contributed equally to this work.

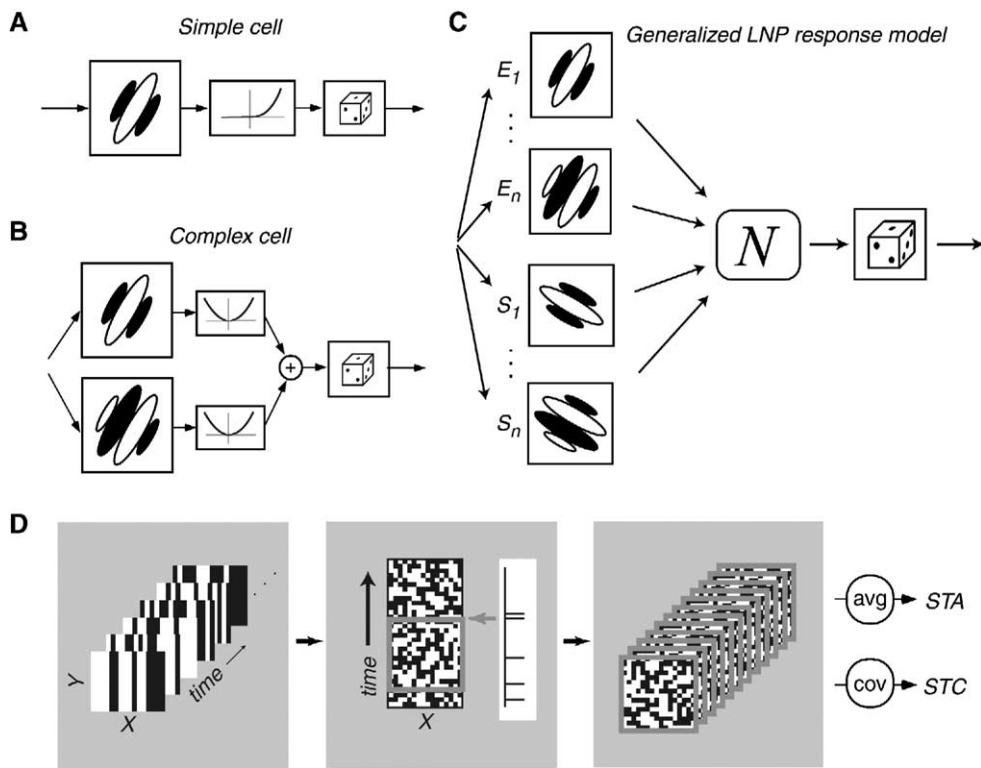


Figure 1. LNP Functional Models for V1 Neurons, and Their Characterization Using Spike-Triggered Analyses

(A) A standard simple cell model, based on a single space-time oriented filter. The stimulus is convolved with the filter, and the output is passed through a half-wave rectifying and squaring nonlinearity. This signal determines the instantaneous rate of a Poisson spike generator. (B) The “energy model” of a complex cell, based on a pair of space-time oriented filters with a quadrature (90°) phase relationship (Adelson and Bergen, 1985). Each filter is convolved with the stimulus, and the responses are squared and summed. The resulting signal drives a Poisson spike generator.

(C) The generalized linear-nonlinear-Poisson (LNP) response model used in this paper. The cell is described by a set of n linear filters (L), which can be excitatory (E) or suppressive (S). The model response is computed by first convolving each of the filters with the stimulus. An instantaneous nonlinearity (N) governs the combination of excitatory and suppressive signals that drives a Poisson spike generator (P).

(D) Spike-triggered analysis. (Left panel) A schematic of the random binary bar stimulus used to drive V1 neurons. The bars were aligned with the neuron’s preferred orientation, and the stimulus array spanned the classical receptive field. (Middle panel) A 2D representation of the stimulus sequence made by taking a slice through the stimulus ensemble in the plane orthogonal to the preferred orientation—each pixel represents the intensity of a bar at a particular location in one frame. The collection of stimulus blocks during the 16 frames (160 ms) before each spike (gray box) forms the spike-triggered stimulus distribution. (Right panel) The STA is a block of pixels, each corresponding to the average of the corresponding pixel values over the distribution. The STC is a matrix whose entries contain the average product of each pair of pixels after the mean has been projected out. See [Experimental Procedures](#) for details.

ceding each spike defines the spike-triggered stimulus ensemble (Figure 1D, right). We recovered a set of linear filters for each neuron by analyzing the statistics of this ensemble. The first such filter was estimated directly via the spike-triggered average, which is simply the ensemble mean. Additional filters were estimated by identifying axes in the stimulus space along which the variance of spike-associated stimuli differs significantly from that expected due to chance. Specifically, we calculated the covariance matrix of the spike-triggered stimulus ensemble and used standard techniques (principal components analysis) to identify stimulus components associated with increases or decreases in the variance of the spike-triggered ensemble relative to the variance of the raw stimuli (see [Experimental Procedures](#)). These filters define the fundamental stimulus selectivity of the cell (more formally, they determine a low-dimensional linear subspace of the stimulus in which the cell’s response is generated) (de Ruyter van

Steveninck and Bialek, 1988; Brenner et al., 2000; Schwartz et al., 2002; Touryan et al., 2002; Aguera y Arcas et al., 2003; Paninski, 2003; Simoncelli et al., 2004; Horowitz et al., 2005).

Note that receptive field components that are linearly combined, such as the excitatory and inhibitory signals arising from positive and negative subregions of a simple cell’s receptive field, are resolved into a single linear filter by this analysis. Linear components that are combined after a nonlinear operation such as rectification or squaring manifest themselves as multiple filters. The influence of each recovered filter on the neuron’s response can be excitatory or suppressive, depending on the way in which its output is incorporated into the nonlinear stage. It is important to realize that the individual filters are unique only up to a linear transformation (i.e., one can form an equivalent model based on an alternative set of filters that are related by an invertible linear transformation) and should not be taken as a literal rep-




Figure 2. Model Filters Recovered for an Example Cell Classified as Simple by Its Response Modulation of 1.51 to an Optimized Drifting Sinusoidal Grating

(A) Pairs of images representing the STA and three excitatory and four suppressive filters recovered from the STC analysis. The left-hand image of each pair represents the filter in the x - t plane represented by the middle panel of Figure 1D. The right-hand image is the filter's spatiotemporal frequency spectrum, in the ω_x - ω_t plane, where ω indicates frequency; the origin is at the center. For display, the contrast of each filter and its associated amplitude spectrum are scaled by the square root of its recovered weight (indicated next to each filter). Weights were independently normalized for the excitatory and suppressive pools, with the largest weight in each pool set to 1.

(B) Pooled excitatory (green) and suppressive (red) filter spatiotemporal envelopes computed as the L^2 -norm (square root of the weighted sum of squares) of the filter values for each x - t pixel. Regions of overlap are indicated by yellow. The temporal profiles of both signals at the level of the highest amplitude (eleventh) bar are plotted to the left.

(C) Pooled excitatory (green) and suppressive (red) frequency spectra as a weighted-sum of the ω_x - ω_t amplitude spectra for each filter. As in (B), regions of overlap are displayed in yellow.

resentation of underlying mechanisms. Nevertheless, the subspace they span is uniquely determined, and the filters in any functionally equivalent model must span the same subspace.

The spatiotemporal filters for a representative simple cell, along with their spatiotemporal frequency spectra, are shown in Figure 2A. The space-time tilt of the STA, or equivalently the concentration of its spatiotemporal

energy in diagonally opposite quadrants of the spectrum, indicates a preference for the direction of a moving stimulus. Were this simple cell adequately described by a single linear filter (as in the standard model of Figure 1A), our analysis would reveal no other filters. However, the STC analysis revealed three additional excitatory filters, each with the same direction preference as the STA, and four suppressive filters, each tuned for the direction opposite that of the excitatory filters. To examine the structure of the excitatory and suppressive elements of the model, we computed separate spatiotemporal and spectral envelopes for the pooled excitatory and suppressive filters by taking the square root of the sum of the squared filters and their spectra (Figures 2B and 2C). The pooled excitatory and suppressive signals overlap completely in space and time; the time courses of excitation and suppression are shown by the traces to the left of Figure 2B. In the frequency domain, the excitatory and suppressive spectra are largely nonoverlapping and concentrated in opposite quadrants, indicating their selectivity for opposite directions of motion.

Now consider a typical complex cell (Figure 3). The energy model (Figure 1B) predicts a flat STA and exactly two significant STC filters. The analysis produced a weak STA and two strong excitatory filters, along with an additional five excitatory and seven suppressive filters. As for the simple cell, all excitatory filters had the same direction preference, and most suppressive filters had the opposite direction preference; the sixth suppressive filter—a relatively weak one—began with the same direction preference as the other suppressive filters, but switched its direction preference to that of the excitation at longer time delays. The six strongest excitatory and suppressive filters for the complex cell appear in pairs, with each member of a pair appearing as a spatial phase-shifted replica of the other. The pooled spatiotemporal excitatory and suppressive envelopes indicate that the suppression was time delayed relative to the excitation for this neuron (indicated by the green band below the yellow core of the receptive field in Figure 3B, and the adjacent time course traces). The excitatory and suppressive frequency spectra were almost completely nonoverlapping (Figure 3C), as was the case for the simple cell (Figure 2C).

We have illustrated the recovered filters for two direction-selective neurons. For nondirectional neurons, the excitatory filters generally (though not invariably) did not have discernible space-time tilt. Suppressive filters, when found, were also typically separable, but they often spanned a different region of the spatiotemporal frequency spectrum than the excitatory filters.

Although the standard model for a simple cell contains only a single linear filter, we recovered between one and three additional excitatory filters for each of the simple cells in our population (Figure 4A, dark gray). Similarly, the energy model predicts two excitatory filters for complex cells, but we recovered additional excitatory filters in many complex cells (Figure 4A, light gray). We also found significant suppressive filters in many cells (41/59). Suppressive filters were more numerous in direction-selective complex cells than in other cells, but it might be that non-direction-selective cells would be more powerfully suppressed by orthogo-

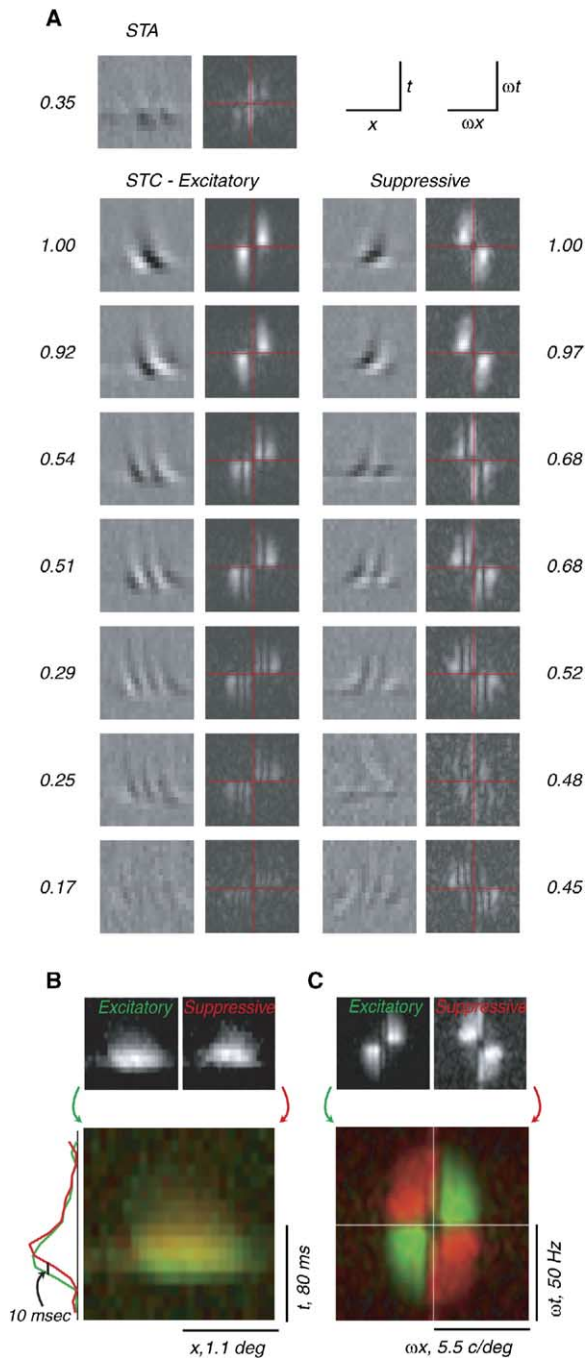


Figure 3. Model Filters Recovered for an Example Cell Classified as Complex by a Relative Modulation of 0.1 to an Optimized Drifting Sinusoidal Grating

(A) The STA, seven excitatory, and seven suppressive filters recovered from the STC shown with the same conventions as in Figure 2A. The recovered weights for each filter are indicated.

(B) Pooled excitatory (green) and suppressive (red) filters, represented as in Figure 2B. The temporal profiles of both signals at the level of the twelfth (middle) bar are plotted to illustrate the delay of suppression relative to excitation.

(C) Pooled excitatory (green) and suppressive (red) frequency spectra represented as in Figure 2C.

nal orientations not present in our stimulus ensemble. (Figure 4B).

It is important to understand that both the quality and number of filters recovered by our analysis depend not only on the strength of their influence on neural response, but also on the number of spikes collected (Aguera y Arcas and Fairhall, 2003; Paninski, 2003), a fact that we verified experimentally (see Figure S1 in the Supplemental Data available online). The values in Figure 4 are therefore a lower bound on the number of filters required to characterize the response of these neurons accurately. In our population analysis, we only included cells for which we collected at least 50 spikes per spatiotemporal dimension (mean 229 spikes per dimension, 55,000 total spikes).

Recovering the Nonlinearity

After recovering a set of linear filters, the model is completed by estimating the nonlinear function that combines the outputs of these filters to produce a firing rate. When the number of filters is small (one or two), this can be done directly by computing the filter responses to the stimulus sequence and evaluating the average number of spikes observed for each binned response combination. Figures 5A–5C show firing rate functions for filters selected from the example simple and complex cells of Figures 2 and 3, indicated by the red curves along the diagonal margins of each plot. The rate functions associated with the STA were always half-wave rectified (Figure 5A, right hand axis), and functions for the STC filters were always symmetric (Figure 5A, left axis; Figures 5B and 5C, both axes). Excitatory STC filters (those corresponding to PCA-derived axes with increased variance) had firing rate functions that increased monotonically with the magnitude of their outputs (Figures 5A and 5B), and suppressive filters (those corresponding to axes with decreased variance) produced firing rate functions that decreased monotonically with the magnitude of their outputs (Figure 5C).

Important regularities emerged when we examined firing rate as a function of the activity of *pairs* of excitatory or suppressive filters; Figures 5A–5C show examples of these joint nonlinearities as grayscale images, where the lightness of each pixel corresponds to the firing rate associated with the associated pair of filter outputs; the axes drawn across the images cross at the pixel for which both filter outputs were 0. These 2D nonlinearities have a characteristic form: contours of constant firing rate are well fit by ellipses with principal axes aligned with the coordinate system, suggesting that the firing rate can be expressed as a function of the weighted sum of squares of the filter responses (Figures 5B and 5C). When the STA is combined with another excitatory dimension, the contours outline a crescent, as would be expected if the STA response were half-wave rectified prior to squaring (Figure 5A).

It is not feasible to directly estimate the complete nonlinear function that maps the filter outputs into a firing rate, since the number of response combinations grows exponentially with the number of filters. Instead, we exploited the regularity in the pairwise nonlinear response functions to reduce the dimensionality of the




Figure 4. Relationship between Directional Selectivity and the Number of Filters Revealed by STC, Excluding the STA

The population included 18 simple (dark gray) and 41 complex cells (light gray). Directionality was determined by the response of each cell to an optimized drifting grating. A direction index of 0 corresponds to a cell that is equally responsive to gratings drifting in both directions along its preferred axis; cells with direction indices that exceed 1 are suppressed below baseline by a grating drifting opposite the preferred direction. (A) The number of excitatory filters was larger on average for complex than for simple cells, but did not depend on the cell's direction selectivity. (B) Suppressive filters were more often manifest in complex than in simple cells. Directionally selective complex cells in particular had many suppressive filters. This analysis only includes cells for which we collected at least 50 spikes per spatiotemporal dimension.

problem. Specifically, we combined the outputs of excitatory and suppressive filters separately, each by a root sum of weighted squares, to form unified excitatory and suppressive signals (\mathcal{E} and \mathcal{S} in Figure 5D). The STA response is half-squared and included in the excitatory pool. We obtained the weight associated with each filter by maximizing the mutual information between the joint excitatory and suppressive signals and the spikes (see Experimental Procedures). Space-time and spectral representations of these pools for the example cells can be seen in Figures 2B and 2C and Figures 3B and 3C. We completed the reconstruction of the nonlinearity with a second stage that combines the excitatory and suppressive pooled signals to produce a firing rate response. The estimated firing rates for each pair of binned \mathcal{E} and \mathcal{S} values are shown in Figure 6A.

Properties of the Suppressive Signal

For those cells in which we found suppressive filters, we quantified the strength of the suppression as the fractional change in response to the strongest excitatory stimulus with and without suppression (Figure 6A). Suppression in many cases was highly effective—for the example complex cell shown in Figure 3, a suppressive stimulus reduced response by 49%. On average, the fractional suppression was 63%.

To examine how excitation and suppression combine, we fit a model to the two-dimensional nonlinearity that describes firing rate as a combination of the pooled \mathcal{E} and \mathcal{S} signals (Figure 6A). We fit the 289 data points (17 equally populated bins along each axis) with a sigmoid excitatory function that was suppressed through both subtractive and divisive terms (see Experimental Procedures). Figure 6A shows the data and model fits plotted as slices of increasing excitation at different levels of suppression for the cell in Figure 3. The subtractive component of suppression causes the

relatively small reduction in response at the lowest values of excitation (left end of each plot). The divisive component captures the change in slope and position along the abscissa that can be seen by comparing curve shapes as the value of suppression increases. In this case, as for almost all of our cells, the divisive component of suppression was numerically more important, but the model fits were significantly worse if we omitted either the divisive or the subtractive term. This model provided a consistently good fit to all our data, accounting for most of its variance (Figure 6C).

Predicting Neuronal Responses

We used a “playback” approach to verify the existence and influence of the unexpected additional filters that emerged from our analysis (Touryan et al., 2002). For a subset of complex cells, we presented stimuli tailored to discriminate the standard energy model from the full LNP model revealed by STC; for each cell, the stimuli were movies of the actual filters derived from the STC analysis. Figure 7A shows the response of an example complex cell to a movie of its first STC component. Both the energy model and the full model accurately predicted the response of the cell. Figure 7B shows the response to the sixth STC component, for which the two models make drastically different predictions: the energy model predicts only a weak response, but a robust response was evoked, as predicted by the full STC model.

We used a related method to verify the influence of suppressive filters. We created compound stimulus movies that were time-aligned mixtures of excitatory and suppressive STC components (Figure 7C). We measured contrast-response for the first excitatory STC component and observed that adding the first suppressive component reduced the response by shifting the contrast-response function down and/or to the right. This behavior was not predicted by a model con-




Figure 5. Estimating the Nonlinearity

(A–C) Firing rate as a function of the output of single filters and as the joint output of two filters. Firing rate is indicated by pixel intensity; overlaid red curves indicate contours of constant firing rate. The graphs along the diagonal margins of each plot show the 1D firing rate functions for each filter, represented as the average response as a function of the projection of the stimulus onto the filter.

taining only the excitatory filters (Figure 7D) or by the energy model (data not shown), but was well described by the full model containing both excitation and suppression (Figure 7E).

Subunit Structure of Complex Receptive Fields

In complex cells, the STC analysis consistently recovered more than just the filter pair predicted by the energy model. Figure 8A shows the filters recovered by STC for an example complex cell. In these cells we observed a specific structural relationship between the filter pairs: the temporal envelopes for all the filters revealed by STC were similar, but the spatial envelopes differed across filter pairs. Figure 8B compares the overall spatiotemporal envelope of the receptive field (shaded) to envelopes for the first two excitatory filters (red) and the last four (green). Shown is a slice across the spatial envelope at the peak temporal offset, 65 ms. The two strongest STC filters were confined to the center of the receptive field, while the additional filters were weak in the middle but robust at the receptive field edge. Mindful that the underlying biological mechanisms of a cell can be linear combinations of the filters revealed by STC, we wondered what types of subunits could produce such results. Simulations confirmed that a model of a complex cell that included five pairs of spatially shifted subunits (Figures 8C and 8D) produced responses and STC filters that are strikingly similar to filters we found in most complex cells (Figures 8E and 8F).

Discussion

Previous reconstructions of spatiotemporal receptive fields for V1 cells have been specific with regard to cell type and have focused primarily on linear or quasilinear descriptions of neural response (Movshon et al., 1978b; Movshon et al., 1978a; Emerson et al., 1987; Jones and Palmer, 1987; McLean and Palmer, 1989; DeAngelis et al., 1993; Lau et al., 2002; Touryan et al., 2002; Livingstone and Conway, 2003). We have explored a more general nonlinear model for V1 neurons that can be applied to cells of any type. The model includes a linear stage incorporating a small set of filters, followed by a nonlinear function that combines the filter outputs to generate a firing rate. The model can be fit to data using spike-triggered techniques and used to predict

Dashed lines indicate the mean response to all stimuli. The 1D rate functions appear to be either half-squared (like the STA [E0] projection in [A]) or full-squared (like the E1 projections in [A] and [B]). 1D and 2D firing rate functions are plotted for (A) the STA [E0] and strongest excitatory filter revealed by STC for the simple cell example of Figure 2, (B) the two strongest excitatory filters revealed by STC for the complex cell example of Figure 3, and (C) the two strongest suppressive filters revealed by STC for the complex cell example of Figure 3.

(D) The elliptical symmetry of the 2D firing rate functions allows for a two-stage representation of the nonlinearity. In the first stage, the output of the excitatory and suppressive signals are pooled separately, each via a weighted sum of squares. In the second stage, a two-dimensional function governs the combination of the excitatory and suppressive signals to produce a firing rate.

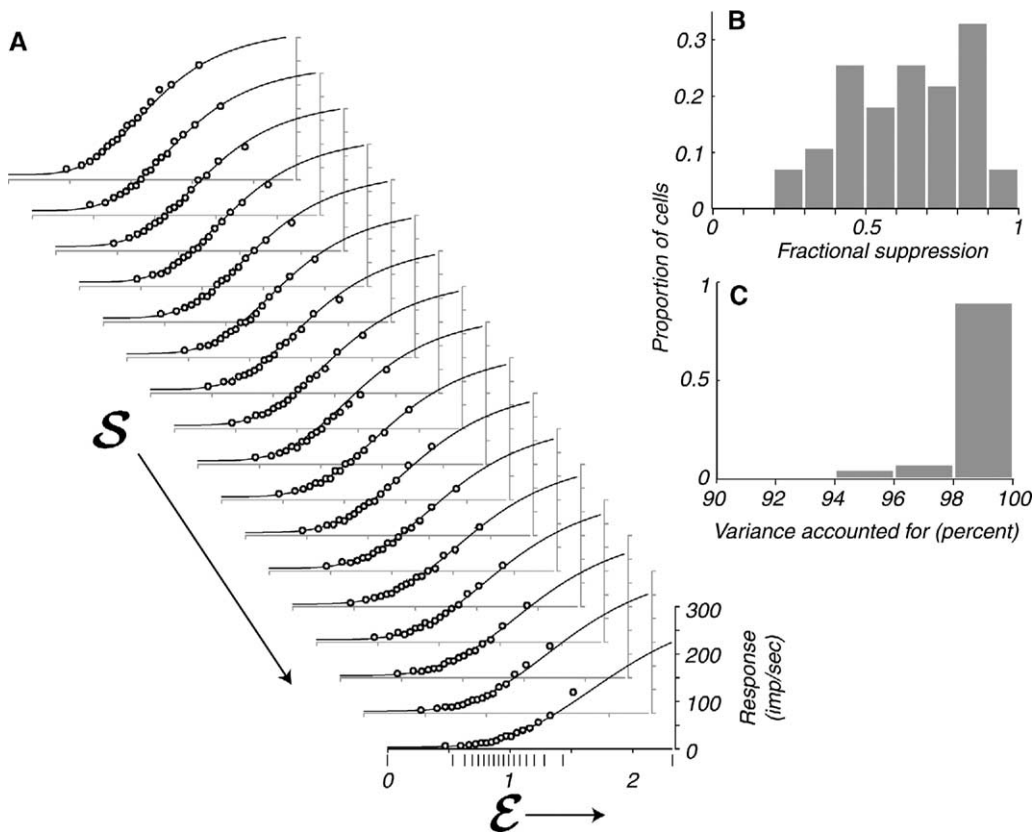


Figure 6. Properties of the Suppressive Signal

(A) Slices through $N(E, S)$, the firing rate function for the joint output of the excitatory and suppressive pooled signals for the complex cell in Figure 3. Each graph shows the influence of increasing excitation at a particular suppression value. Suppression is minimal for the uppermost graph and increases toward the bottom. Smooth curves show the best fit of the combined model of excitation and suppression described in the text. The data are unevenly distributed along the abscissa because they are unequally binned—the bin boundaries are indicated by the vertical strokes below the abscissa of the bottommost graph in (A) and the data point corresponds to the centroid of data in each bin. The reduction in response at the right-hand point (i.e., the largest value of E) in the topmost to the bottommost graph is used to measure fractional suppression, defined as $1 - (\text{maximal excitation with maximal suppression} / \text{maximal excitation with minimal suppression})$, in this case 0.49. (B) The distribution of fractional suppression for the 16 directional and 24 nondirectional cells with significant suppressive filters; the remaining 18 cells are omitted. (C) Distribution of the fraction of the variance in the data accounted for by the fits to the model for the same cells.

responses to arbitrary stimuli. Most macaque V1 cells required substantially more filters than predicted by standard models of these neurons. For simple cells, we always recovered more than the single excitatory linear filter predicted by the standard model of these cells (Figure 1A). For complex cells, in many cases we recovered more than the two excitatory filters predicted by the energy model (Figure 1B). More than half the cells of each type had one or more significant suppressive filters whose existence is also not predicted by standard models. Touryan et al. (2002) used a related method to characterize complex cells in cat area 17. They found only two excitatory filters for most cells, and no suppressive filters. While it is possible that cat and monkey cortex use different computational strategies, the likely reason for the difference in the number of filters found in their study and ours is that we recorded more spikes from our neurons; in the Supplemental Data we present an analysis showing how importantly the number of spikes influences the number of filters revealed by an STC analysis.

Before discussing the basis for the unexpected extra filters, we consider whether they might arise artifactually. Small involuntary eye movements are inescapable in preparations like ours (Forte et al., 2002). We have estimated the eye movements that occurred during data acquisition by analyzing short segments of data—the movements are small and do not produce artifactual filters in simulation. We also simulated the effects of published eye movement records from a very similar preparation (Forte et al., 2002) on our STC results. We conclude that the number of artifactual filters produced by such movements is small and inconsistent with our results (details of this analysis may be found with Figure S2).

We also wondered whether the unexpected filters arose from deviations from the Poisson spiking assumed by the model. Refractoriness, which amounts to a brief suppression after every spike, is known to produce suppressive filters in a Poisson-based STC analysis (Schwartz et al., 2002; Aguera y Arcas and Fairhall, 2003). Similarly, intrinsic bursting spike patterns could

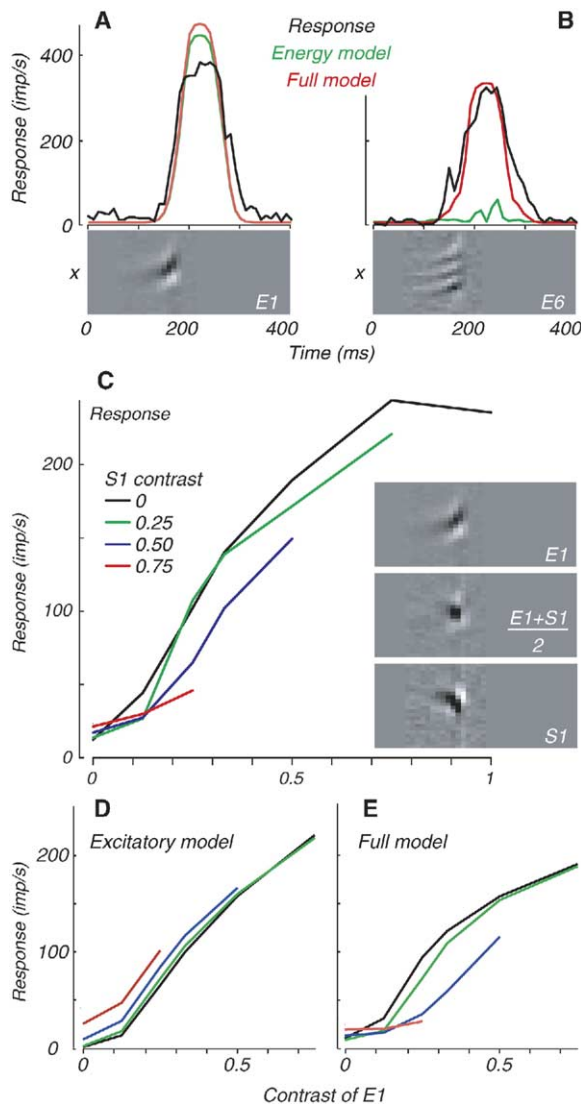


Figure 7. Experimental Validation of the Unexpected Filters for the Complex Cell Shown in Figure 3

(A and B) Validation of the influence of unexpected excitatory filters. Shown are the stimuli presented to the cell as x - t plots, the response of the cell (black), the predicted response of the energy model (the STA and first two filters, red) and the prediction of the full model including all excitatory and suppressive filters (green). (A) The stimulus presented was a movie of the strongest excitatory filter preceded and followed by periods of gray. Both models predict a similar response to this stimulus. (B) The energy model and the full model predict very different responses to a movie of the sixth excitatory filter. The full model better predicts the actual response of the cell.

(C–E) Validation of the suppressive filters. Stimuli were weighted combinations of the strongest excitatory (E1) and suppressive (S1) filter, each scaled such that the absolute value of the largest intensity pixel had a value of 1 (defined as a contrast of 1). The inset in (C) shows x - t plots of these filters (top and bottom) and a combination of the two at equal magnitude (middle). (C) shows the mean response to each stimulus as a function of the contrast of the E1 movie, with line colors indicating the contrasts of the added S1 movie. (D and E) Predicted responses of the energy model and of the full STC model to the same set of stimuli. Model predictions for all panels share a common contrast gain adjustment as described in the Experimental Procedures.

create artifactual excitatory subunits. A key feature of either artifact is that it appears as a time-delayed replica or a temporal derivative of the most potent excitatory filters, a pattern that we almost never observed. Almost all our filters have similar time courses, but their spatial structures are distinct from one another (Figures 2, 3, and 8). Moreover, in direction-selective cells, suppressive filters have the opposite x - t slope to excitatory ones (Figures 2 and 3). In rare cases, such as the sixth suppressive filter for the complex example of Figure 3, we found spatiotemporal structure that partly resembled time-delayed replicas of strong excitatory filters; we never saw such replicas among the excitatory filters.

What neural mechanisms produce the filters that we have found? Our model includes all feedforward, feedback, and collateral processes affecting a V1 neuron. Even in retina, a similar technique can resolve multiple linear filters in ganglion cells (J.W. Pillow, E.P. Simoncelli, and E.J. Chichilnisky, personal communication). The filters we recovered in V1, however, differ in spatiotemporal properties from these retinal subunits, and we therefore believe that they are of cortical origin and do not simply reproduce nonlinear processing of signals in the retina. A likely source of multiple filters is the convergence of signals of intracortical origin. The large receptive fields of complex cells are thought to arise from such a convergence, creating multiple spatially distributed subunits (Hubel and Wiesel, 1962; Hubel and Wiesel, 1968; Movshon et al., 1978a; Cavanaugh et al., 2002a). Our analysis shows that this kind of distributed subunit structure can give rise to filters very much like the ones we routinely observed (Figure 8). This analysis also highlights the importance of distinguishing the structure of the filters revealed with STC from the neural machinery from which they arise (Figure 8F versus Figure 8C).

Finding extra filters in simple cells is perhaps even more unexpected than in complex cells. The standard feedforward quasilinear model of these cells has been accepted for some time (Movshon et al., 1978b; De Valois et al., 1982; Carandini et al., 1997), and the revision needed to account for the extra filters is significant. As for complex cells, the most attractive possibility is that these filters arise from intracortical circuits. As proposed by Hubel and Wiesel, the core of simple cell receptive fields seems to be related to the spatial arrangement of inputs from the LGN (Alonso et al., 2001); these inputs seem likely to reveal themselves primarily in the STA. The additional filters we find in simple cells are “complex-like,” in that they are full-rectified and would therefore contribute nonlinearly to simple cell activity. Alonso et al. (2001) found that the LGN inputs to simple cells are concentrated in the center of the receptive field, so it may be that the nonlinear inputs we have found form a phase-insensitive “fringe” of excitation of intracortical origin that complements the direct feedforward input from the LGN. This is also consistent with the idea that simple and complex cells may form a continuum (Chance et al., 1999; Mechler and Ringach, 2002; Priebe et al., 2004). In our analysis, the phase sensitivity (or “simplicity”) of a simple cell is determined by the relative weight of its STA and STC filters in determining responses—cells with strongly

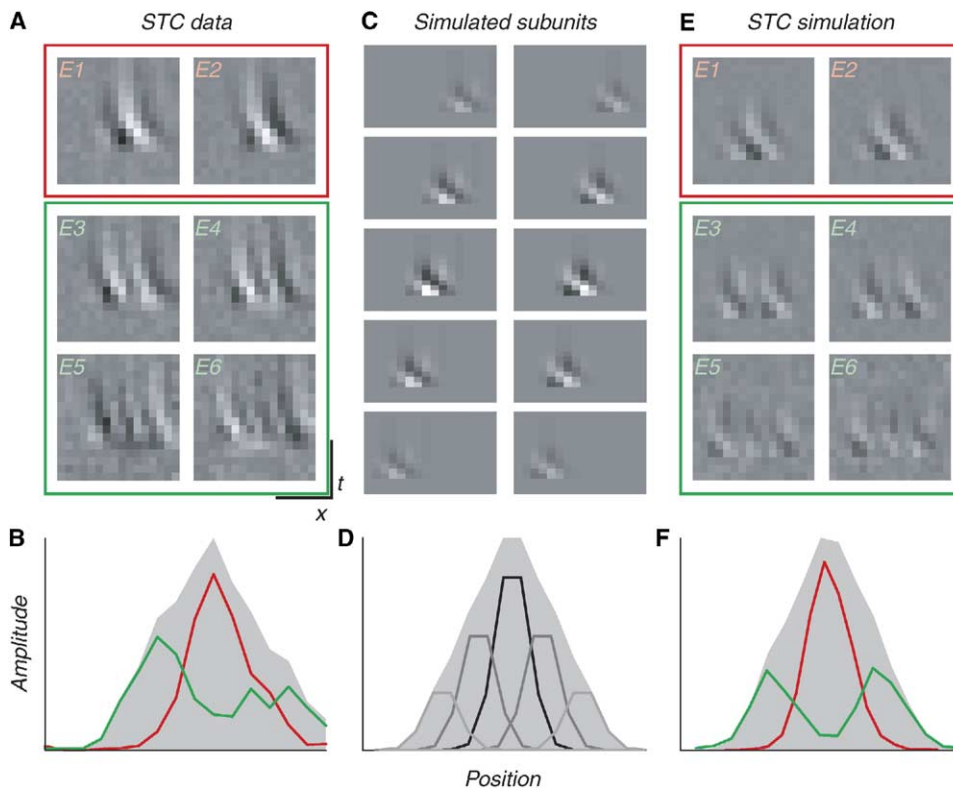


Figure 8. Relationship between STC Filters and Complex Cell Subunits

(A) the six excitatory filters revealed by STC for an example complex cell (the weights of the filters were 1, 0.98, 0.74, 0.48, 0.68, and 0.45). Note the similarity to the example cell of Figure 3.

(B) a cross-section of the spatial profile of the receptive field, computed by taking the square root of the weighted summed squared combination of the pixel values for each of the filters. The cross-section was taken at the peak temporal offset, $t = 65$ ms before a spike. The spatial profile is shown for all the filters (gray), the two strongest filters (red), and the remaining four filters (green).

(C) We used the ten subunits (five quadrature pairs) illustrated to create a simulated complex cell and analyzed its responses using STC.

(D) The sensitivity profiles of the five pairs of subunits and of their envelopes, computed as in (B).

(E) The filters revealed by using STC analysis on the responses of the simulated complex cell. Note the similarity between the STC filters for this simulated cell and the real data in (A).

(F) The spatial profile of the STC filters of the model and their envelope, computed as in (B). Simulated data were generated by convolving each of the subunits with the random binary bar stimulus used in the experiment and then combining the signals via a weighted sum of squares with weights of 0.33, 0.66, 1, 0.66, and 0.33 applied to the filter pairs, as shown from top to bottom in (B).

dominant STA components will be “classic” simple cells, while those with more balanced STA and STC weights will have intermediate properties.

In many cells we found suppressive filters. Suppression is of course a well-known feature of cortical circuits, but many commonly postulated forms of suppression would not manifest themselves as suppressive STC filters of the kind we observed. In particular, the “inhibitory side lobes” of linear receptive fields would be absorbed into the STA as negatively weighted regions and would not appear as separate filters. The suppressive signals had a predominantly divisive impact on neuronal response (Figure 6) and are thus broadly in agreement with models that postulate gain control by divisive mechanisms (Geisler and Albrecht, 1992; Heeger, 1992b; Carandini et al., 1997). However, the data are not consistent with the specific prediction of these models that the suppressive signal should be untuned for stimulus features. In particular, the prevalence and strength of suppression in the nonpreferred direction in

direction-selective cells (e.g., Figures 2 and 3) suggests that it does more than set gain, and contributes directly to stimulus selectivity. One mechanism that might contribute in that way is the one usually assigned to the receptive field surround, which is believed to overlap with the classical receptive field and is known to have selective tuning for a number of stimulus features (Knierim and van Essen, 1992; Levitt and Lund, 1997; Cavanaugh et al., 2002b). We conclude that our suppressive filters represent a mechanism that enhances neuronal selectivity by reducing response gain for non-optimal stimuli.

We have presented evidence for profound deviations from standard models of V1 neurons based on a general functional model estimated with spike-triggered techniques. Additional important properties of these neurons might be revealed by extending the analysis to cover the second spatial dimension as well as other stimulus attributes like color, binocularity, and center-surround organization. These extensions would be re-

quired to build a model that could predict the response to any arbitrary stimulus. But even when confined to examining a limited number of stimulus attributes, spike-triggered techniques are sensitive tools for revealing previously unsuspected aspects of cortical computation.

Experimental Procedures

Electrophysiology

We recorded from isolated single units in primary visual cortex (V1) of adult macaque male monkeys (*Macaca fascicularis* and *Macaca nemestrina*). After surgical preparation (see Cavanaugh et al., 2002a, for details), anesthesia was maintained throughout the experiment with continuous intravenous administration of 4–12 μ g/kg/hr of sufentanil citrate. During experiments, the animal was artificially respired and body temperature was maintained with a heating pad. Vital signs (heart rate, lung pressure, EEG, ECG, body temperature, and end-tidal CO_2) were monitored continuously. Vecuronium bromide (Norcuron, Organon) was administered intravenously (0.15 mg/kg/hr) to prevent involuntary slow drifts of the eyes, and gas-permeable contact lenses were used to protect the corneas throughout the experiment. Supplementary lenses chosen by direct ophthalmoscopy were used to make the retinas conjugate with a screen 165–180 cm distant, and refractive correction was checked by adjusting the lens power to maximize the resolution of recorded units. All experiments were performed in compliance with the National Institutes of Health Guide for the Care and Use of Laboratory Animals and within the guidelines of the New York University Animal Welfare Committee.

Neurons were located in the operculum and in the calcarine sulcus with receptive fields centered between 3° and 20° from the fovea. Single-unit activity was recorded using quartz-glass microelectrodes (Thomas Recordings, Giessen, Germany). Signals were amplified, band-pass filtered, and fed into a time-amplitude window discriminator. Spike arrival times and stimulus synchronization pulses were stored with a resolution of 0.25 ms.

At the end of the experiment, animals were killed with an overdose of sodium pentobarbital (>60 mg/kg) and perfused with 4% paraformaldehyde. Identification of the recording location was made through histological identification of electrolytic lesions made at suitable locations along the electrode tracks by passing 1–2 μ A of current for 2–5 s through the tip of the electrode.

Stimuli

Stimuli were generated by a Silicon Graphics Octane 2 workstation and presented on a gamma-corrected monitor with a refresh rate of 100 Hz and a mean luminance of 33 cd/m². All stimuli were presented monocularly to the cell's preferred eye.

Upon encountering a cell, the initial characterization involved a determination of the best direction, spatial frequency, and temporal frequency of drifting grating stimuli. The size of the classical receptive field was defined as the size at which an optimized full-contrast sinusoidal grating saturated the response without impinging upon the suppressive surround. Stimuli used in the spike-triggered characterization were extended temporal sequences in which each frame contained a set of bars randomly assigned as black or white (Figure 1D). The orientation of the bars was aligned with the cell's preferred orientation, and the stimulus array was confined to the classical receptive field. The number of bars (8–32) was chosen such that at least five bars fell within a 75% contour drawn around the most active region in the receptive field. A new frame was displayed every 10 ms, and the stimulus was presented for a total duration of 15–80 min, depending on the responsiveness of the cell. We set a minimal criterion of 50 spikes per spatiotemporal dimension for our population. For a few cells with extremely low firing rates, characterization required an unreasonable amount of time for data collection, and these cells were abandoned.

Recovering Linear Filters with Spike-Triggered Analysis

We define the spike-triggered stimulus block, $S_n(x, t)$, as the set of bar intensities, relative to the mean stimulus intensity, in the 16

frames preceding the n th spike (Figure 1D, middle panel). Each of these spike-triggered stimuli corresponds to a D -dimensional vector, indexed by the parameter pair (x, t) , where D lies between 64 and 512 (16 times the number of bars in the stimulus). The spike-triggered average is computed as

$$A(x, t) = \frac{1}{N} \sum_n S_n(x, t),$$

where N indicates the number of spikes. If one assumes that the neural response is generated by convolution with a single linear filter followed by an instantaneous asymmetric (e.g., half-squaring) nonlinearity and Poisson spike generation, the STA provides an unbiased estimate of the linear filter (de Boer and Kuyper, 1968; Chichilnisky, 2001; Paninski, 2003). But the STA will fail to produce a description of receptive field properties if the nonlinearity is symmetric (e.g., squaring) or if the response depends on more than a single axis within the stimulus space. In order to handle these situations, one can examine higher-order statistical properties of the spike-triggered stimulus ensemble. A number of authors have examined the second-order statistics of spike-triggered stimuli by computing the spike-triggered covariance (STC) (de Ruyter van Steveninck and Bialek, 1988; Brenner et al., 2000; Schwartz et al., 2002; Touryan et al., 2002; Aguera y Arcas et al., 2003; Horwitz et al., 2005).

Before computing the STC, we first eliminate the component corresponding to A from each member of the stimulus ensemble. Specifically, we compute the normalized (unit vector) STA, \hat{A} , and define

$$S'_n(x, t) = S_n(x, t) - \left[\sum_{x, t} S_n(x, t) \cdot \hat{A}(x, t) \right] \cdot \hat{A}(x, t)$$

The spike-triggered covariance matrix is then computed as

$$C(x_1, t_1, x_2, t_2) = \frac{1}{N-1} \sum_n S'_n(x_1, t_1) S'_n(x_2, t_2).$$

This differs from the traditional calculation of covariance [in which $A(x, t)$ would be subtracted from each $S_n(x, t)$], but ensures that the axes obtained in the STC analysis will be orthogonal to the STA, thereby avoiding unwanted interactions between the STA and STC analyses. We also found that it greatly simplified the description of the subsequent nonlinear portion of the model. The matrix C , with the parameter pairs $\{x_1, t_1\}$ and $\{x_2, t_2\}$ specifying the row and column indices, fully represents the variance of the spike-triggered stimulus ensemble in all possible directions within the D -dimensional stimulus space. Geometrically, the surface swept out by a vector whose length is equal to the variance along its direction is a hyperellipse, and the principal axes of this hyperellipse, along with the variance along each axis, may be recovered using principal components analysis. Specifically, the principal axes of this ellipse correspond to the eigenvectors of the covariance matrix, and the variance along each of these axes is equal to the corresponding eigenvalue.

In the absence of any relationship between the stimulus and the spikes (and in the limit of infinite data), the spike-triggered ensemble would be a randomly selected subset of all stimuli, and the variance of this subset in any direction would be identical to that of the full stimulus set. In an experimental setting, the finiteness of the spike-triggered ensemble produces random fluctuations in the variances in different directions. We are interested in recovering those axes of the stimulus space along which the neuron's response leads to an increase or decrease in the variance of the spike-triggered ensemble that is greater than what is expected from this random fluctuation due to finite sampling.

We tested a nested sequence of hypotheses to determine the number and identity of axes corresponding to significant increases or decreases in variance. We began by assuming that there were no such axes (i.e., that the neuron's response was independent of the stimulus). We used a Monte Carlo simulation to compute the distribution of minimal and maximal variances under this hypothesis. Specifically, we randomly time-shifted the spike train relative to the stimulus sequence, performed our STA/STC analysis on the resulting spike-triggered stimulus ensemble, and extracted the

minimum and maximum eigenvalues. Based on 500 such calculations, we estimated the 99% confidence interval for both the largest and smallest eigenvalues. We then asked whether the eigenvalues obtained from the true spike-triggered ensemble lay within this interval. If so, we concluded that the hypothesis was correct. Otherwise, we assumed the largest outlier (either the smallest or largest eigenvalue) had a corresponding axis (eigenvector) with a significant influence on neural response. We added this axis to a list of significant axes and proceeded to test the hypothesis that all remaining axes were insignificant.

The STC analysis can provide an unbiased estimate of the filters in an LNP model, but this is only guaranteed when the raw stimulus ensemble is Gaussian distributed (Paninski, 2003; Bialek and de Ruyter van Steveninck, 2005). But the low contrast of Gaussian stimuli leads to low firing rates in V1 neurons, which produces noisy estimates of the model components, as well as a reduction in the number of significant filters. We therefore chose to use higher-contrast binary stimuli. For neurons with many excitatory subunits, we occasionally found suppressive filters that differed from the others in that they contained only a few isolated nonzero bars. Simulations confirm that such filters can arise as artifacts when binary stimuli are used. The problem arises specifically because the distribution of raw binary stimuli tapers (i.e., the variance decreases) as one moves in particular directions away from the origin in the stimulus space. When the excitatory filters lie near these special directions, the spike-triggered ensemble will necessarily have reduced variance (compared with the full ensemble of raw stimuli), and spurious suppressive filters will result (Paninski, 2003).

This problem can largely be corrected by adjusting the values of the raw stimulus components such that the variances in all directions are approximately equivalent (conditional whitening) before estimating the suppressive STC components. Specifically, we computed the STA and excitatory STC filters as described above. For each raw stimulus block, we computed the response of these filters and combined them to generate a single pooled excitatory response value (see “Estimating the Nonlinearity” below for details). We binned the raw stimuli into ten equal-sized subsets according to their pooled excitatory response value and computed the covariance matrix for each subset (C_n for the n th subset). The stimuli in each subset were then whitened by multiplying by

$$E_e E_e^T + E_o E_n D_n^{-1/2} E_n^T E_o^T$$

where E_e is a matrix containing the (orthogonal) excitatory filters, E_o contains an orthogonal basis for the remainder of the stimulus space, and E_n and D_n are the eigenvectors and eigenvalues of the conditional covariance matrix, C_n . After this conditional whitening, the stimuli were recombined and PCA analysis was applied to the spike-triggered set in order to re-estimate the filters.

To independently verify that our procedure was not subject to significant artifacts due to binary stimulation, we recorded responses to both binary and Gaussian white noise stimuli for six cells that responded with sufficient vigor to the low-contrast Gaussian stimuli. The number of filters recovered by STC analysis was similar for the two stimulus types, and the filters were statistically equivalent (Figure S3).

Estimating the Nonlinearity

For individual filters, the nonlinear function that maps filter output to firing rate is estimated directly by taking the ratio of the number of spikes to the number of stimuli for each (binned) filter output value (Figures 5A–5C, diagonal marginal graphs). Similarly, firing rate as a function of the responses of two filters is estimated by taking the ratio of the joint (two-dimensional) counts of the number of spikes to the number of stimuli (Figures 5A–5C). As mentioned in Results, this direct method is not feasible for estimating the nonlinear function that combines all filter outputs. In this case, we defined the firing rate nonlinearity in two stages. First, excitatory and suppressive filter outputs were combined into pooled excitatory and suppressive signals, respectively, by taking the square root of a weighted sum of their squares (with the STA half-squared). Weights for each pool were obtained by maximizing the mutual information between the two pooled signals and the spikes. The use of mutual information as an optimization criterion is advanta-

geous, as it makes no assumptions regarding the form of interaction between the excitatory and suppressive signals. Because the weights applied to each pool are individually normalized, the relative strengths of excitation and suppression cannot be determined.

The second stage of the nonlinearity combines the pooled excitatory and suppressive signals to generate a firing rate and was estimated directly by taking the ratio of the two-dimensional binned counts of the number of spikes to the total number of stimuli presented. Because the data were not uniformly distributed across this two-dimensional space, we adjusted individual bin sizes along each axis to maintain a uniform distribution of data points across each of the marginals—an example can be seen on the abscissa of Figure 6A.

Parametric Model for the Interaction of Excitatory and Suppressive Signals

We fit a simple parametric model to the binned second-stage nonlinear firing rate function. For each bin, we assumed a value of excitation (\mathcal{E}) and suppression (\mathcal{S}) equal to the center of mass of the data in that bin. We fit the model with a function that combined a standard sigmoidal Naka-Rushton excitatory term (exponent ζ , and coefficients α , β , and γ), modified by both subtractive and divisive suppressive terms (coefficients δ and ϵ):

$$R = \alpha + \frac{\beta \mathcal{E}^\zeta - \delta \mathcal{S}^\zeta}{\gamma \mathcal{E}^\zeta + \epsilon \mathcal{S}^\zeta + 1} \quad (1)$$

Models lacking one of these suppressive terms provided fits that were significantly worse for the 41 cells with significant suppressive axes. We used STEPIT (Chandler, 1969) to minimize the mean-squared error between the measured firing rates and model predictions; in most cases the fits accounted for more than 99% of the variance in the data, and in all cases more than 94%.

Classification of V1 Response Types

We classified cells as simple or complex based upon their response to full-contrast drifting sinusoidal gratings optimized for direction, spatial frequency, temporal frequency, position, and size. We presented gratings for an integer number of cycles and removed the response to the first cycle to eliminate transient onset effects. We took *relative modulation* as the ratio of the magnitude of the vector average response at the grating temporal frequency to the baseline-subtracted mean response. Baseline was taken as the response to a blank (mean gray) screen.

We determined direction selectivity by comparing neuronal responses to the two directions of movement of otherwise optimal drifting gratings. The directional index was defined as $1 - (\text{nonpreferred response} / \text{preferred response})$, with baseline subtracted from both responses.

Predicting the Responses to Arbitrary Stimuli

We computed the predicted response of the energy model by first convolving the stimulus with the STA and the two strongest filters revealed by STC. Similarly, the predicted response of the excitatory model was computed by convolving the stimulus with the STA and all the excitatory filters recovered by STC. For both models, we pooled excitatory signals by a weighted sum of squares. To correct for discretization by binning and to better estimate poorly sampled bins, we used the nonlinear function given in Equation 1 (with the suppressive terms omitted) to relate the output of the relevant filters to firing rate. We predicted response from the full model by convolving the stimulus with each of the recovered excitatory and suppressive filters (including the STA). The excitatory and suppressive signals were combined separately, each via the square root of a weighted sum of squares. Firing rate was computed from the pooled \mathcal{E} and \mathcal{S} signals using the fit of Equation 1 to the data.

For stimuli presented at lower contrast energies than the bar stimulus used in the initial characterization (e.g., Figure 7), the model did not accurately predict firing rates without an additional gain parameter. We simulated this gain adjustment by applying a single scale factor to the pooled \mathcal{E} and \mathcal{S} signals before converting them into firing rates. For the example case in Figure 7, this parameter had a value of 5.

Supplemental Data

The Supplemental Data for this article can be found at <http://www.neuron.org/cgi/content/full/46/6/945/DC1>.

Acknowledgments

We are grateful to Neot Doron and Mian Hou for help with histology; Adam Kohn, Arnulf Graf, James Hedges, and Yasmine El-Shamayleh for help during the experiments; and to Adam Kohn and Jonathan Pillow for helpful discussions. This work was supported by HHMI Investigatorships to E.P.S. and J.A.M. and by a grant to J.A.M. from the National Institutes of Health (EY 2017).

Received: August 23, 2004

Revised: March 14, 2005

Accepted: May 6, 2005

Published: June 15, 2005

References

- Adelson, E.H., and Bergen, J.R. (1985). Spatiotemporal energy models for the perception of motion. *J. Opt. Soc. Am. A* 2, 284–299.
- Aguera y Arcas, B., and Fairhall, A.L. (2003). What causes a neuron to spike? *Neural Comput.* 15, 1789–1807.
- Aguera y Arcas, B., Fairhall, A.L., and Bialek, W. (2003). Computation in a single neuron: Hodgkin and Huxley revisited. *Neural Comput.* 15, 1715–1749.
- Alonso, J.M., Usrey, W.M., and Reid, R.C. (2001). Rules of connectivity between geniculate cells and simple cells in cat primary visual cortex. *J. Neurosci.* 21, 4002–4015.
- Bialek, W., and de Ruyter van Steveninck, R. (2005). Features and dimensions: Motion estimation in fly vision (<http://arxiv.org/abs/q-bio/0505003>).
- Brenner, N., Bialek, W., and de Ruyter van Steveninck, R. (2000). Adaptive rescaling maximizes information transmission. *Neuron* 26, 695–702.
- Carandini, M., Heeger, D.J., and Movshon, J.A. (1997). Linearity and normalization in simple cells of the macaque primary visual cortex. *J. Neurosci.* 17, 8621–8644.
- Cavanaugh, J.R., Bair, W., and Movshon, J.A. (2002a). Nature and interaction of signals from the receptive field center and surround in macaque V1 neurons. *J. Neurophysiol.* 88, 2530–2546.
- Cavanaugh, J.R., Bair, W., and Movshon, J.A. (2002b). Selectivity and spatial distribution of signals from the receptive field surround in macaque V1 neurons. *J. Neurophysiol.* 88, 2547–2556.
- Chance, F.S., Nelson, S.B., and Abbott, L.F. (1999). Complex cells as cortically amplified simple cells. *Nat. Neurosci.* 2, 277–282.
- Chandler, J.D. (1969). Subroutine STEPIT: find local minima of a smooth function of several parameters. *Behav. Sci.* 14, 81–82.
- Chichilnisky, E.J. (2001). A simple white noise analysis of neuronal light responses. *Network* 12, 199–213.
- de Boer, E., and Kuyper, P. (1968). Triggered correlation. *IEEE Trans. Biomed. Eng.* 15, 169–179.
- de Ruyter van Steveninck, R., and Bialek, W. (1988). Real-time performance of a movement-sensitive neuron in the blowfly visual system. *Proc. R. Soc. Lond. B Biol. Sci.* 234, 269–276.
- De Valois, R.L., Albrecht, D.G., and Thorell, L.G. (1982). Spatial frequency selectivity of cells in macaque visual cortex. *Vision Res.* 22, 545–559.
- DeAngelis, G.C., Ohzawa, I., and Freeman, R.D. (1993). Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. I. General characteristics and postnatal development. *J. Neurophysiol.* 69, 1091–1117.
- Emerson, R.C., Citron, M.C., Vaughn, W.J., and Klein, S.A. (1987). Nonlinear directionally selective subunits in complex cells of cat striate cortex. *J. Neurophysiol.* 58, 33–65.
- Forte, J., Peirce, J.W., Kraft, J.M., Krauskopf, J., and Lennie, P. (2002). Residual eye-movements in macaque and their effects on visual responses of neurons. *Vis. Neurosci.* 19, 31–38.
- Geisler, W.S., and Albrecht, D.G. (1992). Cortical neurons: isolation of contrast gain control. *Vision Res.* 32, 1409–1410.
- Heeger, D.J. (1992a). Half-squaring in responses of cat striate cells. *Vis. Neurosci.* 9, 427–443.
- Heeger, D.J. (1992b). Normalization of cell responses in cat striate cortex. *Vis. Neurosci.* 9, 181–197.
- Horowitz, G.D., Chichilnisky, E.J., and Albright, T.D. (2005). Blue-yellow signals are enhanced by spatiotemporal luminance contrast in macaque V1. *J. Neurophysiol.* 93, 2263–2278.
- Hubel, D.H., and Wiesel, T.N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.* 160, 106–154.
- Hubel, D.H., and Wiesel, T.N. (1968). Receptive fields and functional architecture of monkey striate cortex. *J. Physiol.* 195, 215–243.
- Jones, J.P., and Palmer, L.A. (1987). The two-dimensional spatial structure of simple receptive fields in cat striate cortex. *J. Neurophysiol.* 58, 1187–1211.
- Knierim, J.J., and van Essen, D.C. (1992). Neuronal responses to static texture patterns in area V1 of the alert macaque monkey. *J. Neurophysiol.* 67, 961–980.
- Lau, B., Stanley, G.B., and Dan, Y. (2002). Computational subunits of visual cortical neurons revealed by artificial neural networks. *Proc. Natl. Acad. Sci. USA* 99, 8974–8979.
- Levitt, J.B., and Lund, J.S. (1997). Contrast dependence of contextual effects in primate visual cortex. *Nature* 387, 73–76.
- Livingstone, M.S., and Conway, B.R. (2003). Substructure of direction-selective receptive fields in macaque V1. *J. Neurophysiol.* 89, 2743–2759.
- McLean, J., and Palmer, L.A. (1989). Contribution of linear spatiotemporal receptive field structure to velocity selectivity of simple cells in area 17 of cat. *Vision Res.* 29, 675–679.
- Mechler, F., and Ringach, D.L. (2002). On the classification of simple and complex cells. *Vision Res.* 42, 1017–1033.
- Movshon, J.A., Thompson, I.D., and Tolhurst, D.J. (1978a). Receptive field organization of complex cells in the cat's striate cortex. *J. Physiol.* 283, 79–99.
- Movshon, J.A., Thompson, I.D., and Tolhurst, D.J. (1978b). Spatial summation in the receptive fields of simple cells in the cat's striate cortex. *J. Physiol.* 283, 53–77.
- Paninski, L. (2003). Convergence properties of three spike-triggered analysis techniques. *Network* 14, 437–464.
- Priebe, N.J., Mechler, F., Carandini, M., and Ferster, D. (2004). The contribution of spike threshold to the dichotomy of cortical simple and complex cells. *Nat. Neurosci.* 7, 1113–1122.
- Rust, N.C., Schwartz, O., Movshon, J.A., and Simoncelli, E.P. (2004). Spike-triggered characterization of excitatory and suppressive stimulus dimensions in monkey V1. *Neurocomputing* 58–60, 793–799.
- Schwartz, O., Chichilnisky, E.J., and Simoncelli, E.P. (2002). Characterizing gain control using spike-triggered covariance. *Adv. Neural Inf. Process. Syst.* 14, 269–276.
- Simoncelli, E.P., Pillow, J.W., Paninski, L., and Schwartz, O. (2004). Characterization of neural response with stochastic stimuli. In *The Cognitive Neurosciences*, M. Gazzaniga, ed. (Cambridge, MA: MIT Press), pp. 327–338.
- Spitzer, H., and Hochstein, S. (1985). A complex-cell receptive-field model. *J. Neurophysiol.* 53, 1266–1286.
- Touryan, J., Lau, B., and Dan, Y. (2002). Isolation of relevant visual features from random stimuli for cortical complex cells. *J. Neurosci.* 22, 10811–10818.