J Neurophysiol 104: 1884–1898, 2010.
First published August 4, 2010; doi:10.1152/jn.01118.2009.

Linear and Nonlinear Contributions to the Visual Sensitivity of Neurons in Primate Lateral Geniculate Nucleus

Samuel G. Solomon,[1] Chris Tailby,[1,2] Soon K. Cheong,[2] and Aaron J. Camp[1]

[1]*Discipline of Physiology, School of Medical Sciences and Bosch Institute, University of Sydney, Sydney, New South Wales; and* [2]*National Vision Research Institute, Carlton, Victoria, Australia*

Submitted 18 December 2009; accepted in final form 1 August 2010

Solomon SG, Tailby C, Cheong SK, Camp AJ. Linear and nonlinear contributions to the visual sensitivity of neurons in primate lateral geniculate nucleus. *J Neurophysiol* 104: 1884–1898, 2010. First published August 4, 2010; doi:10.1152/jn.01118.2009. Several parallel pathways convey retinal signals to the visual cortex of primates. The signals of the parvocellular (P) and magnocellular (M) pathways are well characterized; the properties of other rarely encountered cell types are distinctive in many ways, but it is not clear that they can provide signals with the same fidelity. Here we study this by characterizing the temporal receptive field of neurons in the lateral geniculate nucleus of anesthetized marmosets. For each neuron, we measured the response to a flickering uniform field, and, from this, estimated the linear and nonlinear receptive fields using spike-triggered average (STA) and spike-triggered covariance (STC) analyses. As expected the response of most P-cells was dominated by the STA, but the response of most M-cells required additional nonlinear components, and these usually acted to suppress cell responses. The STC analysis showed stronger suppressive axes in suppressed-by-contrast cells, and both suppressive and excitatory axes in ON-OFF cells. Together, the STA and the STC analyses form a model of the temporal response to a large uniform field: under this model, the information that was provided by suppressed-by-contrast cells or ON-OFF cells approached that provided by the P- and M-cells. However, whereas P- and M-cells provided more information about luminance, the nonlinear cells provided more information about the contrast energy. This suggests that the nonlinear cells provide complimentary signals to those of P- and M-cells, with reasonably high fidelity, and may play an important role in normal visual processing.

INTRODUCTION

Most ganglion cells in the primate retina and its major target, the dorsal lateral geniculate nucleus of the thalamus (LGN), belong to the parvocellular (P) or magnocellular (M) pathway (Leventhal et al. 1981). All P- and M-cells show receptive fields of the center-surround type and can be reasonably well characterized by linear models (Benardete and Kaplan 1997a,b, 1999; Derrington and Lennie 1984): each neuron responds as if it computes a weighted sum of the intensity of the pattern imaged on its receptive field. Deviations from linearity are most clearly evident in M-cells, all of which exhibit contrast gain controls (Benardete et al. 1992; Kaplan and Shapley 1986; Solomon et al. 2002) and some of which show more severe nonlinearities (Crook et al. 2008b; Kaplan and Shapley 1982; Petrusca et al. 2007; Solomon et al. 2006). Despite these functional differences, the responses of P- and M-cells to achromatic images often appear redundant (Levitt et al. 2001; van Hateren et al. 2002).

Address for reprint requests and other correspondence: S. Solomon, Anderson Stuart Bldg., F13, Univ. of Sydney, Sydney, NSW 2006, Australia (E-mail: samuels@physiol.usyd.edu.au).

Although most studies of primate retina and LGN focus on the work done by the P- and M-pathways, other ganglion cell types project to the LGN (Dacey et al. 2003; Rodieck and Watanabe 1993; Szmajda et al. 2008) and therefore might have direct access to conscious visual sensation. Some of these neurons receive substantial input from S-cones and have predominantly linear receptive fields that seem suited to conveying chromatic information (Crook et al. 2009; Derrington et al. 1984; Lankheet et al. 1998). Other rarely but consistently encountered cell types respond to the retinal image in a fundamentally nonlinear manner (Casagrande 1994; de Monasterio 1978; White et al. 2001). Recent work has characterized the discharge of some of these cells to sinusoidal modulation (Crook et al. 2008a; Petrusca et al. 2007; Tailby et al. 2007) and shows important differences between the receptive fields of these cells and those of the P- and M-pathways. These nonlinear cells probably form part of the koniocellular pathway (Casagrande 1994; White et al. 2001), which in turn is often identified with the sluggish W-pathway of the cat visual system, whose label suggests a pathway of low fidelity. However, it remains unclear whether the information about the stimulus that these cells provide is comparable to that provided by the P- and M-pathways or how their nonlinear signals complement that of P- and M-cells.

Spike-triggered covariance (STC) analyses have recently shown important aspects of nonlinear processing in visual cortex (Chen et al. 2007; Horwitz et al. 2005; Rust et al. 2005; Touryan et al. 2002), but we know much less about what may be already present in the subcortical cells that provide their input, especially in primates. STC analysis showed nonlinearities in an unidentified cell in the macaque retina (Pillow and Simoncelli 2006); two similar analyses in the LGN of macaque are suggestive: although one found few nonlinearities in the (multiunit) response of the blue-yellow pathway (Horwitz et al. 2005), another found nonlinear mechanisms in both M- and P-cells (Sincich et al. 2009). In the retina of salamanders, STC analysis has shown several distinct classes of receptive field (Fairhall et al. 2006). Using STC analyses, we establish here that there are strong nonlinearities in the temporal receptive fields of linear cells (particularly those in the M-pathway) and the very nonlinear cells. The nonlinearities modulate the sensitivity of M-cells and P-cells; in the very nonlinear cells, they endow a code for contrast energy, a very different signal to those provided by P- and M-cells.

METHODS

Surgery

Experiments were undertaken as part of a larger series on eight adult male common marmosets (*Callithrix jacchus jacchus*), weighing

between 350 and 430 g. All procedures conformed to the guidelines approved by Animal Ethics Committee of the University of Sydney. Each animal was initially sedated with an intramuscular injection of 12 mg/kg Alfaxan (Jurox) and 3 mg/kg Diazepam (Roche). We gave preoperative intramuscular injections of 0.2 mg/kg atropine (Pfizer) to reduce lung secretions and dexamethasone (0.3 mg/kg; Maine Pharmaceuticals) to reduce inflammation. Subsequent surgery was performed under supplemental local anesthesia (Lignocaine 2%; Astra Zaneca). A femoral vein was cannulated, the trachea was exposed, and an endotracheal tube was inserted. The head was placed in a stereotaxic frame, and a craniotomy was made over the LGN.

Postsurgical anesthesia was maintained by continuous intravenous infusion of sufentanil citrate (4–12 μg/kg/h, Sufenta Forte, Janssen Cilag) in physiological solution (sodium lactate, Baxter International) with added dexamethasone (0.4 mg/kg/h; Mayne Pharma) and Synthamin 17 (225 mg/kg/h, Baxter International). The ECG and EEG were monitored continuously. Muscular paralysis was induced and maintained by continuous infusion of pancuronium bromide (0.3 mg/kg/h; Astra Zaneca). The animal was artificially ventilated so as to keep end-tidal CO_2 near 33 mmHg. At any sign of the anesthesia becoming less effective, the dose of sufentanil citrate was increased. Rectal temperature was kept near 38°C with the use of a heating blanket. Additional antibiotic and anti-inflammatory cover was given daily by intramuscular injection of 25 mg Noricillin (Norbrook), and 0.1 mg dexamethasone. The pupils were dilated with atropine sulfate, and the corneas were protected with high-permeability contact lenses that remained in place for the duration of the experiment. No artificial pupils were used. Supplementary lenses (with power determined by maximizing the spatial resolution of parvocellular cells) were used to focus the eyes at a distance of 114 cm.

A small incision was made in the dura, and a guide tube containing the electrodes was inserted and positioned above the LGN (most measurements were obtained with single electrode recordings: paralyene-coated tungsten, FHC, 9–12 MOhm; 3 nonlinear cells were identified in a larger set of tetrode recordings: Thomas Recordings tetrodes, 2–5 MOhm). At the end of the experiment, the monkey was killed with intravenous 500 mg/kg sodium pentobarbitone (Lethobarb, Verbac Australia).

Visual stimuli and recording

A front-silvered mirror was used to bring the receptive field onto the center of a cathode ray tube monitor (ViewSonic G810, 100 Hz refresh rate; or Sony G520, refresh rate 120 or 120.4 Hz), viewed from 114 cm. Visual stimuli were generated by a G5 Power Macintosh computer using custom software (EXPO; P. Lennie); they were drawn with 8-bit resolution using commands to OpenGL. For each phosphor, we determined the relationship between the output of the video card and the photopic luminance; the inverse of this relationship was applied to the image that was sent to the video card. The stimulus was a drifting sinusoidal grating or a uniform field modulated in time; all stimuli modulated around the mean luminance (45–55 cd/m^2) and were presented within a circular window with hard edges (diameter 8° unless otherwise specified), outside of which the screen (20 × 15°) was held at the mean luminance. The output of the three monitor phosphors could be adjusted to produce modulations of specific cone photoreceptors, using knowledge of the spectral radiance distribution of the monitor phosphors, the peak sensitivity of marmoset cone photoreceptors, and the spectral absorption characteristics of the optic media and macular pigment (Blessing et al. 2004; Tailby et al. 2008b). In most experiments, the full calibration was used; in three, we did not have the full spectral distribution of the monitor phosphors and therefore we used conversion matrices that were derived from previous calibrations. The responses to S cone modulation of M- and P-cells encountered in all animals were nevertheless very small, encouraging us to believe that we substantially isolated the signals of the S cone from those of the LM cone. The S cone–isolating stimulus produced 60–80% contrast in S cones and <5% contrast in LM class cones. The LM cone–isolating stimulus produced >60% contrast in LM class cones and <5% contrast in S cones. The receptive fields of all cells were located within 15° of the fovea. The analogue signals from the electrodes were amplified, filtered, and sampled at 48 kHz by the same computer that generated the visual stimulus. Putative spikes were displayed on a monitor, and templates for discriminating spikes were constructed by analyzing multiple traces. The timing of waveforms that cohered to the template was recorded with an accuracy of 0.1 ms. Off-line analysis was performed using Matlab (MathWorks, Natick, MA).

Cell identification

For each cell, we determined the sign of response (ON, OFF) to achromatic and cone-isolating modulation of spatially uniform fields and the tuning for temporal frequency, spatial frequency, and contrast using drifting achromatic gratings. Along with the pattern of transitions between eye representations and in some cases subsequent histological reconstructions, these measures were used to classify most cells as part of the P- or M-pathway (Derrington and Lennie 1984; Dreher et al. 1976; White et al. 2001). We encountered several types of cells that could not be functionally defined as P- or M-cells, including the very nonlinear suppressed-by-contrast cells and ON-OFF cells described here.

During presentation of drifting gratings, the response of both suppressed-by-contrast and ON-OFF cell types is best characterized by changes in the mean rate. Figure 1, *C* and *D,* shows for ON-OFF and suppressed-by-contrast cells how the mean rate and the modulated rate (at the stimulus frequency) depend on the spatial frequency of the drifting grating. For the ON-OFF cell, spatial frequencies up to ~3 cycles/° increase the mean rate and produce a smaller modulated component. The high maintained discharge of the suppressed-by-contrast cell was reduced over a similar range of frequencies, and at very low spatial frequencies, the gratings also modulated the discharge (see Tailby et al. 2007 for very similar tuning functions obtained from LGN of macaque). Figure 1, *E* and *F,* shows that the characteristic changes in mean rate for these types of cells are apparent over a range of temporal frequencies.

Spike-triggered average and STC

The primary visual stimulus used here was a uniform field whose luminance was updated at the refresh rate of the monitor and drawn from a Gaussian distribution with zero mean and an SD of 0.3 (where 1 is a full increment from the background and −1 is a full decrement). To determine the temporal properties of mechanisms that contributed to the response of cells, we estimated the average (STA) and covariance (STC) of the spike-triggered stimulus ensemble (Brenner et al. 2000; de Ruyter van Steveninck and Bialek 1988; Schwartz et al. 2006).

The luminance of each frame was stored and rearranged into a series of stimulus vectors, *s*, one for each frame in the sequence. The column vector *s*, where s_t is the luminance on frame *t,* extends 50 frames into the past. The occurrence of each spike was found, and these spikes were binned into a vector *r* with the same temporal resolution as the stimulus sequence and aligned so that the time covered by the bin corresponded to the last frame in each vector *s*.

The spike-triggered stimulus ensemble is the matrix of vectors *s* that was associated with a nonzero bin in *r*; each vector *s* was weighted by the number of spikes in the associated time bin of *r*. The *STA* is the average of this weighted stimulus ensemble

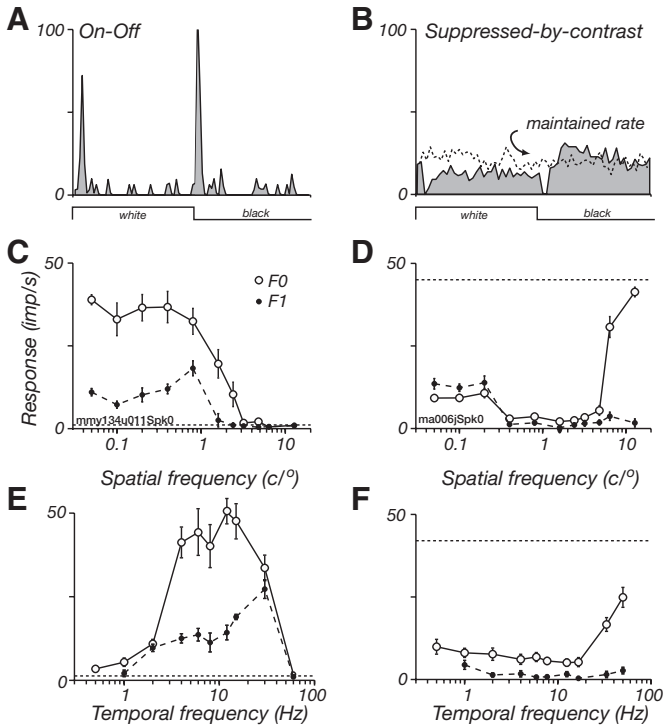

FIG. 1. Identification of ON-OFF and suppressed-by-contrast receptive fields in marmoset lateral geniculate nucleus. *A, C, E*: ON-OFF cell. *B, D, F*: suppressed-by-contrast cell. *A, B*: peristimulus time histograms (PSTHs) of the discharge rate during presentation of a large uniform field whose luminance was modulated in time by a 0.5 Hz square wave. Also shown is the maintained discharge rate, calculated in the same way, but for a blank screen held at the mean luminance. *C, D*: responses of the same cells to drifting sinusoidal gratings of varying spatial frequency (temporal frequency 5 Hz). The open symbols show the mean discharge rate; the closed symbols show the amplitude of modulation at the temporal frequency of the grating; and the dashed lines show the maintained discharge rate. *E, F*: the response of the same cells to gratings of varying temporal frequency (spatial frequency 0.8 and 2 cycles/°, respectively). Conventions as in *C*. Michelson contrast in all cases 1.0. Error bars are ± SE.

$$STA = \frac{1}{m}\sum_{i=1}^{N} r_i \times s_i \tag{1}$$

where m is the total number of spikes and N is the number of vectors in the spike-triggered stimulus ensemble. *STA* can be considered an estimate of the linear temporal receptive field. To quantify the variability associated with this estimate (as in Fig. 2), we divided the stimuli and responses into 10 segments of equal length and performed the analysis separately on each segment. To estimate the varibility in the STA, we estimated the signal-to-noise (*SNR*) ratio of the segment STAs. The SNR was defined as ratio of the amplitude of the average waveform (the STA) to twice the SD of its noise

$$SNR = \frac{\max(STA) - \min(STA)}{2 \times SD(STA_\varepsilon)} \tag{2}$$

where $STA\varepsilon$ is the difference between the average *STA* and the 10 segment *STA*s, across all segments and all time points (to concentrate this measure on the relevant time period we used only the 20 frames before a spike).

We wanted to know whether other aspects of the stimulus sequence determined the spiking response of the cell under study, and here we projected *STA* out of the stimulus ensemble [The reader should note that this is not always done in this type of analysis (Fairhall et al. 2006)]. First we made a unit vector of *STA*, $\widetilde{A}$

$$\widetilde{A} = \frac{STA}{|STA|} \tag{3}$$

We then projected $\widetilde{A}$ out of the spike-triggered stimulus ensemble (S) to make a new ensemble $\widetilde{S}$

$$\widetilde{S} = S - S\widetilde{A}^T\widetilde{A} \tag{4}$$

where ^{T} denotes matrix transpose. This ensures that the axes of the STC will be orthogonal to *STA*. Finally we calculated the covariance of this ensemble

$$C = \frac{\widetilde{S}^T\widetilde{S}}{m} \tag{5}$$

and decomposed the covariance matrix C to obtain eigenvalues and their associated eigenvectors.

This analysis aims to find axes in the stimulus space along which the variance is smaller or larger than expected of a spherical Gaussian distribution; that is, the axes along which the neuron is more responsive or less responsive than expected. To place confidence intervals, we repeated the process 1,000 times, each iteration shifting the times of the action potentials relative to the stimulus by a random value constrained to exceed 1 s.

Information captured by the model

To examine the relationship between projection of the stimulus on *STA* and discharge rate, we first normalized *STA* to unit variance, then

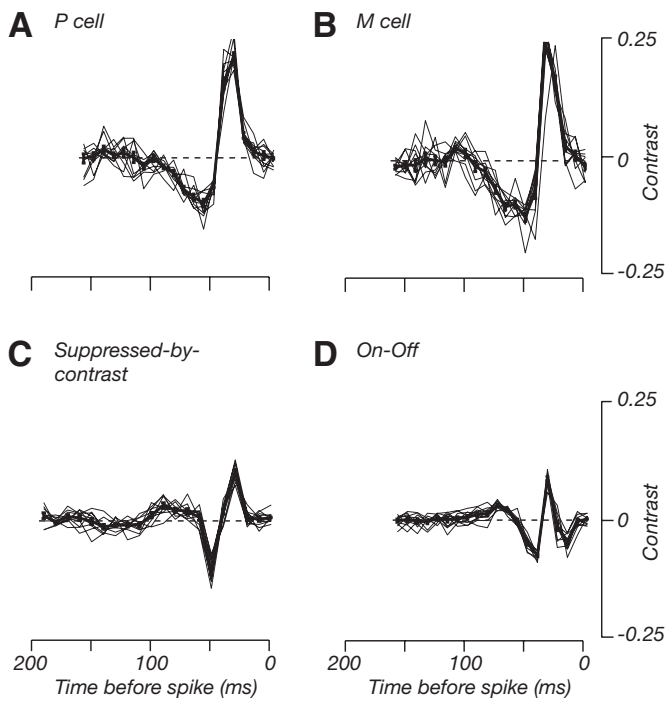

FIG. 2. Temporal receptive fields of lateral geniculate nucleus of the thalamus (LGN) neurons provided by spike-triggered average (STA). *A*: parvocellular (P) cell. *B*: magnocellular (M) cell. *C*: suppressed-by-contrast cell. *D*: ON-OFF cell. The stimulus was a large field, whose luminance was randomly drawn from a Gaussian distribution on every screen refresh. The curves show the weighted average luminance in each frame preceding each action potential, the STA. The thick lines show the STA calculated from the entire stimulus; the thin lines show the STA calculated from each of 10 segments of the stimulus. For clarity, only the 20 stimulus frames before each spike are shown (200 ms in *C* and 160 ms elsewhere). The horizontal dashed lines show the mean luminance; upward excursions show when the average luminance was higher than the background. Error bars (usually too small to see) are ±SE.

determined the projection of every stimulus (both those that led to a discharge and those that did not) onto the normalized *STA*, and finally sorted the projections into 40 equally spaced bins that spanned the observed range. The information captured by *STA* or each eigenvector can be estimated by making two probability distributions: $P(v)$, the probability that the projection v of any stimulus was found in each bin, and $P(v|spike)$, the probability that a stimulus that led to a spike was found in the same bins (in each case the sum of the probability distribution is unity). The information, I, in bits per spike, is

$$I = \sum_{i=1}^{N} P(v_i|spike) \times \log_2\left[\frac{P(v_i|spike)}{P(v_i)}\right] \quad (6)$$

where N is the number of bins and i is the index to that bin (Adelman et al. 2003; Aguera y Arcas and Fairhall 2003; Fairhall et al. 2006; Sharpee et al. 2004; Sincich et al. 2009). The estimate of I carries an upward bias: to estimate that bias, we repeated the analysis for the eigenvector associated with the most insignificant eigenvalue. The I returned by this analysis was small but not zero (for large fields: μ 0.023, SD 0.031; $n = 63$). The values of I reported and used throughout are the difference between this "noise I" and that provided by the relevant vector, calculated independently for each neuron.

To characterize the interaction between the STA and an eigenvector, or two eigenvectors, we estimated the information in their joint probability distribution (Fairhall et al. 2006). This is provided by substituting the appropriate joint probability distributions for $P(v)$ and $P(v|spike)$ in *Eq. 6*. The information in the joint distribution also carries bias, and to estimate that, we performed the calculation for the STA and most insignificant eigenvector. For each cell, we defined the bias as the difference in the information provided by this joint distribution and that provided by the STA alone: the bias is larger than for the single vectors above because there are fewer observations in each bin; on average, it was 0.077 (SD = 0.100, $n = 63$). The information in the joint distribution reported here is the difference between the measured information and this estimate of bias.

Descriptive function for the relationship between STA and the discharge rate

To characterize the input–output function of the STA, we again binned the projections of every stimulus onto the STA. We calculated the average spike rate (in impulses/s) of the stimuli associated with each bin. We fit a cumulative normal distribution (Chichilnisky 2001), with free parameters defining the mean and SD of the distribution and another providing an overall scale factor. Most of the stimuli have projections of near-zero amplitude, so different bins have different numbers of observations. In fitting the cumulative normal, we wanted to give less weight to those bins that had less reliable observations. Our time bins were small enough (8–10 ms) that there was usually only one or zero impulses in them, so we made the simplifying assumption that there was only ever one or no impulses (whereby a sample from a Poisson distribution can be approximated by binary random draw)

$$SD = \sqrt{\frac{P \times (1 - P)}{N}} + \frac{1}{2 \times N} \quad (7)$$

where P is the probability of an impulse associated with the projections falling in the relevant bin, and N is the number of trials contributing to that estimate. Our fitting procedure (*lsqnonlin* in the Matlab environment) used this value to normalize the square error between the prediction of the cumulative normal and the observed firing rate, so bins associated with less reliable observations contributed relatively less to the overall fit.

Decoding trains of impulses to recover luminance and energy

A complementary analysis is that of "decoding," where we try to reconstruct the stimulus given the observed action potentials. We used standard methods of optimal linear decoding, described in detail elsewhere (Warland et al. 1997). Briefly, we divided the stimulus sequence into 11 equally long blocks. We used 10 of these blocks (30,462 frames) to train a decoder, which recovered the filters (which resemble but are not usually identical to the STA) that minimized the squared difference between the observed and predicted stimuli. To estimate how well these filters had done, we used them to predict the sequence of luminance in the 11th block (3,046 frames). We repeated this process for each of the 11 blocks, always training the decoder on the remaining frames. To characterize the performance, we divided the predicted and actual sequences into windows of 100 frames for Fourier analysis. The frequency spectrum of the predicted sequence, $x(\omega)$, is the signal provided by the neuron, and the frequency spectrum of the error between the predicted and actual sequences, $r(\omega)$, is the noise associated with that signal. The signal-to-noise ratio was estimated, at each temporal frequency from the fundamental to the Nyquist limit, from the ratio of signal power to noise power at that frequency, $\log SNR = \log_2[1 + x(\omega)^2/r(\omega)^2]$. To summarize this, we integrated the signal-to-noise ratio across the frequency range, which is an estimate of the mutual information between the actual and predicted sequences of luminance, I_r

$$I_r = \frac{1}{2}\sum_{\omega=1}^{N} \log SNR(\omega) \quad (8)$$

where N is the highest harmonic measurable (here, the 50th). In a parallel set of analyses, the stimulus sequences were first squared, and from this, we recovered the mutual information between the predicted and observed contrasts. Despite their inaccuracies, we refer to analysis of the raw sequence as luminance decoding and of the squared sequences as energy decoding. These estimates are biased because the sequence is of finite length, and therefore for each cell, we performed the same estimates after shifting the trains of impulses by 3,046 frames. This preserves the structure of the impulse trains and the stimulus but not the relationship between the two. The estimates of signal-to-noise ratio and mutual information for single cell reconstructions that are shown in Fig. 10 are the differences between that obtained for real and shifted trains of impulses.

RESULTS

In the following, we describe measurements from 28 P-cells (16 ON, 12 OFF) and 11 M-cells (7 ON, 4 OFF) in three adult marmosets. From these and another five animals, we also made measurements from nine suppressed-by-contrast cells and four ON-OFF cells. These cells were readily distinguished from the ON- and OFF-subtypes of P- and M-cells by their discharge during presentation of a uniform field whose luminance modulated in time. The ON-OFF cell has a low maintained discharge rate, which increases in response to both light increments and light decrements (Fig. 1*A*). The suppressed-by-contrast cell (Rodieck 1967) or uniformity detector (Cleland and Levick 1974) has a high maintained discharge rate, which is suppressed by both light increments and decrements (Fig. 1*B*). The marmosets that were studied here were all male and thus all dichromatic, having one cone photoreceptor sensitive to short wavelengths (the S cone) and another sensitive to the long and medium wavelengths (an LM cone). The receptive fields of P- and M-cells studied here showed very weak or no response to modulation of the S cone; we characterized four additional

neurons that had distinctive blue-ON receptive fields with clear excitatory responses to increments in S cone activity. In seven of eight suppressed-by-contrast cells and three ON-OFF cells where appropriate measurements were made, the discharge was also sensitive to modulation of the S cone. We will return to this in DISCUSSION.

Linear contributions to temporal receptive fields in LGN

For each cell, activity was measured during presentation of a large uniform field (usually 8° in diameter), the luminance of which was chosen from a Gaussian distribution and updated every 8 or 10 ms (the rate at which the monitor was refreshed). From the responses to the flickering field, we obtained the average and covariance of the stimuli that preceded spikes (where each stimulus was considered to be the 50 frames before each spike). The STA characterizes the linear behavior of receptive fields, whereas the STC helps characterize the nonlinear behavior (de Ruyter van Steveninck and Bialek 1988; Fairhall et al. 2006; Rust et al. 2005; Schwartz et al. 2006).

The thick lines in Fig. 2, *A–D,* shows the STA for four example neurons. The STA of the ON-center P- and M-cells are biphasic and resemble those described previously in macaque (Benardete and Kaplan 1997a, 1999). Their shapes imply that the cells will be more likely to generate action potentials when the image on its receptive field changes rapidly from dark to light. The analysis also returned STAs in the nonlinear cells: these are more complicated and generally carry more lobes than those of P- and M-cells. This suggests that the linear component of the temporal receptive field might be limited to very specific temporal sequences and perhaps only a limited range of temporal frequencies. To establish the robustness of the STA, we divided the stimulus and response sets into 10 equal segments and calculated an STA for each of those segments. The thin lines in Fig. 2, which show the STA obtained from each of these shorter segments, are all close to that obtained from the entire dataset, shown by the thick line, suggesting that the STA reflects mechanisms within the temporal receptive field that are stable over time. This was the case for all neurons in our dataset; to quantify the variability across the population of cells, we calculated a signal-to-noise ratio for the segment STAs by comparing them to the STA obtained from the entire dataset (*Eq. 2*, METHODS). For the neurons in Fig. 2, *A–D,* these were, respectively, 7.1, 6.1, 6.7, and 6.3. Across all P-cells, they were on average 6.3 (SD 1.9, $n = 28$), across M-cells, they were 8.3 (SD 3.6, $n = 11$), across suppressed-by-contrast cells, they were 4.7 (SD 2.1, $n = 9$), across ON-OFF cells, they were 9.2 (SD 9.2, $n = 4$), and across blue-ON cells, they were 7.3 (SD 3.2, $n = 4$).

Nonlinear contributions to receptive fields in LGN

The STA shows the average temporal pattern of luminance that is associated with a spike. It will therefore fail to capture elements of a receptive field that responds equivalently to increments and decrements. STC analysis can overcome this by finding axes in the spike-triggered stimulus space that are conspicuously elongated or constricted.

Figure 3, *A–D,* shows the result of the STC analyses on the neurons in Fig. 2 (see METHODS). Briefly, we project the STA out of the spike-triggered stimulus ensemble and identify axes in the resultant ensemble that have significantly large or small variance. The STC analysis returns eigenvalues (which reflect the variance along the axes) and their associated eigenvectors (which reflect the stimulus structure along those axes). We sorted the eigenvalues and these are plotted in the middle panels of Fig. 3.

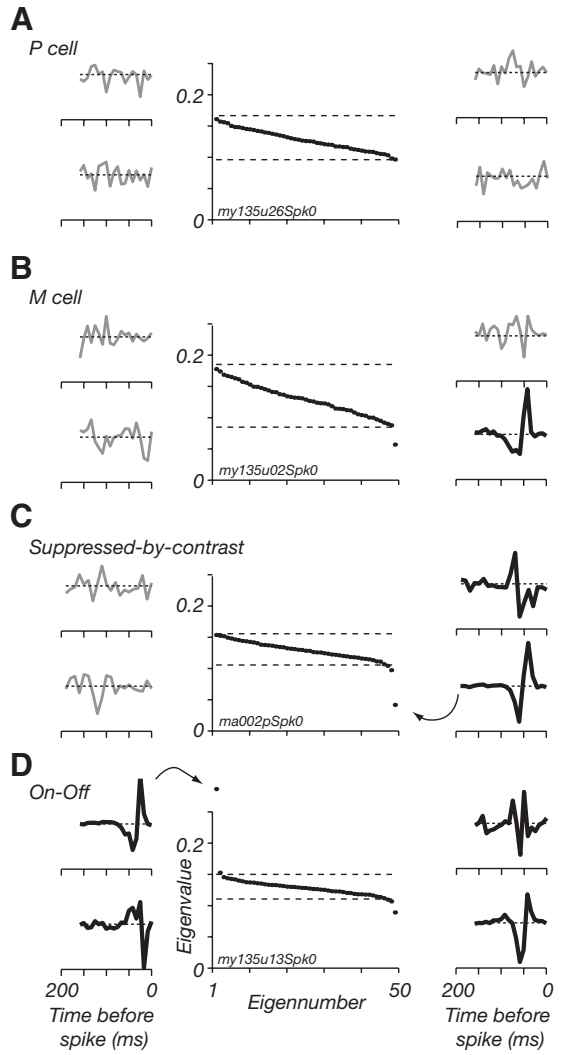

FIG. 3. Temporal receptive fields of LGN neurons provided by spike-triggered covariance (STC). *A*: P cell. *B*: M cell. *C*: suppressed-by-contrast cell. *D*: ON-OFF cell. Same cells as in Fig. 2. The stimulus was a large uniform field modulated in time. Filled symbols in *middle panels* show the eigenvalues returned by the analysis, sorted by their eigenvalue. Dashed lines show 99% confidence limits on these values obtained by 1,000 bootstrap resamples of the discharge; eigenvalues that lie outside the bounded region are unlikely to have arisen by chance. *Left*: the temporal profile of the eigenvectors associated with each of the 2 largest eigenvalues (that for the largest eigenvalue is shown at *top*). *Right*: the temporal profile of the eigenvectors associated with the 2 smallest eigenvalues (that for the lowest eigenvalue is shown at *bottom*). Thick lines show eigenvectors that are associated with significant eigenvalues.

Figure 3 also shows the eigenvectors associated with the four most extreme eigenvalues: in the *left panels* are the two vectors with maximal values (the top one has the largest eigenvalue); and in the *right panels* are the two with minimal values (the bottom one has the smallest eigenvalue). To determine which eigenvectors reflected significant directions in the spike-triggered stimulus ensemble, we repeated the analysis

1,000 times, but each time shifted the spike train relative to the stimulus by a random amount, thereby destroying the relationship between the two (see METHODS). The dashed lines in the middle plot show the 99% confidence limits on the eigenvalues returned from these bootstraps; eigenvalues outside of this region represent significant deviations from the expected variance, and the eigenvectors that correspond to them are shown by thicker lines in the panels on the *left* and *right*.

As expected from their generally linear behavior, none of the eigenvalues recovered for the example P-cell were significant, and it is clear that their associated eigenvectors lack obvious structure. For the example M-cell, an eigenvalue "pops-off" the bottom, lying below the lower bound of the confidence limit of the distribution. This corresponds to an axis of unexpectedly low variance within the spike-triggered ensemble, indicating that stimuli that project onto this axis are associated with a reduced probability of spiking: suppression. In the suppressed-by-contrast cell, more than one eigenvalue was significantly small (indicating several suppressive axes), and in the ON-OFF cell, both significantly small and large eigenvalues (indicating excitatory axes) were observed.

Across 28 P-cells tested, 9 showed significantly suppressive eigenvalues (5 ON cells and 4 OFF cells) and 4 showed significantly excitatory ones (2 ON, 2 OFF; 3 cells showed both excitatory and suppressive eigenvalues). Among 11 M-cells, all 7 ON-cells showed suppressive eigenvalues, as did 2 of 4 OFF-cells; 3 ON-cells and 3 OFF-cells showed excitatory eigenvalues. All nine suppressed-by-contrast cells showed suppressive eigenvalues—seven of these showed more than one significantly suppressive axis, something not seen in the M-cells—and four had excitatory eigenvalues. Among four ON-OFF cells, three had significant suppressive and excitatory eigenvalues; one other responded weakly to the stimulus and showed no clear evidence of either axis. Three of the four blue-ON cells showed suppressive eigenvalues, although only one of the associated eigenvectors showed clear structure, and the same neuron also showed an excitatory eigenvalue.

Important trends become apparent when considering the Fourier spectra of the STA and the eigenvectors. The upper row of Fig. 4 shows the Fourier spectra of the STA: as expected, M-cells prefer higher temporal frequencies and P-cells are more sensitive to lower frequencies; the STA of suppressed-by-contrast cells is more variable. The *bottom row* of Fig. 4 shows that suppressive eigenvectors were also reasonably homogenous within functional classes of LGN cells. The major suppressive eigenvector found in all M-ON cells, all suppressed-by-contrast cells, and some P-ON cells is band-pass and sensitive to high temporal frequencies. The suppressive eigenvectors of four P-OFF cells were responsive to lower temporal frequencies and less band-pass; clearly defined suppressive vectors were not found in the four M-OFF cells studied. This suggests that *1*) nonlinearities in ON- and OFF-cells can be quite different (Chichilnisky and Kalmar 2002) and *2*) at least for P-ON, M-ON, and suppressed-by-contrast cells, the suppressive mechanism has similar characteristics.

Impact of linear and nonlinear mechanisms on spiking probability

To see the relationship between the projection of the stimulus on the STA and the probability of generating a spike, we first calculated the projection of all 50-frame stimulus sequences in our ensemble onto the STA and placed these into equally spaced bins. For the same bins, we also determined the number of stimuli that were associated with a spike and for each bin we could therefore calculate the proportion of stimuli that were associated with a spike, providing an input–output function for the STA. Spike rate cannot be less than zero, and for most cells, there was a threshold below which all stimuli were associated with a lack of discharge. For P- and M-cells, the input–output function generally grew monotonically with projection on the STA, and in M-cells particularly, it saturated at large projection values; the input–output function of the STA in Fig. 5*A* shows this for the example M-cell. For ON-OFF cells, this function was also usually monotonic (Fig. 5*D*). For most suppressed-by-contrast cells, the function was not monotonic (an STA is nevertheless recovered because the

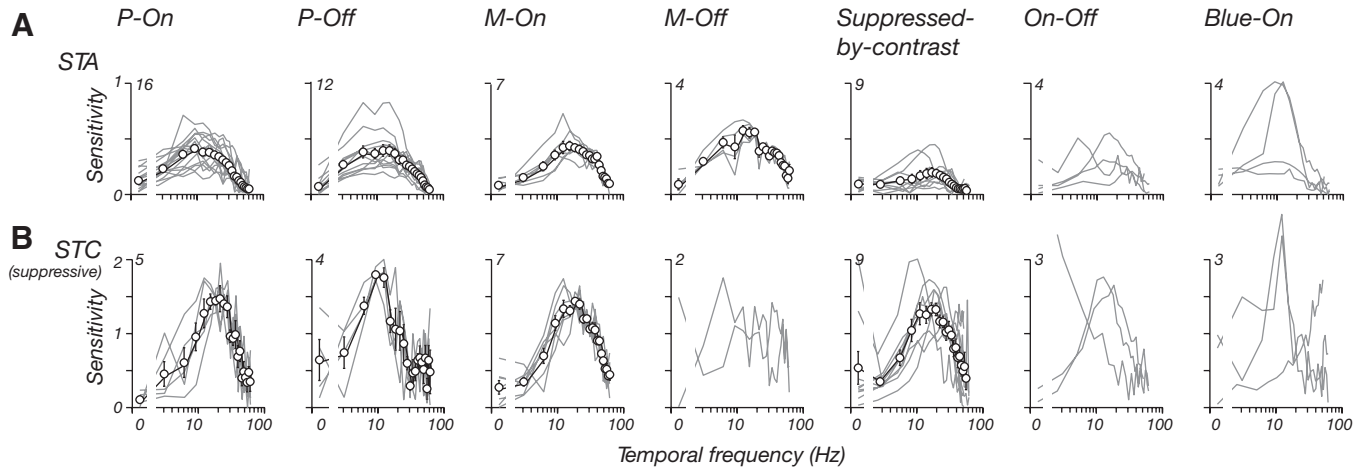

FIG. 4. Population modulation-transfer functions provided by STA and STC analyses. For each cell, the temporal profiles obtained as in Figs. 2 and 3 were subject to Fourier analysis. *A*: each panel shows for each identified class of LGN neuron the modulation transfer function of individual STAs (thin lines) and the average across cells (open symbols, thick lines). *B*: same as *A*, but for the most suppressive eigenvector returned by the STC analysis. Only cells with significant eigenvectors are shown. Conventions as in *A*. In *A* and *B*, some responses were obtained with a refresh rate of 120 Hz and some were obtained with a refresh rate of 100 Hz; in each case, the 1st 25 frequencies were used; to plot the average modulation transfer function for each of 25 bins, we computed the geometric average frequency across cells. Error bars are ±SE.

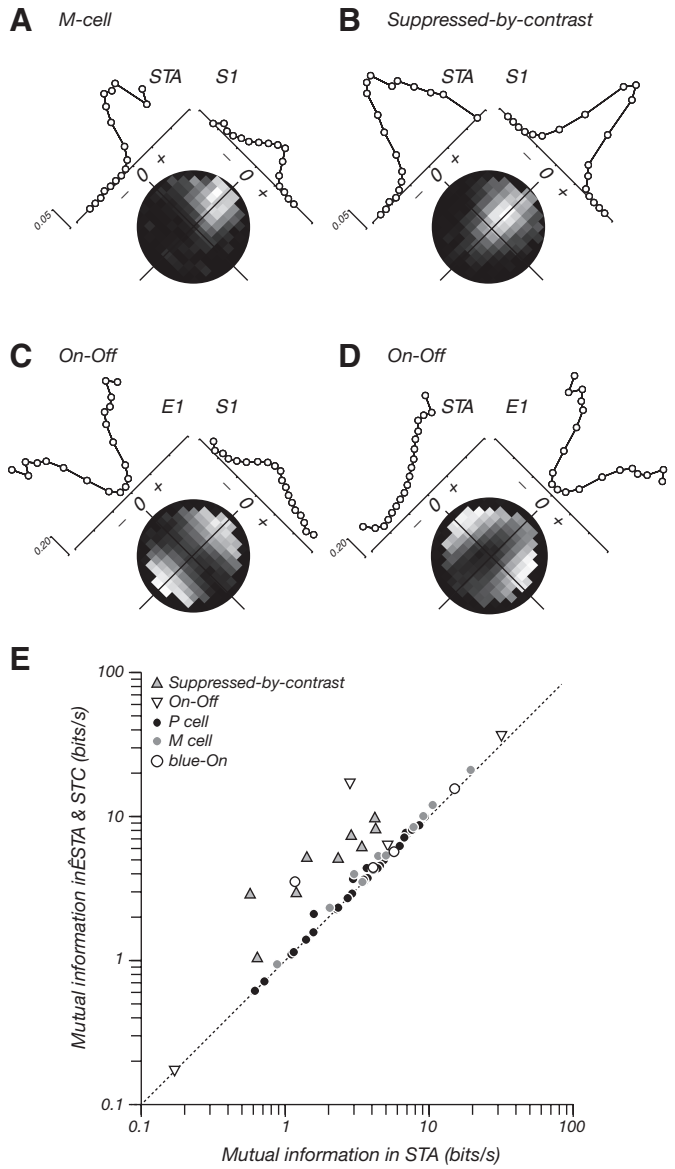

FIG. 5. Input–output functions for STA and STC analyses. *A–D*: firing rate for the example cells is shown as a function of the projection onto 1 vector or the joint projection onto 2 vectors. In the images, the probability of a spike is indicated by pixel intensity; the graphs along the diagonal margins show the average probability for all projections onto 1 of the vectors. For display purposes, pixels in the joint functions were set to 0 if they contained <5% of the number of stimulus projections found in the most populated bin. *A*: M-cell. The average discharge rate of the M-cell was 4.9 impulses/s; the mutual information between the STA and the spike rate was 0.90 bits/spike; for the most suppressive eigenvector (S1) was 0.18; for the joint function was 1.08. *B*: suppressed-by-contrast cell. Average discharge rate: 12.9; STA: 0.22 bits/spike; S1: 0.32; joint: 0.53. *C* and *D*: ON-OFF cell. Average discharge rate: 27.6; STA: 0.10 bits/spike; S1: 0.05; E1: 0.48. Joint function for E1 and S1 (*C*): 0.46 bits/spike. Joint function for STA and E1 (*D*): 0.62. *E*: comparison of information rates (bits/s) for the STA and that for the STA and all the significant eigenvectors revealed by STC analysis. Large distances above the positive diagonal represent cells in which the eigenvectors contributed a substantial fraction of the information available. Average information rates when eigenvectors were included: P-cells: μ 3.92, SD 2.55; M-cells 7.03, SD 5.72; suppressed-by-contrast cells 5.40, SD 2.80.

center-of-mass of the function is not zero and thus the mean of the spike-triggered ensemble is not the same as that of the entire stimulus set).

We made the same calculations for the significant eigenvectors that were returned by the STC analysis. The nature of this analysis means that the relationship between the STC and the discharge rate is always of the same basic shape—excitatory eigenvectors resemble parabolas and suppressive eigenvalues inverted parabolas. Figure 5 shows this for the most excitatory (E1) and suppressive (S1) eigenvectors in the example cells.

This analysis does not tell us how the mechanisms characterized by the eigenvectors interact with each other or with that characterized by the STA (although the eigenvectors and the STA are all by definition orthogonal and independent, their influence on the discharge may not be). We can establish this relationship with the kinds of joint input–output functions that are shown in Fig. 5, *A–D* (Fairhall et al. 2006; Rust et al. 2005; Schwartz et al. 2006). Each joint function is shown by a grayscale image, in which the lightness of each pixel corresponds to the probability of a spike, for each pairing of STA and eigenvector output. In each case, the images resemble what we would expect if the two mechanisms are independent (separable), which would be the product of the two marginal distributions: this was the case for almost all of the pairs of eigenvectors and pairs of eigenvectors and STA in our sample.

Information provided by linear and nonlinear mechanisms

One way of estimating how much of the discharge is captured by the STA or the eigenvectors is to calculate the mutual information between the projection onto the relevant vector and the probability of a spike (see METHODS). To do this, we made two probability distributions: one of the projection of all stimuli onto the STA or eigenvector and one of the projections that were associated with a spike. The mutual information between these (*Eq. 6*) tells us how much the observation of a spike reduces our uncertainty about the stimulus and therefore allows comparison of the informativeness of the STA and eigenvectors; the values for the example cells can be found in the legend to Fig. 5.

The STA of M-cells provided the most bits/spike, on average, 0.70 (SD 0.21), which is substantially more than that of P-cells (μ 0.42, SD 0.30; $P < 0.05$, Students *t*-test), which in turn gave more than suppressed-by-contrast cells (0.15, SD 0.13). The STA of the ON-OFF cells provided intermediate levels of information, and those of the blue-ON cells (also measured using achromatic stimuli) approached that of the M-cells.

The information provided by the eigenvectors was distributed differently. In the three ON-OFF cells with significant excitatory eigenvectors, the primary one provided 0.18 (SD 0.26) bits/spike, which was considerably more than that provided by excitatory eigenvalues in the four P- and six M-cells (respectively μ 0.03, SD 0.02, and μ 0.02, SD 0.01). Similarly, the information provided by the suppressive eigenvectors in the nine suppressed-by-contrast cells (μ 0.15, SD 0.11) was higher than that in the nine P-cells and nine M-cells (μ 0.05, SD 0.06, and μ 0.08, SD 0.05, respectively) or three blue-ON cells (μ 0.08, SD 0.07). All this suggests that the significant eigenvectors of P- and M-cells are much less informative than the STA but that in the nonlinear cells they can be just as informative.

We quantified the separability of the joint functions (shown in Fig. 5, *A–D*) by calculating the information in them, in the same way as we did for the STA and eigenvectors above (see METHODS). If the information in the joint function was greater

than the sum of the information provided by the individual functions alone, this would provide evidence for a lack of independence and the presence of synergistic interactions (Fairhall et al. 2006). We express the increase in information provided by the joint input–output function as a percentage of the information that was provided by the individual vectors alone: in two of eight suppressed-by-contrast cells, the most suppressive eigenvector and STA were synergistic (providing an increase in information of 11 and 53%; the information in the joint function of the latter neuron was, however, only 0.03 bits); in the three ON-OFF cells with significant eigenvectors, the most excitatory ones were synergistic with the STA (12, 16, and 7%). The last of these is shown in Fig. 5*D*. Across all other significant eigenvectors, including those in P-, M-, and blue-ON cells, only one provided a synergy fraction of >5%. We conclude that the mechanisms studied here are largely independent and that the information that they provide can be captured by the sum of the information provided individually.

The eigenvectors were more important in the nonlinear cells than in P- and M-cells, and in addition, the spike rate of the nonlinear cells was higher during stimulation with a large field (During stimulation, the average discharge rates of suppressed-by-contrast cells and ON-OFF cells, respectively 19.7 and 26.8 impulses$\cdot s^{-1}$ were greater than that of P- and M-cells, respectively, 10.2 and 9.8 impulses$\cdot s^{-1}$.) The result is that, when the STA and the eigenvectors were considered together, the information rates of each cell class, expressed in bits/s, was similar. This is shown in Fig. 5*E*, which compares for each cell the information captured by the STA and the sum of the information from the STA and all the significant eigenvectors. The STA of the P- and M-cells provides more information, but these cells lie near the unity line; the suppressed-by-contrast cells are shifted above it, bringing their information rates close to those of the P-and M-cells when the eigenvectors are included.

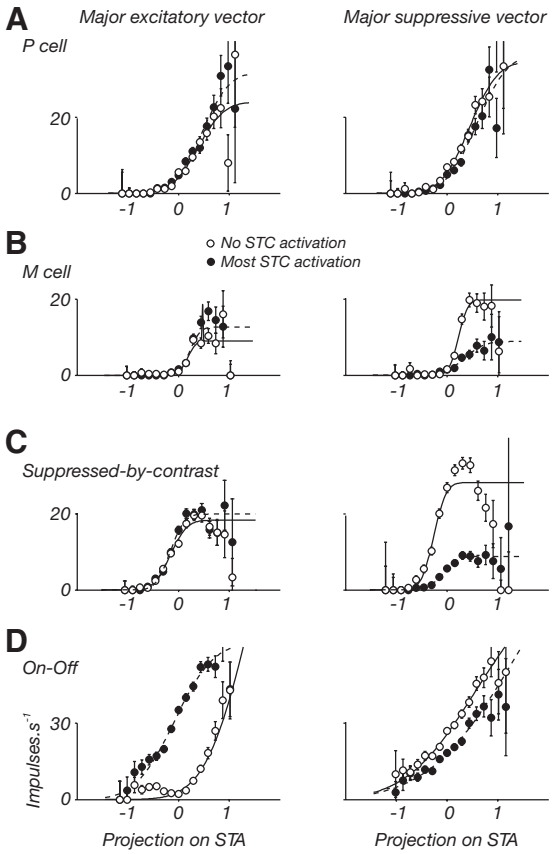

FIG. 6. Impact of excitatory and suppressive eigenvectors on the output of the linear receptive field. Same cells as in Figs. 2 and 3. Each panel shows the relationship between the average discharge rate and projection of the stimulus on the STA. *Left panels* compare this relationship for stimuli that had large projections on the major excitatory axis shown by STC analysis (filled symbols) and stimuli that had small projections on that axis (open symbols). *Right panels* show equivalent functions but for small and large projections on the most suppressive STC eigenvector. Lines show the best fitting cumulative normal distribution in each case. Error bars are ±binomial SD.

Impact of the nonlinear mechanisms on the linear ones

Although Fig. 5, *A–D,* shows that the influence of the STA and the eigenvectors is largely independent, it does not tell us how the eigenvectors exert an impact on the output of the linear receptive field. To examine this, we sorted the entire stimulus ensemble by the absolute magnitude of their projection onto an eigenvector, thereby grouping together positive and negative projections of similar amplitude. We derived six bins that provided equal numbers of stimuli in each. From the sequences in these bins, we calculated the discharge rate as a function of the projection of those sequences on the STA (as in Fig. 5). This gave us a family of input–output functions for the STA; each obtained at a relatively constant level of activation of the selected nonlinear filter, examples of which are seen in Fig. 6, *A–D*. In each case, the *left panel* shows functions for the most excitatory eigenvector and the *right panel* shows functions for the most suppressive eigenvector. For clarity, only two of the six functions are shown: that obtained for the smallest projections onto the eigenvector (open symbols) and that obtained for the largest (filled symbols).

The *left panels* show that, as expected, only in the ON-OFF cell does rate depend strongly on the projection of the stimulus onto the excitatory eigenvector—larger projections onto the eigenvector effectively shift the input–output function of the STA to the left.

The impact of the suppressive eigenvector was most pronounced in the suppressed-by-contrast cell and was less in the M-cell and ON-OFF cell; no effect was observed in the P-cell. In most cells, the input–output function of the STA, under minimal and maximal projections onto the suppressive eigenvector, seemed to be vertically scaled versions of each other. We quantified this impression by fitting cumulative normal distributions to the input–output functions, first allowing all parameters to vary with projection on the STC eigenvector and then repeating the fitting procedure but allowing only a vertical scale factor or a horizontal scale factor to change with projection (see METHODS). Curves were fit simultaneously to the full family of six input–output functions, but for clarity, we plot only the two extremes. In the nine M-cells with significantly suppressive vectors, allowing only vertical scaling increased the error between the prediction and the fit to 170% (SD 93) of that obtained when all parameters were allowed to vary. Horizontal scaling increased it to 558% (SD 463), which was significantly greater ($P < 0.05$, paired Student's *t*-test). In nine P-cells, these were 140 and 260% (SD 39 and 112; $P < 0.005$), and in nine suppressed-by-contrast cells, they were 129 and

1,651% (SD 31 and 840; $P < 0.005$). Similar results were obtained for blue-ON cells and the excitatory eigenvalue in P- and M-cells, but in ON-OFF cells, neither vertical or horizontal scaling alone adequately described the curves.

Figure 7 uses the vertical scale factors returned by the fitting procedure to show the impact of the significant suppressive and excitatory eigenvectors on the visual sensitivity of our population of LGN cells. This vertical scale factor is the response gain, and Fig. 7*A* compares the response gain of the STA for stimuli that were orthogonal, or nearly so, to the most suppressive eigenvector, with the response gain of the STA when stimuli were close to that eigenvector. It is clear that the suppressive eigenvectors in M-cells and suppressed-by-contrast cells are associated with a substantial reduction in response gain (and therefore in discharge rate also), suggesting they are important regulators of the output of the linear receptive field. Those of most P-cells are less clear, suggesting that in these cells the suppressive vectors have less impact on the output of the linear receptive field. Figure 7*B* shows a similar plot for the most excitatory eigenvector; except for some of the ON-OFF cells, the significant vectors generally have modest effect on the linear receptive field.

Impact of stimulus size on responses of P- and M-cells

It is well known that, in M-cells, the shape of the temporal filter depends on the stimulus, reflecting the spatiotemporal sensitivity of both the classical linear receptive field and an (nonlinear) extraclassical receptive field to that stimulus (Benardete and Kaplan 1999; Camp et al. 2009). One way to manipulate the contribution of the extraclassical field is to change the size of the stimulus (smaller stimuli should recruit less of it). If the nonlinear mechanisms shown by the STC analysis reflect the action of an extraclassical receptive field, they should be less prevalent when the stimulus is small. We therefore made measurements of responses to the same sequence of flickering fields for small (average 0.75° diam), and large (4–8°) fields in 21 P-cells (12 ON, 9 OFF) and 7 M-cells (3 ON, 4 OFF). The size of the smaller stimulus was our best estimate of the stimulus that encompassed most of the center and surround of the classical receptive field, yet minimally extended beyond it (the remainder of the screen was held at the mean luminance). We estimated this size using size-tuning curves obtained for both modulated fields and drifting gratings.

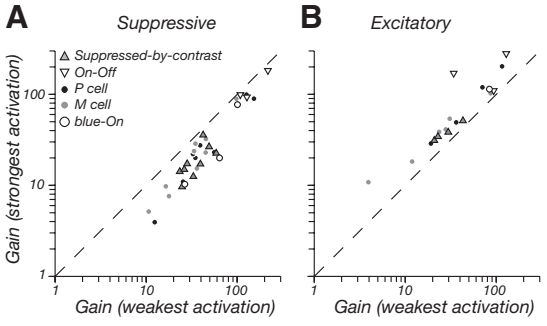

FIG. 7. Impact of excitatory and suppressive eigenvectors on the response gain of the STA. Cumulative normal functions were fit to the input–output function of the STA as in Fig. 6, with only the vertical scale factor (response gain) allowed to vary with projection on the STC vector. *A*: impact of most suppressive STC vector. Comparison of response gains obtained for stimuli with small projections on the STC vector (abscissa) and those with large projections (ordinate). *B*: same as *A*, but for the most excitatory STC vector.

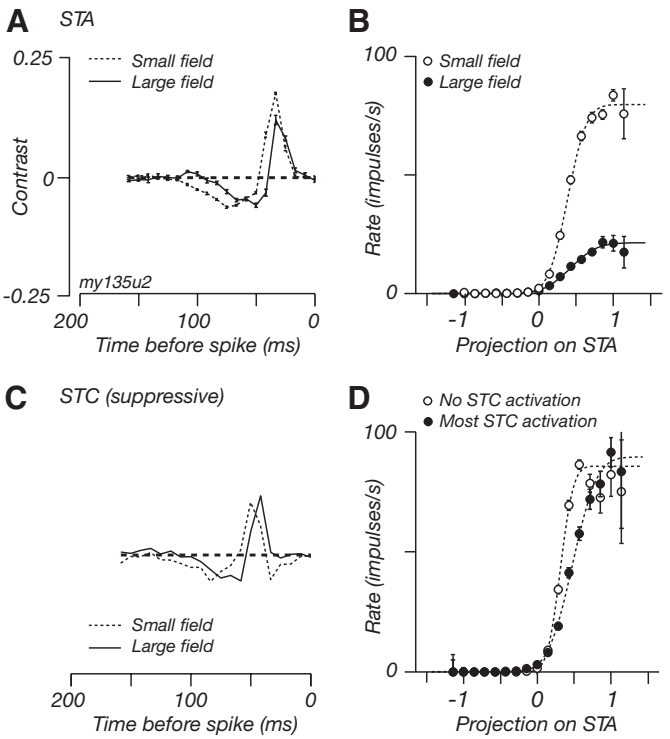

FIG. 8. Responses of an M-cell during stimulation with small and large fields. Same M-cell as Figs. 2, 3, 5, and 6. Measurements were made during stimulation with the same temporal sequence, confined to a small uniform field (diameter 1°) or a large field (diameter 4°). The remainder of the screen was held at the mean luminance. *A*: temporal profile of the STA. Error bars are ±SE. Same format as Fig. 2. *B*: input–output functions for the STAs shown in *A*. Error bars are ±binomial SD. *C*: temporal profile of the suppressive eigenvectors. One significant eigenvector was found during stimulation with a small field, and 1 was found during stimulation with a large field. *D*: input–output function for the STA for small and large projections on the most suppressive STC eigenvector obtained in small fields. The impact of the eigenvector is to change the slope (sensitivity) of the STA. Equivalent functions for the large field can be found in Fig. 6*B*. Error bars are ±binomial SD.

The dashed lines in Fig. 8, *A* and *B*, show the STA and its input–output function for the example M-ON cell, obtained using small fields; for comparison, the solid lines show those obtained with a large field. The STA is slightly broader in small fields, which probably reflects reduced action of a contrast gain control and reduced antagonism from the classical (linear) receptive field surround (Benardete and Kaplan 1999). The relationship between STA activation and spike rate is nevertheless similar up to a vertical scale factor.

Changing the size of the stimulus changed the discharge of cells during the flickering field: in the 7 M-cells, confining the stimulus to a smaller field increased the overall firing rate [from 5.5 (SD 3.1) in a large field to 21.7 (SD 7.2) impulses/s in a small field; $P < 0.001$, paired Student's *t*-test]; in the 21 P-cells, these were 9.4 (SD 4.6) and 11.8 (SD 5.6) impulses/s ($P = 0.08$). Furthermore, for both P- and M-cells, the information rate of the STA (bits/spike) increased when using small fields: in P-cells by 49% (from 0.41, SD 0.31 to 0.61, SD 0.50; $P < 0.01$) and in M-cells by 68% (from 0.77, SD 0.23 to 1.29, SD 0.17; $P < 0.01$). Because changing the size of the stimulus changed the discharge rate and the information per spike, the total information rate of the STA was a good deal larger in small fields: for M-cells, it increased from 4.1 to 27.1 bits/s; for P-cells, it increased from 3.2 to 7.8.

Accordingly, the STC analysis of responses to small fields produced somewhat different results to that obtained for large fields. Among the 21 P-cells tested with both, the STC showed significant eigenvalues in two cells when using small fields and in seven when using large fields. Among each of the three M-ON cells tested with small fields, we found significant suppressive axes, but these were very different to those found with large fields—for small sizes, the eigenvectors were broader and relatively more sensitive to lower frequencies (an example is shown in Fig. 8*C*); in the four M-OFF cells, the excitatory eigenvalues became more prominent when small fields were used (data not shown).

The way in which the eigenvectors of M-cells interacted with the STA also changed with stimulus size. Figure 8*D* shows, separately for stimuli with weak (open symbols) and strong (closed symbols) projections onto the suppressive eigenvector, the input–output function of the example cell obtained using small fields. For a small field, the change in the input–output function that the STC brings about is better described as scaling of the horizontal axis than as a scaling of the vertical axis: in other words, stimuli that have large projections onto the suppressive eigenvector change the slope of the input–output function but not its maximal rate. This is quite different to the impact of the suppressive eigenvector during stimulation with a large field, which is shown in the *right panel* of Fig. 6*B*. In seven M-cells, we obtained responses with both sizes: in four cells the small fields yielded significantly suppressive eigenvectors, and in these, the predictions of the horizontal scale model were always better than that of the vertical scale model [increasing the error in the fit to 124% (SD 18) vs. 186% (SD 66)]; in the same cells for large fields, the predictions of the horizontal scale model were always worse than those of the vertical scale model [481%, (SD 528) vs. 145% (SD 49)]. In summary, suppressive eigenvectors found during stimulation with large fields have a large impact on the response gain of the linear receptive field, whereas those found during stimulation with small fields have more of an impact on its sensitivity.

Reconstruction of stimulus luminance and energy

The characteristic receptive field properties of the very nonlinear cells make it unlikely that their signals are useful for reconstructing the spatial position or luminance polarity (increment, decrement) of the local image. It is more likely that they represent the presence or absence of spatiotemporal contrast. Such an early representation of contrast energy would be complementary to that of the largely linear receptive fields of P- and M-cells, where a lack of change in spike rate might represent either the presence of a gray field or any other pattern that is orthogonal to the receptive field of the cell. The analyses above show that there is information about contrast energy in the signals of the very nonlinear cells, but we do not know if these signals can be decoded by later processing. The simplest decoder is a linear one, and in the following, we use a well-established method of linear decoding (Bialek et al. 1991; Pillow et al. 2008; Warland et al. 1997) to quantify the information it can provide. Although the linear decoder does not allow for nonlinearities in retinal or LGN processing, it does allow us to place a lower bound on the information that might be extracted by later processing.

We used the spikes obtained during presentation of the flickering field to train the decoder, so that it learned the filter that best predicted the stimulus, given the train of spikes (it minimized the square error between the actual and predicted sequences). We assessed the performance of this decoder by trying to predict a separate sequence of frames on the basis of the sequence of spikes elicited by those frames.

Cells early in the visual pathway can provide substantially accurate reconstructions of the temporal profile of flickering fields like those we have used (see, e.g., Bialek et al. 1991; Passaglia and Troy 2004; Pillow et al. 2008; Warland et al. 1997). The *top panels* in Fig. 9, *A* and *B,* show a small segment of reconstruction obtained from one P-ON cell and one M-ON cell in response to a small flickering field. The gray line shows the luminance profile of the stimulus; the thicker black line shows the reconstruction of this stimulus sequence from the trained decoder. The reconstruction does not always faithfully track the actual stimulus, especially during periods of very slow or very rapid flicker, but there is obvious congruence

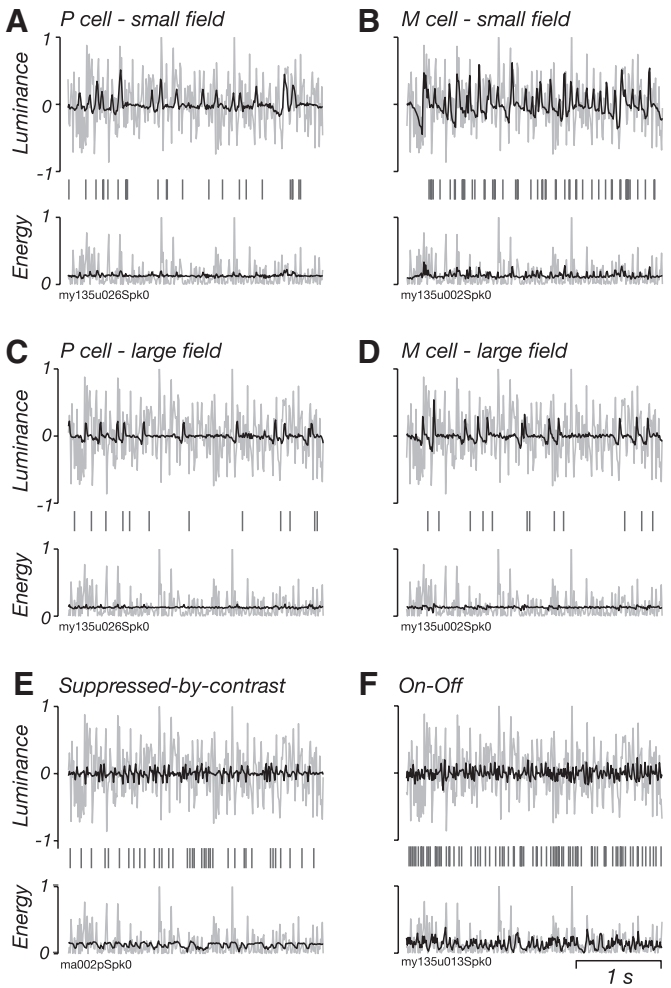

FIG. 9. Decoding of luminance and energy by LGN neurons. *A, C*: P cell. *B, D*: M cell. *E*: suppressed-by-contrast cell. *F*: ON-OFF cell. Same cells as in Figs. 2, 3, and 6. *Top panels* in each of *A–F* show the luminance profile of the stimulus (thin line) and that of the reconstruction (thick line) predicted by an optimal linear decoder trained on responses to separate periods of stimulation. *Bottom panels* show the energy (the square of the luminance) of the stimulus (thin line) and that of the reconstruction (thick lines). Vertical dashes between the plots indicate the time of arrival of action potentials.

between the two sequences. Figure 9, *C* and *D,* shows for the same cells, the reconstruction during flicker of the same temporal sequence presented within a larger field. The discharge rate is lower, and the reconstruction is considerably worse.

The *top panels* in Figs 9, *E* and *F,* show luminance reconstructions obtained from one suppressed-by-contrast cell and one ON-OFF cell. The reconstructions less clearly track the stimulus, as expected given that the algorithm is trying to reconstruct the luminance of the stimulus, and we know that the cell responds in a similar way to both increments and decrements. We therefore also examined the quality of reconstructions obtained by squaring the stimulus before training the decoder. The signals of the nonlinear cells were much more useful in decoding this contrast energy; this is clear in the *bottom panels* in Fig. 9, *E,* and *F*, where the reconstruction tracks the stimulus. In this task, the signals of the P- and M-cells (Fig. 9, *A–D, bottom panels*) were less useful, which is not surprising because they respond quite differently to luminance increments and decrements that nevertheless have the same energy.

Quality of reconstructions for luminance and energy

To quantify the correlation between the actual stimulus and its reconstruction we expressed the difference between them as the signal-to-noise ratio at each temporal frequency (Warland et al. 1997) (see METHODS). Figure 10*A* shows this analysis for the luminance reconstructions obtained for large fields. As expected from the shape of their STA, over the population, P-cells gave generally better reconstructions at low frequencies and M-cells gave better reconstructions at high frequencies. The ON-OFF cell and the suppressed-by-contrast cell performed poorly at low frequencies but at high frequencies matched the performance of P- and M-cells. Figure 10*B* shows equivalent plots for the same cells but for reconstructions of the stimulus energy. Those provided by the P-cell and M-cell were poor; the reconstructions provided by the ON-OFF cell and suppressed-by-contrast cell were much better.

The integral of the signal-to-noise ratio across a linear scale of temporal frequency is an estimate of the mutual information between the reconstruction and the actual stimulus. These information rates (in bits/s) are plotted in Fig. 10*C*, which compares for each cell the information provided about luminance with that provided about energy. For P- and M-cells, the information rates for energy are much less than those about luminance. This is not the case for the ON-OFF cells and particularly for the suppressed-by-contrast cells, both of which provide similar information rates for energy and luminance. As is expected from the reconstructions in Fig. 9, the decoder recovered more information about luminance from P- and M-cells when small sizes were used (data not shown). For P-cells, the average information increased from 1.30 bits/s (SD 0.94) in large sizes to 3.19 in small (SD 3.35, $n = 21$; $P < 0.01$, paired Students *t*-test); for M-cells, the impact of size was more pronounced, increasing the information rate from 1.51 (SD 0.81) to 9.77 (SD 2.42, $n = 7$; $P < 0.0001$).

To characterize the temporal resolution of the reconstructions, we found the point on the high temporal frequency limb at which the reconstruction had fallen to one half its maximal signal-to-noise ratio (Fig. 10*D*). For luminance modulations, P-cells resolved on average slightly lower frequencies than

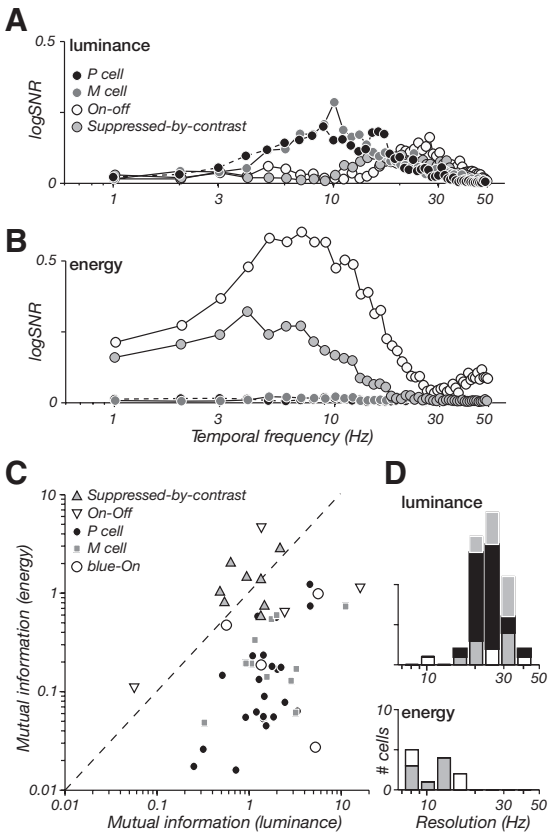

FIG. 10. Signals for luminance and energy in LGN cells. *A*: fidelity of luminance reconstructions for each of the neurons in Fig. 9. The ordinate shows the base–2 logarithm of the signal-to-noise ratio (logSNR) at each temporal frequency. *B*: same as *A*, but for reconstructions of energy. *C*: comparison of mutual information rates (in bits/s) provided by individual cells for luminance and energy. Information rates were determined as one half the integral of the logSNRs shown in *A* and *B*. Average information rates for luminance: P-cells: μ 1.52, SD 1.11; M-cells 2.66, SD 2.98; suppressed-by-contrast cells 1.02, SD 0.61. Average information rates for energy: P-cells: μ 0.18, SD 0.28; M-cells 0.27, SD 0.23; suppressed-by-contrast cells 1.30, SD 0.72. *D*: temporal frequency at which the high-frequency limb of the logSNR curves declined to one half the maximal logSNR for luminance (*top*) and energy (*bottom*).

M-cells (averaging 24.8 and 28.8, respectively; $P < 0.05$, Student's *t*-test), but the difference was small. We found no difference between ON and OFF cells of each class. Where we could characterize them, the nonlinear cells were not distinguished from P- and M-cells on the basis of frequency resolution of luminance reconstructions. The energy reconstructions of ON-OFF and suppressed-by-contrast cells were strong over much lower temporal frequencies (resolving at 12.8 and 10.6 Hz, respectively) than the luminance reconstructions provided by P- and M-cells.

In summary, Figs. 9 and 10 confirm that a linear decoder can recover information about the field's luminance from P- and M-cells. The same decoder can capture substantial information about its contrast energy from ON-OFF and suppressed-by-contrast cells.

DISCUSSION

Nonlinear mechanisms in P- and M-cells

The receptive fields of P-cells are well characterized by a linear model (Benardete and Kaplan 1997a, b; Benardete et al.

1992). The analysis of covariance here shows that some P-cells nevertheless require additional nonlinearities, particularly suppressive ones, and that when present the suppressive mechanisms of P-ON and P-OFF cells differ in temporal sensitivity (Fig. 4).

The receptive fields of most M-cells exhibited stronger nonlinearities, and the mechanisms of suppression differed between M-ON and M-OFF cells, providing further evidence for functional asymmetries between the ON- and OFF-pathways (Chichilnisky and Kalmar 2002; Zaghloul et al. 2003). We note that the temporal structure along many of these suppressive axes resembled the STA, shifted in time. That they resemble the STA is not surprising, because nearly all retinal ganglion cells show biphasic impulse response functions (see Fig. 4), but it might also suggest that the suppressive axes simply reflect non-Poisson properties of neurons, such as burst responses or a refractory period, which causes a brief reduction in discharge rate after each spike (Aguera y Arcas and Fairhall 2003; Fairhall et al. 2006; Pillow and Simoncelli 2003) We think that this cannot provide a complete explanation for the suppressive axes, because reducing the size of the stimulus led to an increase in the average discharge rate in both P- and M-cells but did not increase the number of cells showing suppressive axes or the number of suppressive axes in cells.

During stimulation with small fields, the influence of the suppressive mechanisms on the output of the STA could be reasonably described by a change in the sensitivity of the output of the STA, the impact that is expected of a contrast gain control (Chander and Chichilnisky 2001; Shapley and Victor 1979). In large fields, the major effect of the suppressive mechanisms was instead to divide the evoked discharge rate by a constant. This is consistent with previous work in marmoset LGN, which shows that stimuli placed in a region around the classical receptive field will bring about a reduction in response gain (Camp et al. 2009; Felisberti and Derrington 2001; Solomon et al. 2002; Webb et al. 2002). The mechanism that gives rise to the sensitivity reduction that is evident during stimulation with small fields is thus either unresponsive in large fields or its functional signature is overwhelmed by the reduction in response gain; regardless, it is clear that under most visual conditions, responsivity is modulated by multiple suppressive mechanisms with different functional characteristics (Camp et al. 2009; Zaghloul et al. 2007).

Because response includes the discharge generated over all projections onto the STA, large fields reduced the discharge rate of M-cells even when there was no drive to the linear receptive field. We therefore expect that the maintained discharge rate can be reduced by peripheral stimulation even when there is no effective stimulus over the classical receptive field. This effect has been observed in previous studies of retina and LGN: stimulation with remote, contrast reversing patterns reduces the maintained discharge of M-cells (Dhruv et al. 2009; Solomon et al. 2006). Although many studies suggest a retinal origin for some of the suppressive effect, we note that the STC analysis reported here might reflect mechanisms that arise within the LGN (Dhruv et al. 2009; Hubel and Wiesel 1966; Levick et al. 1972; Murphy and Sillito 1987; Webb et al. 2002). We did not make recordings from retinal afferents, so we do not know which components are already established in the retina (Benardete and Kaplan 1999; Enroth-Cugell and Jakiela 1980; Passaglia et al. 2001, 2009; Sincich et al. 2009; Solomon et al. 2006). Recent work in macaque, in which retinal afferents and their LGN targets were recorded simultaneously, may help (Sincich et al. 2009). In both P- and M-cells, mechanisms much like those shown by the STC analyses here were already apparent in the retinal signals: the primary transformation of visual signals in LGN was conferred by the temporal summation of retinal afferent spikes. However, that work explored activity during temporal modulation of a small spot confined to the classical receptive field, which is unlikely to engage strong suppressive mechanisms in LGN, and therefore we do not know if processing within LGN provides additional mechanisms.

Nonlinear mechanisms in uncommon receptive fields

Suppressed-by-contrast cells and ON-OFF cells have receptive fields that are by definition very nonlinear, and these nonlinearities seem to have been captured by the STC analysis. Suppressed-by-contrast cells in particular showed strong suppressive eigenvectors, whereas ON-OFF cells usually had the most complicated of the receptive fields we studied, displaying both strong excitatory and suppressive eigenvectors. Unfortunately, although STC analysis recovers a representation of the spike-triggered stimulus ensemble using a set of orthogonal filters, these eigenvectors are not guaranteed to be the fundamental axes of that subspace. Because of this, it is not straightforward to relate the profile of the eigenvectors to the underlying neural machinery (Rust et al. 2005; Schwartz et al. 2006).

Suppressed-by-contrast cells are known to exist in the retina (Cleland and Levick 1974; Rodieck 1967; Tailby et al. 2007), so it is likely that the STC analyses show mechanisms in place there. Similarly, the mechanisms that provide the excitatory eigenvalues in ON-OFF cells may be those that confer the "Y-like" frequency-doubled responses that have been extensively studied in cat retina (Enroth-Cugell and Robson 1966; Hochstein and Shapley 1976), as well as similar cell types in the retina of salamanders (Fairhall et al. 2006), guinea pig (Demb et al. 1999), and primates (Crook et al. 2008a, b; Petrusca et al. 2007). Where studied, the nonlinear responses of these cell types are thought to reflect rectification at the bipolar cell synapse (Demb et al. 2001). Figure 4 shows that, in ON-OFF cells, the excitatory eigenvectors are sensitive to very high temporal frequencies, consistent with the idea that the nonlinearity is early in visual processing.

The presence of an STA in suppressed-by-contrast cells was surprising, because there are few indications of stimuli that increase the discharge rate above that for a blank screen (Rodieck 1967; Tailby et al. 2007). As with the suppressive axes, the STA was usually quite transient, and the input–output function often nonmonotonic; if suppression arises in half-wave rectified mechanisms sensitive to increments or decrements, the STA may reflect an imbalance in the strength of these mechanisms.

Functional characteristics of parallel pathways through the LGN

We used achromatic stimuli to characterize the linear and nonlinear mechanisms in all classes of LGN cells. For step changes in achromatic contrast, the dominant mode of response for suppressed-by-contrast cells was a transient

suppression of the discharge at each change in contrast (see examples in Figs. 1*B* and 11). However, many of the suppressed-by-contrast cells responded in a strikingly different way when the step changes in contrast were largely restricted to one class of cone photoreceptor. Figure 11 shows the response of a representative suppressed-by-contrast cell: the discharge rate was low during the entire ON-phase of S cone modulation but not the OFF-phase. Strong transient suppression at both phases of LM cone modulation was accompanied by a sustained reduction in rate during the OFF-phase. We saw the same patterns during cone-isolating modulation in seven of the eight suppressed-by-contrast cells tested. For the cell in Fig. 11, the responses during achromatic modulation were reasonably well described by simply summing the two cone-isolating responses but that was not always the case. This suggests that the suppressed-by-contrast cell gets sustained input from a half-wave rectified "+S-LM" cone-opponent mechanism. An additional transient suppression must arise in other mechanisms that are dominated by LM cone input. We do not know if this convergence arises in the retina or thalamus.

The response of the suppressed-by-contrast cell during achromatic modulation is different to that of blue-ON or blue-OFF cells that we have encountered previously, where achromatic and cone-isolating modulation bring about changes in discharge that are reasonably symmetric around the maintained rate (Tailby et al. 2008a). In LGN, both blue-ON cells and blue-OFF cells (much like P- and M-cells) can be suppressed by achromatic stimuli presented in regions outside the classical receptive field (Solomon et al. 2002; Tailby et al. 2008b); we may therefore expect to find the kinds of eigenvectors that are expressed in the P- and M-cells in blue-ON and blue-OFF cells also. In one of the four blue-ON cells encountered here, the suppressive filters were as strong and transient as those in M-cells or suppressed-by-contrast cells; however, the responses of this cell to step contrast changes showed no signs of transient suppression. In three of the ON-OFF cells, we saw transient excitatory responses to both phases of S cone modulation, similar in shape but much smaller than the response to achromatic modulation (we cannot exclude the possibility that this may simply reflect residual LM cone modulation).

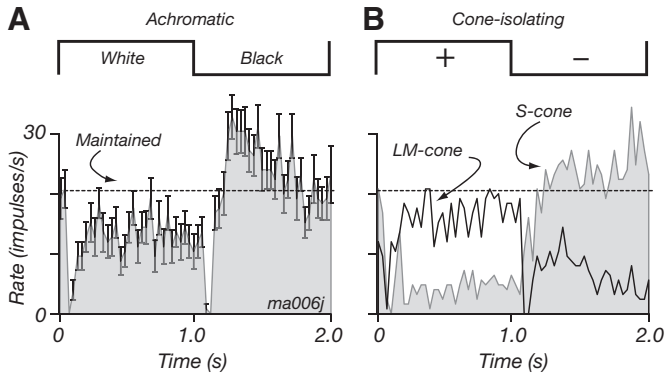

FIG. 11. Cone inputs to a suppressed-by-contrast cell. Same cell as Fig. 1*B*. Each panel shows the average PSTH accumulated during square-wave temporal modulation of a large uniform field. *A*: response to achromatic modulation (dashed line shows the maintained discharge): activity is transiently suppressed at steps to white or black. *B*: response to modulations substantially restricted to the S-cone or to the LM-cone. LM-cone modulation induces transient suppression at both phases and a sustained suppression during the OFF phase. S-cone modulation leads to a reduction in the discharge during the ON phase. Error bars are ±SE.

It is often thought that because the sampling density of S cones is too low to support high acuity spatial vision, S cone input to a receptive field is indicative of a role in color vision. It is not yet clear to us if the functional diversity shown above is best characterized by multiple subclasses of blue-ON and blue-OFF cells or by diversity in the expression of S cone input to the receptive fields of the nonlinear cells. We favor the latter: whereas the functional properties of the very nonlinear cells do not rule out a role in color vision, such a role would require reconciling their very nonlinear responses with those of more linear cells, particularly in trichromats where their signals need to be combined with those of red-green opponent cells in the P-pathway. However, an alternative functional role for the substantial input from S cones remains unclear; whatever the role, we suspect it may point to an early evolutionary origin of S cones pathways (Mollon 1989) and the very nonlinear pathways, both of which are found in primates and other mammals.

Linear decoding of stimulus features

A simple linear decoder was able to reconstruct the time course of either luminance or contrast energy by selectively drawing on the spikes of different cell classes. ON-OFF and suppressed-by-contrast cells provide information about contrast energy that is comparable to the information that P- and M-cells provide about luminance. These information estimates are lower bounds. We expect more information could be recovered from the spike trains of all cell classes by decoders that worked with more complicated combination rules, considered simultaneously recorded spike trains from neighboring cells, and took into account higher-level features of the spike train (Lundstrom and Fairhall 2006; Passaglia and Troy 2004; Pillow et al. 2008). Nevertheless, our analysis shows that even with the simplest decoder, the cortex may have access to a reasonably high-fidelity representation of contrast energy as well as luminance; this might be a more efficient strategy than superimposing the two signals onto the spike trains of a single pathway (Balasubramanian and Sterling 2009). We speculate that such an early code for contrast energy might be useful in initiating coarse bottom-up maps for salience (Itti et al. 1998) or in providing a neural measure of contrast energy against which the responses of other cells can be referenced (Heeger 1992; Tailby et al. 2007; Troy et al. 1989). The distinctiveness of these reconstructions of contrast energy might help distinguish their contributions to subsequent processing or behavior—signals were robust at temporal frequencies up to ~15 Hz, which is much lower than the frequencies to which the luminance reconstructions of P- and M-cells extended.

ACKNOWLEDGMENTS

We thank P. R. Martin for assistance in experiments and anonymous reviewers for constructive criticism and comments.

GRANTS

This work was supported by National Health and Medical Research Council of Australia Grants 457337 and 511967.

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the authors.

REFERENCES

Adelman TL, Bialek W, Olberg RM. The information content of receptive fields. *Neuron* 40: 823–833, 2003.

Aguera Y Arcas B, Fairhall AL. What causes a neuron to spike? *Neural Comp* 15: 1789–1807, 2003.

Balasubramanian V, Sterling P. Receptive fields and finctional architecture in the retina. *J Physiol* 587: 2753–2767, 2009.

Benardete EA, Kaplan E. The dynamics of primate M retinal ganglion cells. *Vis Neurosci* 16: 355–368, 1999.

Benardete EA, Kaplan E. The receptive field of the primate P retinal ganglion cell, I: linear dynamics. *Vis Neurosci* 14: 169–185, 1997a.

Benardete EA, Kaplan E. The receptive field of the primate P retinal ganglion cell, II: nonlinear dynamics. *Vis Neurosci* 14: 187–205, 1997b.

Benardete EA, Kaplan E, Knight BW. Contrast gain control in the primate retina: P cells are not X-like, some M cells are. *Vis Neurosci* 8: 483–486, 1992.

Bialek W, Rieke F, de Ruyter van Steveninck RR, Warland D. Reading a neural code. *Science* 252: 1854–1857, 1991.

Blessing EM, Solomon SG, Hashemi-Nezhad M, Morris BJ, Martin PR. Chromatic and spatial properties of parvocellular cells in the lateral geniculate nucleus of the marmoset (*Callithrix jacchus*). *J Physiol* 557: 229–245, 2004.

Brenner N, Bialek W, de Ruyter van Steveninck R. Adaptive rescaling maximizes information transmission. *Neuron* 26: 695–702, 2000.

Camp AJ, Tailby C, Solomon SG. Adaptable mechanisms that regulate the contrast response of neurons in the primate lateral geniculate nucleus. *J Neurosci* 29: 5009–5021, 2009.

Casagrande VA. A third parallel visual pathway to primate area V1. *Trends Neurosci* 17: 305–310, 1994.

Chander D, Chichilnisky EJ. Adaptation to temporal contrast in primate and salamander retina. *J Neurosci* 21: 9904–9916, 2001.

Chen X, Han F, Poo MM, Dan Y. Excitatory and suppressive receptive field subunits in awake monkey primary visual cortex (V1). *Proc Natl Acad Sci USA* 104: 19120–19125, 2007.

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

Chichilnisky EJ, Kalmar RS. Functional asymmetries in ON and OFF ganglion cells of primate retina. *J Neurosci* 22: 2737–2747, 2002.

Cleland BG, Levick WR. Properties of rarely encountered types of ganglion cells in the cat's retina and an overall classification. *J Physiol* 240: 457–492, 1974.

Crook JD, Davenport CM, Peterson BB, Packer OS, Detwiler PB, Dacey DM. Parallel ON and OFF cone bipolar inputs establish spatially coextensive receptive field structure of blue-yellow ganglion cells in primate retina. *J Neurosci* 29: 8372–8387, 2009.

Crook JD, Peterson BB, Packer OS, Robinson FR, Gamlin PD, Troy JB, Dacey DM. The smooth monostratified ganglion cell: evidence for spatial diversity in the Y-cell pathway to the lateral geniculate nucleus and superior colliculus in the macaque monkey. *J Neurosci* 28: 12654–12671, 2008a.

Crook JD, Peterson BB, Packer OS, Robinson FR, Troy JB, Dacey DM. Y-cell receptive field and collicular projection of parasol ganglion cells in macaque monkey retina. *J Neurosci* 28: 11277–11291, 2008b.

Dacey DM, Peterson BB, Robinson FR, Gamlin PD. Fireworks in the primate retina: in vitro photodynamics reveals diverse LGN-projecting ganglion cell types. *Neuron* 37: 15–27, 2003.

de Monasterio FM. Properties of ganglion cells with atypical receptive-field organization in retina of macaques. *J Physiol* 41: 1435–1449, 1978.

de Ruyter van Steveninck R, Bialek W. Real-time performance of a movement sensitive neuron in the blowfly visual system: coding and information transfer in short spike sequences. *Proc R Soc Lond Ser B* 234: 379–414, 1988.

Demb JB, Haarsma L, Freed MA, Sterling P. Functional circuitry of the retinal ganglion cell's nonlinear receptive field. *J Neurosci* 19: 9756–9767, 1999.

Demb JB, Zaghloul K, Haarsma L, Sterling P. Bipolar cells contribute to nonlinear spatial summation in the brisk-transient (Y) ganglion cell in mammalian retina. *J Neurosci* 21: 7447–7454, 2001.

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

Derrington AM, Lennie P. Spatial and temporal contrast sensitivities of neurones in lateral geniculate nucleus of macaque. *J Physiol* 357: 219–240, 1984.

Dhruv NT, Tailby C, Sokol SH, Majaj NJ, Lennie P. Nonlinear signal summation in magnocellular neurons of the macaque lateral geniculate nucleus. *J Neurophysiol* 102: 1921–1929, 2009.

Dreher B, Fukada Y, Rodieck RW. Identification, classification and anatomical segregation of cells with X-like and Y-like properties in the lateral geniculate nucleus of Old World primates. *J Physiol* 258: 433–452, 1976.

Enroth-Cugell C, Jakiela HG. Suppression of cat retinal ganglion cell responses by moving patterns. *J Physiol* 302: 49–72, 1980.

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

Fairhall AL, Burlingame CA, Narasimhan R, Harris RA, Puchalla JL, Berry MJ II. Selectivity for multiple stimulus features in retinal ganglion cells. *J Neurophysiol* 96: 2724–2738, 2006.

Felisberti F, Derrington AM. Long-range interactions in the lateral geniculate nucleus of the New-World monkey, *Callithrix jacchus*. *Vis Neurosci* 18: 209–218, 2001.

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

Hochstein S, Shapley RM. Linear and nonlinear spatial subunits in Y cat retinal ganglion cells. *J Physiol* 262: 265–284, 1976.

Horwitz GD, Chichilnisky EJ, Albright TD. Blue-yellow signals are enhanced by spatiotemporal luminance contrast in macaque V1. *J Neurophysiol* 93: 2263–2278, 2005.

Hubel DH, Wiesel TN. Effects of varying stimulus size and color on single lateral geniculate cells in Rhesus monkeys. *Proc Natl Acad Sci USA* 55: 1345–1346, 1966.

Itti L, Koch C, Niebur E. A model of saliency-based visual attention for rapid scene analysis. *IEEE T Pattern Anal* 20: 1254–1259, 1998.

Kaplan E, Shapley RM. The primate retina contains two types of ganglion cells, with high and low contrast sensitivity. *Proc Natl Acad Sci USA* 83: 2755–2757, 1986.

Kaplan E, Shapley RM. X and Y cells in the lateral geniculate nucleus of macaque monkeys. *J Physiol* 330: 125–143, 1982.

Lankheet MJ, Lennie P, Krauskopf J. Temporal-chromatic interactions in LGN P-cells. *Vis Neurosci* 15: 47–54, 1998.

Leventhal AG, Rodieck RW, Dreher B. Retinal ganglion cell classes in the Old World monkey: morphology and central projections. *Science* 213: 1139–1142, 1981.

Levick WR, Cleland BG, Dubin MW. Lateral geniculate neurons of cat: retinal inputs and physiology. *Invest Ophthalmol Vis Sci* 11: 302–311, 1972.

Levitt JB, Schumer RA, Sherman SM, Spear PD, Movshon JA. Visual response properties of neurons in the LGN of normally reared and visually deprived macaque monkeys. *J Neurophysiol* 85: 2111–2129, 2001.

Lundstrom BN, Fairhall AL. Decoding stimulus variance from a distributional neural code of interspike intervals. *J Neurosci* 26: 9030–9037, 2006.

Mollon JD. Tho' she kneel'd in that place where they grew. (The uses and origins of primate color vision.). *J Exp Biol* 146: 21–38, 1989.

Murphy P, Sillito AM. Corticofugal feedback influences the generation of length tuning in the visual pathway. *Nature* 329: 727–729, 1987.

Passaglia CL, Enroth-Cugell C, Troy JB. Effects of remote stimulation on the mean firing rate of cat retinal ganglion cells. *J Neurosci* 21: 5794–5803, 2001.

Passaglia CL, Freeman DK, Troy JB. Effects of remote stimulation on the modulated activity of cat retinal ganglion cells. *J Neurosci* 29: 2467–2476, 2009.

Passaglia CL, Troy JB. Information transmission rates of cat retinal ganglion cells. *J Neurophysiol* 91: 1217–1229, 2004.

Petrusca D, Grivich MI, Sher A, Field GD, Gauthier JL, Greschner M, Shlens J, Chichilnisky EJ, Litke AM. Identification and characterization of a Y-like primate retinal ganglion cell type. *J Neurosci* 27: 11019–11027, 2007.

Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, Simoncelli EP. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* 454: 995–999, 2008.

Pillow JW, Simoncelli EP. Dimensionality reduction in neural models: an information-theoretic generalization of spike-triggered average and covariance analysis. *J Vis* 6: 414–428, 2006.

Pillow JW, Simoncelli EP. Biases in white noise analysis due to non-Poisson spike generation. *Neurocomputing* 52–54 109–115, 2003.

Rodieck RW. Receptive fields in the cat retina: a new type. *Science* 157: 90–92, 1967.

Rodieck RW, Watanabe M. Survey of the morphology of macaque retinal ganglion cells that project to the pretectum, superior colliculus, and parvicellular laminae of the lateral geniculate nucleus. *J Comp Neurol* 338: 289–303, 1993.

Rust NC, Schwartz O, Movshon JA, Simoncelli EP. Spatiotemporal elements of macaque v1 receptive fields. *Neuron* 46: 945–956, 2005.

Schwartz O, Pillow JW, Rust NC, Simoncelli EP. Spike-triggered neural characterization. *J Vis* 6: 484–507, 2006.

Shapley R, Victor JD. The contrast gain control of the cat retina. *Vision Res* 19: 431–434, 1979.

Sharpee T, Rust NC, Bialek W. Analyzing neural responses to natural signals: maximally informative dimensions. *Neural Comp* 16: 223–250, 2004.

Sincich LC, Horton JC, Sharpee TO. Preserving information in neural transmission. *J Neurosci* 29: 6207–6216, 2009.

Solomon SG, Lee BB, Sun H. Suppressive surrounds and contrast gain in magnocellular-pathway retinal ganglion cells of macaque. *J Neurosci* 26: 8715–8726, 2006.

Solomon SG, White AJ, Martin PR. Extraclassical receptive field properties of parvocellular, magnocellular, and koniocellular cells in the primate lateral geniculate nucleus. *J Neurosci* 22: 338–349, 2002.

Szmajda BA, Grunert U, Martin PR. Retinal ganglion cell inputs to the koniocellular pathway. *J Comp Neurol* 510: 251–268, 2008.

Tailby C, Solomon SG, Dhruv NT, Majaj NJ, Sokol SH, Lennie P. A new code for contrast in the primate visual pathway. *J Neurosci* 27: 3904–3909, 2007.

Tailby C, Solomon SG, Lennie P. Functional asymmetries in visual pathways carrying S-cone signals in macaque. *J Neurosci* 28: 4078–4087, 2008a.

Tailby C, Szmajda BA, Buzas P, Lee BB, Martin PR. Transmission of blue (S) cone signals through the primate lateral geniculate nucleus. *J Physiol* 586: 5947–5967, 2008b.

Touryan J, Lau B, Dan Y. Isolation of relevant visual features from random stimuli for cortical complex cells. *J Neurosci* 22: 10811–10818, 2002.

Troy JB, Einstein G, Schuurmans RP, Robson JG, Enroth-Cugell C. Responses to sinusoidal gratings of two types of very nonlinear retinal ganglion cells of cat. *Vis Neurosci* 3: 213–223, 1989.

van Hateren JH, Ruttiger L, Sun H, Lee BB. Processing of natural temporal stimuli by macaque retinal ganglion cells. *J Neurosci* 22: 9945–9960, 2002.

Warland DK, Reinagel P, Meister M. Decoding visual information from a population of retinal ganglion cells. *J Neurophysiol* 78: 2336–2350, 1997.

Webb BS, Tinsley CJ, Barraclough NE, Easton A, Parker A, Derrington AM. Feedback from V1 and inhibition from beyond the classical receptive field modulates the responses of neurons in the primate lateral geniculate nucleus. *Vis Neurosci* 19: 583–592, 2002.

White AJR, Solomon SG, Martin PR. Spatial properties of receptive fields in the lateral geniculate nucleus of the marmoset *Callithrix jacchus. J Physiol* 533: 519, 2001.

Zaghloul KA, Boahen K, Demb JB. Different circuits for ON and OFF retinal ganglion cells cause different contrast sensitivities. *J Neurosci* 23: 2645–2654, 2003.

Zaghloul KA, Manookin MB, Borghuis BG, Boahen K, Demb JB. Functional circuitry for peripheral suppression in Mammalian y-type retinal ganglion cells. *J Neurophysiol* 97: 4327–4340, 2007.