

Preserving Information in Neural Transmission

Lawrence C. Sincich,¹ Jonathan C. Horton,¹ and Tatyana O. Sharpee^{2,3}

¹Beckman Vision Center, University of California, San Francisco, California 94143, ²Computational Neurobiology Laboratory, Salk Institute for Biological Studies, La Jolla, California 92037, and ³Center for Theoretical Biological Physics, University of California, San Diego, La Jolla, California 92037

Along most neural pathways, the spike trains transmitted from one neuron to the next are altered. In the process, neurons can either achieve a more efficient stimulus representation, or extract some biologically important stimulus parameter, or succeed at both. We recorded the inputs from single retinal ganglion cells and the outputs from connected lateral geniculate neurons in the macaque to examine how visual signals are relayed from retina to cortex. We found that geniculate neurons re-encoded multiple temporal stimulus features to yield output spikes that carried more information about stimuli than was available in each input spike. The coding transformation of some relay neurons occurred with no decrement in information rate, despite output spike rates that averaged half the input spike rates. This preservation of transmitted information was achieved by the short-term summation of inputs that geniculate neurons require to spike. A reduced model of the retinal and geniculate visual responses, based on two stimulus features and their associated nonlinearities, could account for >85% of the total information available in the spike trains and the preserved information transmission. These results apply to neurons operating on a single time-varying input, suggesting that synaptic temporal integration can alter the temporal receptive field properties to create a more efficient representation of visual signals in the thalamus than the retina.

Introduction

To completely characterize the signal transformation performed by an individual neuron, all of its inputs and its output spike train need to be recorded simultaneously. Because the number of cells providing input to any one neuron is quite high in most neural systems, the opportunities for making such recordings are rare, particularly in vertebrates. We capitalized on the fact that the receptive field center of many relay neurons in the lateral geniculate nucleus (LGN) of the primate thalamus receive their primary input from one ganglion cell in the retina (Cleland et al., 1971; Kaplan and Shapley, 1984; Sincich et al., 2007). The synapses formed by each ganglion cell axon are large, numerous, and occupy the proximal dendrites of LGN neurons (Conley and Fitzpatrick, 1989; Wilson, 1989). This anatomical arrangement allows the EPSPs generated by each ganglion cell spike to be recorded with an extracellular electrode (Hubel and Wiesel, 1961; Bishop et al., 1962; Kaplan and Shapley, 1984). Thus, in principle, it is possible to recreate the spike trains occurring on both sides of the synapse from such recordings.

The macaque visual system offers several experimental advan-

tages for examining the signal transformation of inputs to outputs. The receptive field centers of primate ganglion cells and LGN neurons dominate the responses to visual stimuli, and they have essentially no spatial structure (Reid and Shapley, 1992, 2002; Croner and Kaplan, 1995; Chichilnisky and Kalmar, 2002). These neurons also exhibit a range of temporal dynamics that may be stimulus dependent (Smirnakis et al., 1997; Usrey et al., 1999; Chander and Chichilnisky, 2001; Solomon et al., 2004). Therefore, to understand the simplest case of neuron-to-neuron signal transformation, we focused on how the representation of a flickering spot of light covering the field center is altered between a ganglion cell and its target LGN neuron. To approximate the spectrum of temporal frequencies that are encountered in natural environments, we constructed a noise stimulus to mimic both the non-Gaussian distribution and the second-order correlations of luminance values observed during natural viewing (Dong and Atick, 1995; Dan et al., 1996; van Hateren, 1997; Simoncelli and Olshausen, 2001; Ringach et al., 2002). We found that LGN spikes represented such stimuli differently than the input ganglion cell spikes, and did so with increased information efficiency in each spike, sometimes to the extent that lossless information transmission was achieved.

Materials and Methods

Physiology and EPSP identification. Experiments were conducted in 6 adult macaques using procedures approved by the UCSF Institutional Animal Care and Use Committee, and in accordance with National Institutes of Health guidelines. Animals were prepared for physiological recordings as described previously (Sincich et al., 2007), and the recordings reported here are the subset ($n = 9$) of the 37 cells with retinal EPSPs in that study which had sufficient data for the extended analysis detailed below. All cells appeared to fire in “tonic” mode, because <4% of the spikes could be classified as “burst” by spike-interval criteria; however, even these spikes had associated EPSPs (see Sincich et al., 2007 for supporting evidence). Extracellular potentials recorded by tungsten micro-

Received Aug. 5, 2008; revised March 22, 2009; accepted March 27, 2009.

This work was supported by National Science Foundation (NSF) Grant IIS-0712852 (L.C.S. and T.O.S.), National Eye Institute Grants EY10217 (J.C.H.) and EY02162 (Beckman Vision Center), National Institute of Mental Health Grant MH068904 (T.O.S.), an Alfred P. Sloan research fellowship (T.O.S.), a Searle scholarship (T.O.S.), and Research to Prevent Blindness (J.C.H.). Computing resources were provided by the following NSF programs: Partnerships for Advanced Computational Infrastructure at the San Diego SuperComputer Center, Distributed Terascale Facility (DTF), and Terascale extensions: Enhancements to the Extensible Terascale Facility. We thank D. Adams, J. Economides, A. Marathe, and C. Jocsion for help with experiments, J. D. Victor for discussions, W. Bialek for comments on the manuscript, Y. S. Liu and D. Kachan for help with information calculations, and O. Barabash for help with a program classifying events in neural firing rate.

Correspondence should be addressed to Dr. Lawrence C. Sincich, Beckman Vision Center, University of California, San Francisco, 10 Koret Way, San Francisco, CA 94143-0730. E-mail: sincich@vision.ucsf.edu.

DOI:10.1523/JNEUROSCI.3701-08.2009

Copyright © 2009 Society for Neuroscience 0270-6474/09/296207-10\$15.00/0

electrodes were amplified 1000×, analog filtered between 0.3 and 3 kHz, and acquired digitally at 25 kHz (Power 1401, Cambridge Electronic Design).

Isolated EPSPs (those that failed to cause an LGN spike) were identified by routine waveform sorting. EPSPs which generated LGN spikes required additional analysis to identify because their waveforms were partially merged with the spike waveforms. A subtraction technique was used to extract these “successful” EPSPs (Sincich et al., 2007). Each cell’s EPSPs exhibited an absolute refractory period, which is strong evidence that the EPSPs were driven by one retinal ganglion cell. EPSPs from subtracted waveforms had shape components identical to the isolated EPSPs for a given cell, indicating that the successfully transmitted EPSPs originated from the same ganglion cell as the EPSPs that failed to drive a spike. Extracellular methods do not artifactually lead to single-input recordings, because we occasionally encountered LGN neurons with EPSPs arising from more than one ganglion cell, as evidenced by superimposed EPSP waveforms [Sincich et al. (2007), their supplemental Fig. 2; see also Weyand (2007)] and absence of a refractory period. Such cells were not analyzed further because it was impossible to assign individual EPSPs to a particular ganglion cell. Stimuli larger than the receptive field center were more likely to lead to EPSPs from multiple ganglion cells. To minimize this problem, the smallest stimulus that drove the neuron robustly was used.

Visual stimulus. Stimuli were provided by an LED beam restricted to the receptive field center, as detailed previously (Sincich et al., 2007). Retinal EPSPs were measurable for 37 of the 240 cells recorded. The choice of stimulus protocol for the LED’s temporal luminance flicker was influenced by several factors. A primary consideration was the limited time available for recordings. The signal-to-noise ratio needed to record EPSPs extracellularly meant that the electrode was often very close to, if not touching, the cell. Recordings rarely lasted more than an hour, because the neuron could be injured with even minor relative motion between the electrode and the cell. This time constraint made it difficult to compare different stimuli in the same cell. We conducted pilot studies that showed that neither correlated Gaussian nor random white noise drove neurons as effectively as luminance flicker that resembled the temporal fluctuations that occur when looking at natural scenes (van Hateren, 1997; van Hateren et al., 2002). Such stimuli follow a 1/frequency power law, which is weighted toward lower temporal frequencies. In addition, the histograms of luminance values follow a non-Gaussian distribution that is often skewed and has nonzero kurtosis (van Hateren, 1997; Simoncelli and Olshausen, 2001).

We created 5 s luminance sequences that had naturalistic temporal frequency power spectra between 0.2 and 80 Hz. These stimuli were derived by first generating a series of independent samples in time from a non-Gaussian distribution that mimics distributions observed in natural scenes (computed by adding a uniformly distributed random variable over the interval [0, 1] to the square of another Gaussian random variable of unit variance and zero mean). We then Fourier-filtered the result to have a mean power spectrum that followed a $1/f_t^\gamma$ distribution, with f_t = temporal frequency, and $\gamma = 0.7$. The stimuli power spectra are available in the paper by Sincich et al. (2007) (their supplemental Fig. 1). To compute information rates for the spike trains, a series of unique stimuli was presented as well as one representative 5 s stimulus which was interleaved between the unique sequences. We note that the stimuli were presented on a black background, rather than a background at the mean stimulus luminance, because higher contrasts could be achieved and it improved responses for OFF cells. Consequently, the stimulus fluctuations were not symmetric around the mean luminance, which led to a positive shift in the probability distributions of the filter projection values (supplemental Fig. 1a–c, available at www.jneurosci.org as supplemental material).

Total information conveyed by single spikes. To compute information about the stimulus carried by single spikes, I_{spike} , we repeated one 5 s random stimulus segment between 76 and 400 times (mean = 188). Based on these repeated presentations, we computed the time-varying average firing rate $r(t)$ for this noise sequence, and the information about the stimulus as follows:

$$I_{\text{spike}} = \frac{1}{T} \int dt \frac{r(t)}{\bar{r}} \log_2 \frac{r(t)}{\bar{r}}, \quad (1)$$

where \bar{r} is the average firing rate (Brenner et al., 2000). The values of mutual information contain a positive bias, which decreases as more data are collected (Treves, 1995; Strong et al., 1998; Brenner et al., 2000). To correct for this bias, we computed information values by including different fractions of the available repetitions (80%, 85%, 90%, 95%, 100%). For a given fraction of the data (except for 100%), 1000 different combinations of the individual repetitions were used. In each case, we computed information according to Equation 1. We then used linear extrapolation to find the projected information value if infinite numbers of spikes could be collected (Strong et al., 1998; Brenner et al., 2000). The linear dependence gave a good fit for all cells included in the analysis.

This information measure makes no assumptions about the number of relevant stimulus features or about the shape of the nonlinear gain function (also known as a static nonlinearity) describing the dependence of spike probability on the relevant stimulus components. Therefore it can be used to quantify the performance of any model of a reduced dimensionality, where only a small number of stimulus features are associated with the firing rate.

Finding relevant stimulus features. The relevant stimulus features were computed with two different methods. First, as a control, we found one relevant stimulus feature using the reverse correlation method (de Boer and Kuyper, 1968; Rieke et al., 1997) and correcting for stimulus correlations (Sen et al., 2001; Theunissen et al., 2001; Ringach et al., 2002; Schwartz et al., 2006). This involves first computing a vector of the spike-triggered average (STA), by averaging all stimuli that elicited a spike, and then multiplying by the inverse of the covariance matrix. These vectors ranged 200 ms before a spike, and were binned every 4 ms. Stimuli were resampled at 250 Hz to match the bin size of the spike analysis. No correction was made to account for the EPSP-spike delay (usually <1 ms), as it was smaller than the bin size used to compute the filters. We will refer to the resultant vector as the decorrelated STA (dSTA). Because multiplication by the inverse of the stimulus covariance matrix often leads to noise amplification at high temporal frequencies, one common way to address this issue is to use a pseudoinverse of the covariance matrix (David and Gallant, 2005). A pseudoinverse is computed by diagonalizing the covariance matrix, as in the process of creating the inverse matrix, but including only some of the eigenvectors multiplied by the inverse of their respective eigenvalues. This effectively filters out poorly sampled components to yield a regularized decorrelated STA (rdSTA). For each cell, we chose the cutoff on the eigenvalues of the stimulus covariance matrix that led to the best predictive power on the test part of the data, which was not used in computing the STA.

Second, we used the method of maximally informative dimensions (MID) to find relevant stimulus features (Sharpee et al., 2004). The algorithm is available online at <http://cml-t.salk.edu>. Relevant stimulus dimensions were found by searching through the space of all possible stimulus combinations to find the dimensions that accounted for maximal information (or variance: (Sahani and Linden, 2003; Machens et al., 2004; Sharpee, 2007)) in the neural response. Information accounted for by a particular dimension \mathbf{v} can be computed as:

$$I(\mathbf{v}) = \int dx P_{\mathbf{v}}(x|\text{spike}) \log_2 \frac{P_{\mathbf{v}}(x|\text{spike})}{P_{\mathbf{v}}(x)}, \quad (2)$$

where $P_{\mathbf{v}}(x)$ is probability distribution of stimulus components along dimension \mathbf{v} and $P_{\mathbf{v}}(x|\text{spike})$ is the analogous probability distribution computed by taking only stimulus segments that lead to a spike (Adelman et al., 2003; Agüera y Arcas et al., 2003; Paninski, 2003; Sharpee et al., 2004; Fairhall et al., 2006). We briefly outline the logic behind this argument as described by Adelman et al. (2003). Let us start with the expression for the overall information carried in the arrival times of single spikes, I_{spike} . Instead of the time average, I_{spike} can be rewritten as an average over all of the presented stimuli:

$$I_{\text{spike}} = \int ds P(s) \frac{P(\text{spike}|s)}{P(\text{spike})} \log_2 \frac{P(\text{spike}|s)}{P(\text{spike})}, \quad (3)$$

where \mathbf{s} represents different vectors in the stimulus space describing presented stimuli. Using Bayes' rule this expression can be further transformed into the following:

$$I_{\text{spike}} = \int ds P(\mathbf{s}) \frac{P(\mathbf{s}|\text{spike})}{P(\mathbf{s})} \log_2 \frac{P(\mathbf{s}|\text{spike})}{P(\mathbf{s})} \\ = \int ds P(\mathbf{s}|\text{spike}) \log_2 \frac{P(\mathbf{s}|\text{spike})}{P(\mathbf{s})}. \quad (4)$$

The last variant of the equation for the overall information carried by the arrival times of single spikes is very similar to the expression for the information with respect to stimulus components along dimension \mathbf{v} . The only difference is that instead of the full, unreduced stimulus \mathbf{s} , we take stimulus components along this direction \mathbf{v} of interest.

Predictive power of reduced models: cross-validation. To estimate a model's predictive power, data from unrepeatable stimuli were separated into training and test datasets (75% for the training dataset and 25% for the test dataset). All relevant dimensions (STAs, dSTAs, rdSTAs, and MIDs) were computed using only unrepeatable parts of the dataset. These data were then further separated into a "training" and "test" subsets to minimize overfitting (David and Gallant, 2005) and to compute error bars on filter estimates (Efron and Tibshirani, 1998). In the case of rdSTA, the optimal cutoff (see above) was selected by optimizing predictive power on the test dataset.

To characterize predictive power of reduced models based on the first MID and the two MIDs considered together, we took these features as estimated from the dataset based on neural responses to unrepeatable stimuli and computed information according to Equation 2 on the dataset based on the responses to repeated stimuli. This separation of data was necessary to ensure that predictive power was not affected by overfitting effects. Even without overfitting, information $I(\mathbf{v})$ along any particular dimension \mathbf{v} is positively biased for a dataset of finite size. This bias was removed according to the same procedure as described above for the case of information conveyed by single spikes, except that the data subsets were obtained from different segments of the stimuli. Nonlinear gain functions shown in supplemental Figure 1g–i (available at www.jneurosci.org as supplemental material) were computed on the test part of the unrepeatable data.

The percentage of information explained was computed as a ratio between information carried by single spikes (as estimated from repeated trials) and information accounted for by relevant dimensions derived from the training part of the unrepeatable trials about a neuron's responses to the repeated trials.

Finding relevant retinal spike train features. We adapted the methods used to find relevant stimulus features to quantify the information carried by different retinal spike sequences. In this case, retinal spike trains were treated as stimuli binned every 4 ms, where a value of 1 was given at each occurrence of an EPSP, and was 0 otherwise. Such stimuli are binary by construction, and do not conform to Gaussian statistics. Because the MID algorithm can be used with non-Gaussian stimuli, we used it to find the relevant features of spike trains as well. The computation began by computing a spike-triggered average by averaging 200 ms segments on retinal spike trains that preceded a thalamic spike. This vector was used to initialize optimization for the first relevant feature \mathbf{v} . The probability distribution of components along dimension \mathbf{v} $P_{\mathbf{v}}(x)$ and $P_{\mathbf{v}}(x|\text{spike})$, which are necessary for computation of information according to Equation 2, were computed by projecting segments of retinal spike trains onto \mathbf{v} . This computation is well defined for binary vectors composed on spike train segments, provided the spike-triggered feature \mathbf{v} has at least one nonzero value. This condition was always fulfilled in the case of optimization of the first feature, which was initialized using the STA. The search for the second relevant feature was initialized using a random segment of the retinal spike train. In cases where this segment had no spikes, several 1 values were added at random. The remaining computation of most informative retinal spike trains followed the algorithm for finding relevant stimulus features.

Results

Spike information gain across the retinogeniculate synapse

In recordings of macaque LGN neurons acquired with flicker stimulation of the receptive field center, isolated EPSPs as well as EPSPs which evoked an LGN spike were recognizable (Fig. 1). We used a subtraction technique to identify EPSP waveforms which were partially merged with the spike waveform (Sincich et al., 2007). This allowed us to reconstruct the complete spike train of the retinal ganglion cell that provided input to the LGN cell (Fig. 1c). From these reconstructed spike trains, it was clear that the LGN neuron's spike rate was a temporally varying fraction of the ganglion cell's spike rate (Fig. 1d). The inter-EPSP intervals had an absolute refractory period, indicating that the input arose from only one neuron. The absence of surround stimulation minimized the contribution from additional ganglion cells representing the field surround. The evidence for this is that larger stimuli often evoked additional EPSPs that superimposed with the predominant EPSP. More than 95% of the LGN spikes were driven by the primary EPSP. The remaining 5% either lacked an EPSP, or were preceded by an EPSP at an interval longer than the ganglion cell's refractory period, and thus could not be definitively considered as driving the LGN spike. Therefore the spike train transformation can be accounted for almost completely in these connected cell pair recordings.

To characterize the efficiency of stimulus encoding across this synapse, we asked whether the rate of information transmission changed between the retinal and LGN spike trains. Because LGN firing rates were always lower (Fig. 2a), one might expect that information rates would also be lower based on encoding with independent spikes, a consequence of the data processing inequality (Cover and Thomas, 1991). To measure the information rates directly, we repeated a single 5 s flicker sequence multiple times and computed the average information carried by single spikes (Strong et al., 1998; Brenner et al., 2000; Reinagel and Reid, 2000). For the ON-center cell pair shown in Figure 1, we found that information increased from 0.815 ± 0.002 bits/spike for the ganglion cell to 1.155 ± 0.003 bits/spike for the LGN neuron. Similar increases were present for all the recorded pairs (Fig. 2b), with LGN neurons averaging 36% more bits per spike than ganglion cells (supplemental Table 1, available at www.jneurosci.org as supplemental material). As a control exercise, simply dropping the appropriate fraction of spikes at random from the RGC spike train to match the LGN spike rate for each cell pair produced no increase in bits/spike, as expected theoretically (Fig. 2c). Finally, we note that the computation of information carried by the arrival times of single spikes does not take into account information associated with the absence of a spike at a particular time (Fairhall et al., 2006). However, information in the "silence" between spikes was minor (mean 18% of the information carried by single spikes), and did not affect any of the results described below (supplemental Fig. 2, available at www.jneurosci.org as supplemental material). Therefore, on a spike-by-spike basis, LGN neurons can reliably signal the differences between many more stimuli than retinal ganglion cells.

Multiplying the information per spike by the firing rate yielded the average rate of information transmission. For the example cell pair in Figure 1 the rate remained remarkably constant: 33.4 ± 3.5 bits/s for the ganglion cell and 34.6 ± 5.0 bits/s for the LGN cell. This perfect information transfer rate was maintained by one-third of the cell pairs we recorded (Fig. 2d), indicating that the neural coding across the retinogeniculate synapse often attains fundamental statistical limits, because information

about the incoming stimuli cannot increase after additional operations on the spike train (Cover and Thomas, 1991). One way for LGN neurons to achieve loss-less information transmission would be to respond with perfect fidelity to retinal inputs, but this is not a possible mechanism given the substantially lower spike rates of LGN neurons. In the next sections, we consider how this increased information transmission rate might be accomplished.

Recorded signals increase information efficiency

To elucidate how LGN neurons encode stimuli more efficiently, we first examined whether LGN spikes encode different features of the stimulus than the retinal spikes. As an initial control analysis, we computed the spike-triggered average filters and their decorrelated versions, all of which revealed no major differences between the LGN neuron featured in Figure 1 and its connected ganglion cell (Fig. 3). As an alternative analysis, we computed the most relevant stimulus features by searching for those that accounted for the maximal amount of mutual information between stimulus variations and the LGN or retinal spikes, respectively (Sharpee et al., 2004, 2006). The result was a set of two maximally informative dimensions (MIDs) which represent the temporal filters for each neuron. For the cell pair illustrated in Figure 1, we found that spike probability was influenced by at least two filters, both differing significantly between the retinal and LGN neuron during the time interval from 25 to 100 ms before a spike, as made evident by subtracting the filters from one another (Fig. 4*a,b*). Root-mean square amplitudes of subtracted filters were significant for the population (MID1: 0.038 ± 0.005 SD; Student's *t* test, $p < 0.001$. MID2: 0.058 ± 0.013 SD, $p < 0.001$) compared with spike-triggered averages (0.018 ± 0.01 SD). The MID2 filters for the featured cell pair had ON peaks shifted by 20 ms compared with the MID1 filters. These results suggest that retinal ganglion cells and LGN neurons encode different features in the visual stimulus.

Beyond the difference in the relevant stimulus features, the firing rate of retinal and LGN neurons could be modulated to a different degree by stimulus components along these features. To characterize this modulation, we compared the spike rate gain functions, defined as the average spike rate resulting from stimuli with the same components along the relevant features, normalized by the mean rate across all stimuli. For the example cell pair, the

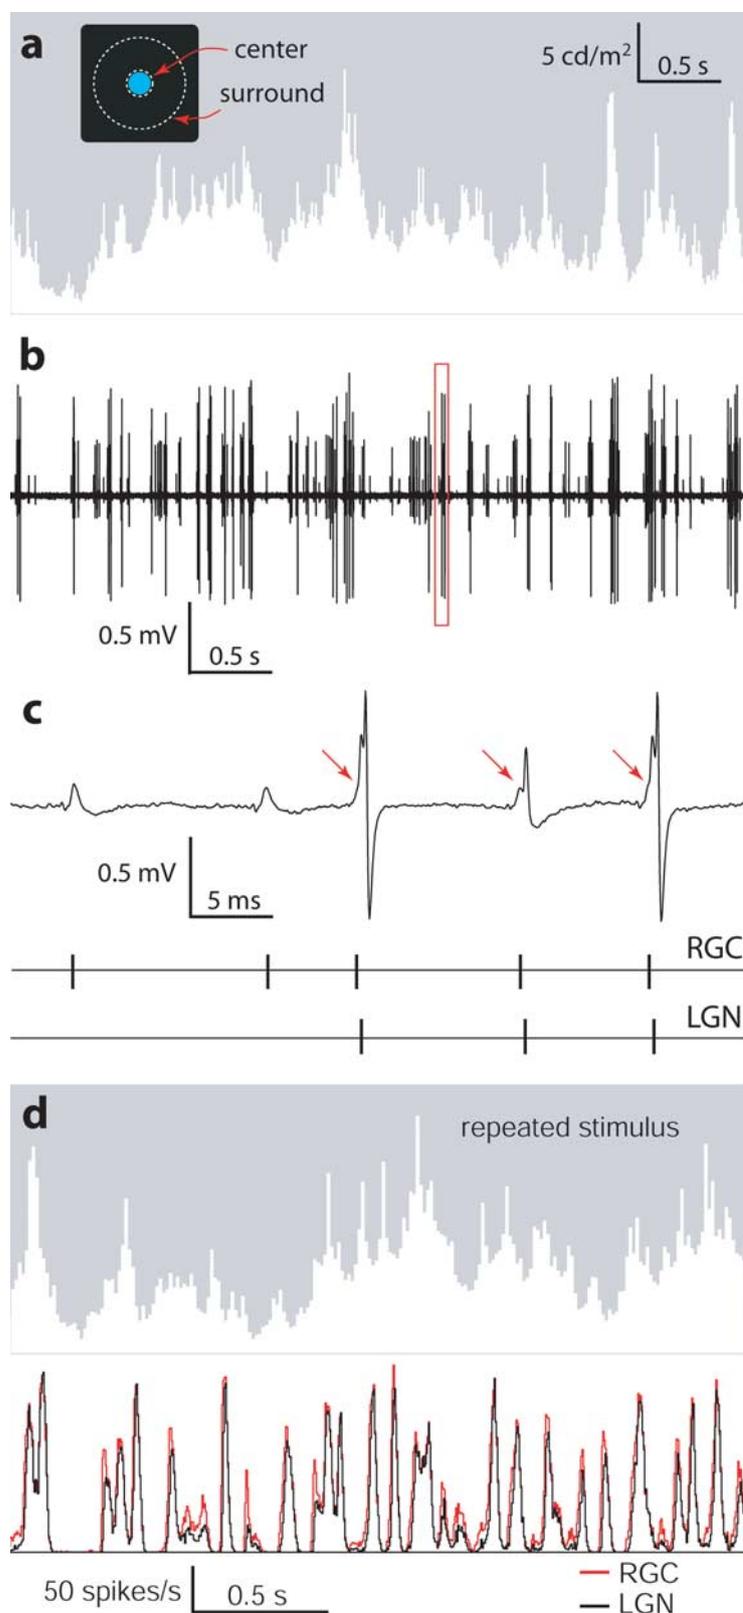


Figure 1. Extracellular recording of retinal input to a single LGN neuron. *a*, Segment of luminance flicker used to stimulate the receptive field center of an ON-center magnocellular LGN neuron. Center diameter = 0.5° , mean center luminance = 6.1 cd/m^2 , surround luminance = 0.002 cd/m^2 . *b*, Raw microelectrode voltage trace recorded in response to the 5 s flicker segment shown in *a*. *c*, Expanded view of the trace outlined in red in *b*, showing 2 isolated retinal EPSPs followed by 3 LGN spikes, one of which is small because it failed to propagate back into the soma and dendrites. Each spike was preceded by an EPSP (red arrows) and was verified by off-line analysis to reconstruct the spike trains. *d*, Peristimulus time histogram of the firing rates of the ganglion cell (RGC, red) and LGN neuron (black) in response to a portion of the repeated stimulus shows that the LGN rates are a variable fraction of the RGC rates. Histogram compiled from 295 stimulus repeats.

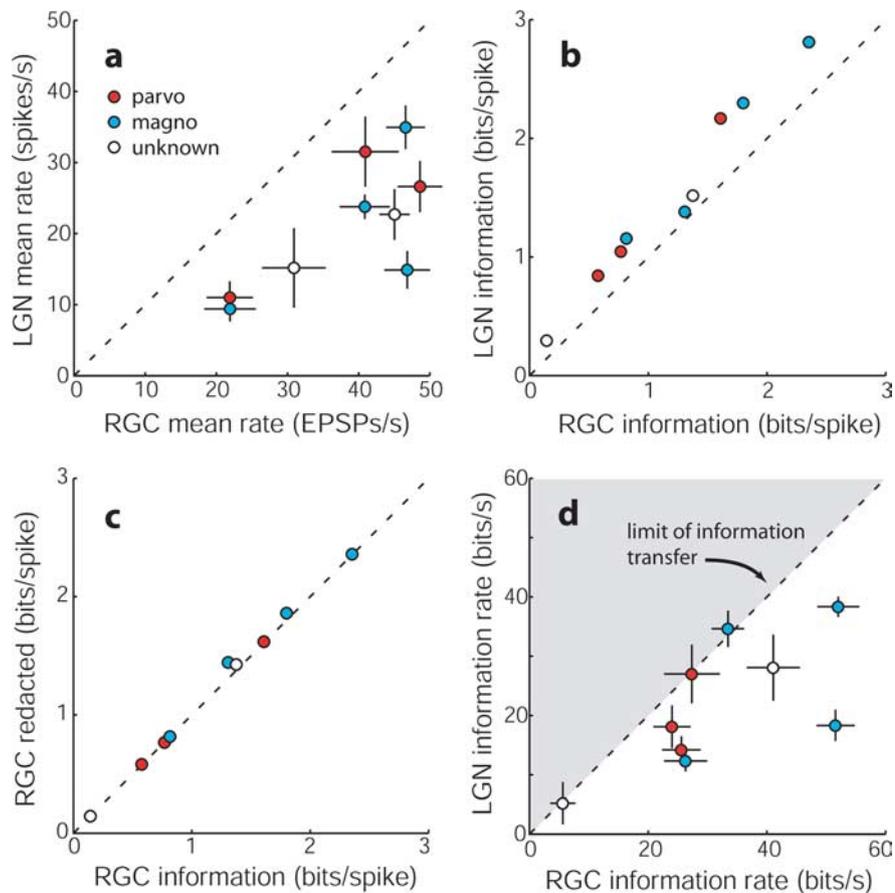


Figure 2. Increased information efficiency from retina to LGN. *a*, Mean retinal ganglion cell (RGC) spike rates are higher than mean LGN spike rates during repeat stimulus presentation. Population averages: RGC = 38.2 spikes/s; LGN = 21.1 spikes/s. Each cell pair is represented by a circle, color coded by LGN neuron type when known. Error bars are ± 1 SD, computed from repeated stimuli. *b*, Individual spikes of LGN cells always carry more information about the stimulus than corresponding input RGC spikes. *c*, Information capacity in redacted RGC spike trains (randomly dropping a fraction of the spikes to equal the number in the LGN spike train) does not differ from observed RGC trains. *d*, Multiplying firing rate (*a*) by bits/spike (*b*) reveals that LGN neurons occasionally attain transmission rates with no loss of retinal information (data points on the unity line). Population averages: RGC = 31.9 bits/s; LGN = 21.8 bits/s. Error bars are ± 1 SD, incorporating the errors when computing bits/spike (too small to plot usefully in *b* and *c*), and the variation in firing rate during repeated stimuli.

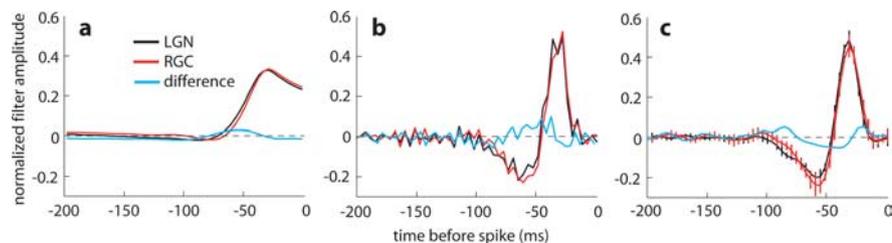


Figure 3. Significant differences did not exist between RGC and LGN filters computed as spike-triggered averages (STAs, in *a*), or as decorrelated STAs (*b*). Regularization of the decorrelated STA (*c*) smoothed noise out of the filters, but did not yield a prominent difference in filter shape. Although these decorrelated filters resembled those found previously using uncorrelated noise techniques (Chander and Chichilnisky, 2001; Reid and Shapley, 2002), they were noisier than the standard STA. Error bars are SEMs, computed from data subsets; in *a* and *b*, they were too small to be usefully plotted. RMS amplitudes of subtracted STAs averaged 0.018 ± 0.01 SD, significantly lower than the RMS values obtained from MID features ($p < 0.001$ for both filters, see Fig. 4*a,b*).

combined MID filters of the LGN neuron yielded a spike rate gain that was always higher than its retinal ganglion cell, reaching 51% greater at its peak (Fig. 4*c,d*). When the MID1 or MID2 filters were considered separately, there was either a modest or negligible difference in spike rate gain functions (supplemental Figs. 1, 3, available at www.jneurosci.org as supplemental material).

Across the population of 9 cells analyzed, the peak spike rate gains of the combined filters were consistently higher for the LGN neuron than the ganglion cell (Fig. 4*e*). The higher spike rate gains can also be shown directly, although less comprehensively, by examining peristimulus time histograms of gains for repeated stimuli. Regardless of cell type, LGN neurons consistently achieved higher gains when stimuli drove cells above the mean firing rate, and lower gains when stimuli suppressed firing (Fig. 4*f*). These results suggest that, by being able to change its instantaneous firing rate to a greater degree, an LGN neuron can signal finer distinctions between stimuli than its connected ganglion cell, and is consistent with the LGN neuron's higher information capacity.

To fully and quantitatively assess a neuron's coding ability, the frequency with which different neural responses are elicited by the stimulus ensemble must be taken into account, as well as their variability. Either Shannon information or Fisher information can be used for this purpose. Of the two quantities, only Shannon information can be used to quantify the accuracy of the transformation between stimuli and spikes either with or without a particular model of the transformation. Results from model-free analysis based on the Shannon information were shown in Figure 2. Here we compute the same information with a reduced model built from the combined MID1 and MID2 filters. In Figure 5 we show that the increased information carried by LGN spikes (like Fig. 2*b*) and the preserved overall rates of information transmission (like Fig. 2*d*) could both be captured by the reduced model. In particular, for the example cell pair, stimulus encoding based on MID filters of the LGN cell provided 1.06 ± 0.02 bits/spike, compared with the corresponding value for the ganglion cell of only 0.74 ± 0.02 bits/spike. The increased information per spike was found for all the recorded pairs (Fig. 5*a*), with the LGN filters averaging 40% more bits per spike than ganglion cells. (A model-based analysis computing Fisher information yielded similar results, see supplemental Fig. 4, available at www.jneurosci.org as supplemental material). In terms of information rate for this cell pair, the reduced model produced a rate of 36.9 ± 0.6 bits/s for the LGN neuron, while its ganglion cell's rate was not significantly different at 34.5 ± 0.8 bits/s. The same third of the population showed a preserved information rate (Fig. 5*c*) as when measured from spikes alone (Fig. 2*d*).

The ratio between the information carried by MID filters and the total information available in the spike trains gives a measure

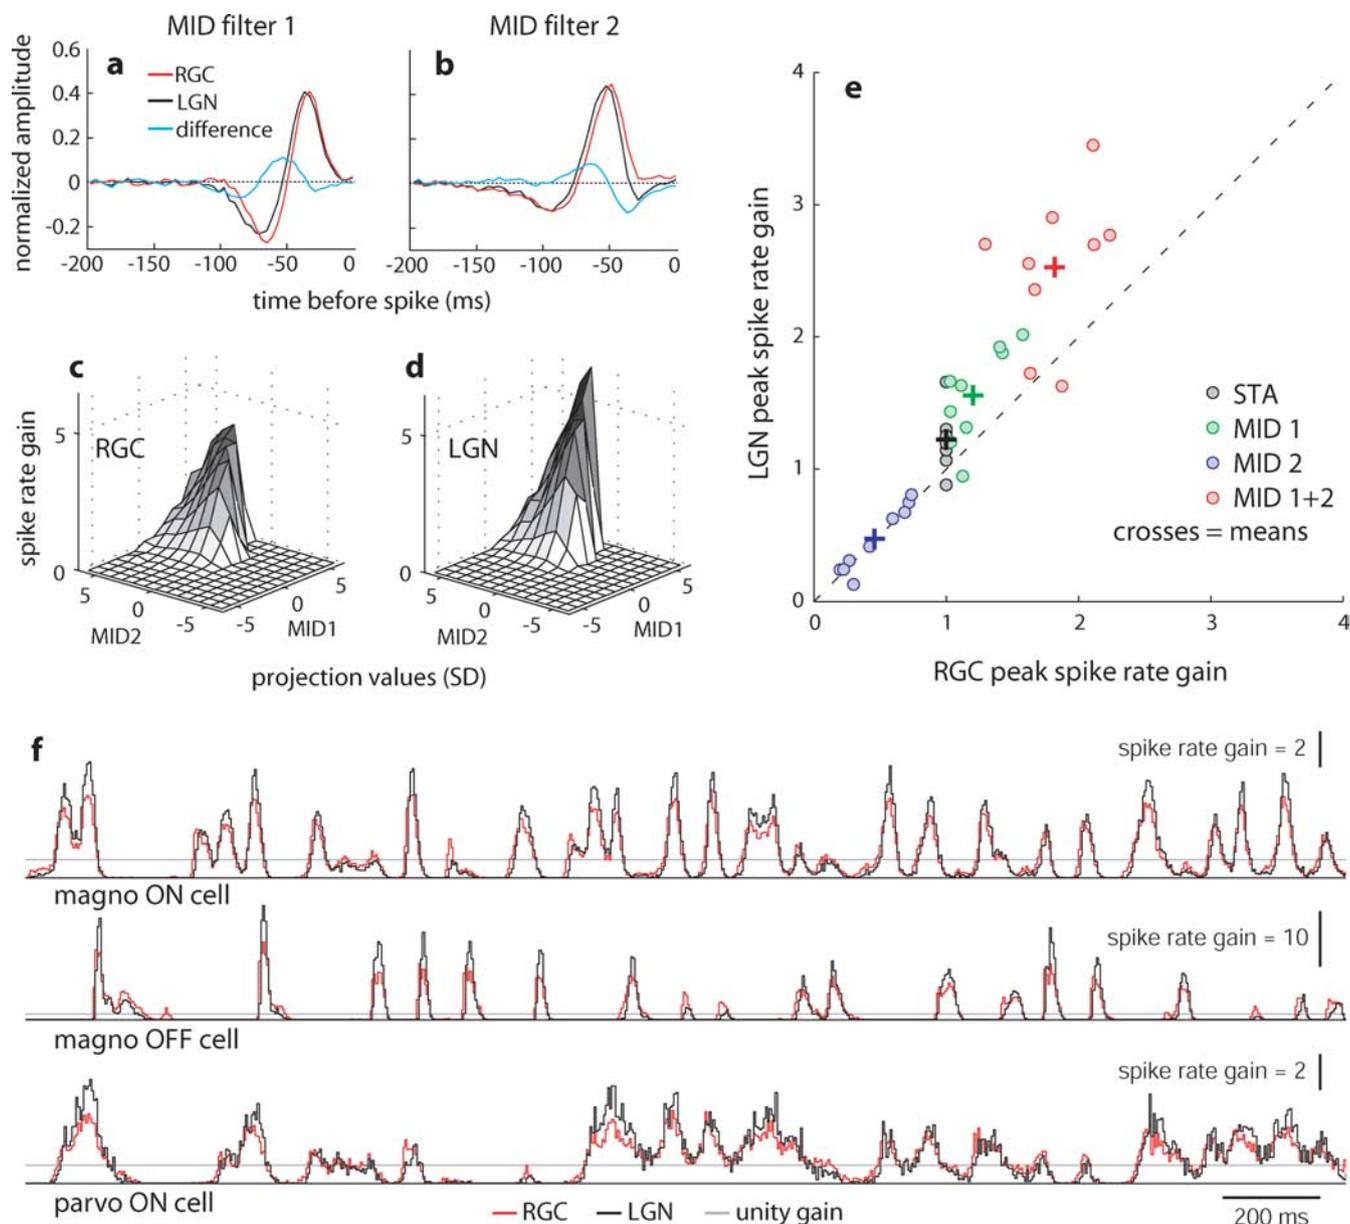


Figure 4. Stimulus representation changes from retina to LGN. *a, b*, Normalized MID filters for the RGC (red) and LGN neuron (black) are temporally altered, as revealed by subtracting the filters (blue). MID1 yields the most information, while MID2 represents an orthogonal stimulus dimension adding maximal information to the first. Error estimates computed over data subsets had average SEMs < 0.01 for each point along all filters (data not shown). *c, d*, Spike rate gain measures how far above the mean firing rate (at gain = 1) any stimulus can drive the cell, plotted as a function of projection value distributions (see Materials and Methods for details). Filter combinations significantly increase the spike rate gain of the LGN neuron (*d*) over the RGC (*c*). Projection values with positive SDs represent stimuli with increasing resemblance to the filters. *e*, Population data comparing RGC and LGN peak spike rate gains for all filters (e.g., the peaks in supplemental Fig. 1*g–i*, available at www.jneurosci.org as supplemental material), normalized to the peak RGC STA rate gain. The combined MID filters exhibited the highest gains. *f*, Peristimulus time histograms of spike rate gains for 3 cells, in response to the same repeated stimuli. Stimulus segment is the same as shown in Figure 1*d*. When LGN gains (black) are above the unity gain line (gray), they exceed RGC gains (red); and when below unity, they are lower than RGC gains, suggesting a greater modulation range for LGN neurons. The information transmission ratios for these cell pairs were: magno ON = 1.0; magno OFF = 0.74; parvo ON = 0.75.

of how well the MID filters capture retinogeniculate information transmission. For the example cell pair, the combined MID model performed better than the MID filters individually, accounting for 91% and 92% of the information carried by single spikes of the retinal and LGN cells respectively (supplemental Table 1, available at www.jneurosci.org as supplemental material). There was no significant difference in the percentage of information accounted for by the combined MID filters across the population ($p = 0.73$, Wilcoxon paired test). The MID filters explained a mean of $\sim 85\%$ of the total information available in the spike trains (Fig. 5*b*), and often nearing 100%. Therefore, as a

reduced model of the retinal and LGN neural activity, the combined MID filters captured most of the information that could be extracted about the stimulus in the spike trains. Because information here was computed indirectly by a stimulus reconstruction procedure, the quantities represent a lower bound on the amount of total information carried in single spikes (Strong et al., 1998; Reinagel et al., 1999; Reinagel and Reid, 2000; Adelman et al., 2003; Fairhall et al., 2006). The remaining $\sim 15\%$ of the information in the spike train was presumably provided by higher-order filters (which we could not compute because of data limitations), or by inputs from other neurons that were undetectable

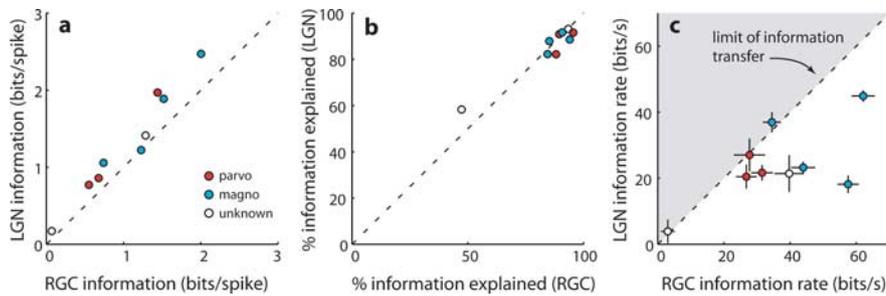


Figure 5. Maximally informative filters account for the increased efficacy of information transfer. *a*, The combined MID filters considered as a model of spike activity manifest the higher information capacity of single LGN spikes seen experimentally across the population (cf. Fig. 2*b*). Cell pairs are represented by circles, color coded by LGN neuron type. *b*, The combined model accounts for most information available in the spike trains (means: RGC = $79 \pm 18\%$; LGN = $83 \pm 15\%$). *c*, The combined filter model also exhibits the same degree of lossless information transfer rate found experimentally (cf. Fig. 2*d*). Error bars are ± 1 SD, incorporating the errors when computing bits/spike (too small to plot usefully in *a* and *b*), and the variation in firing rate during repeated stimuli.

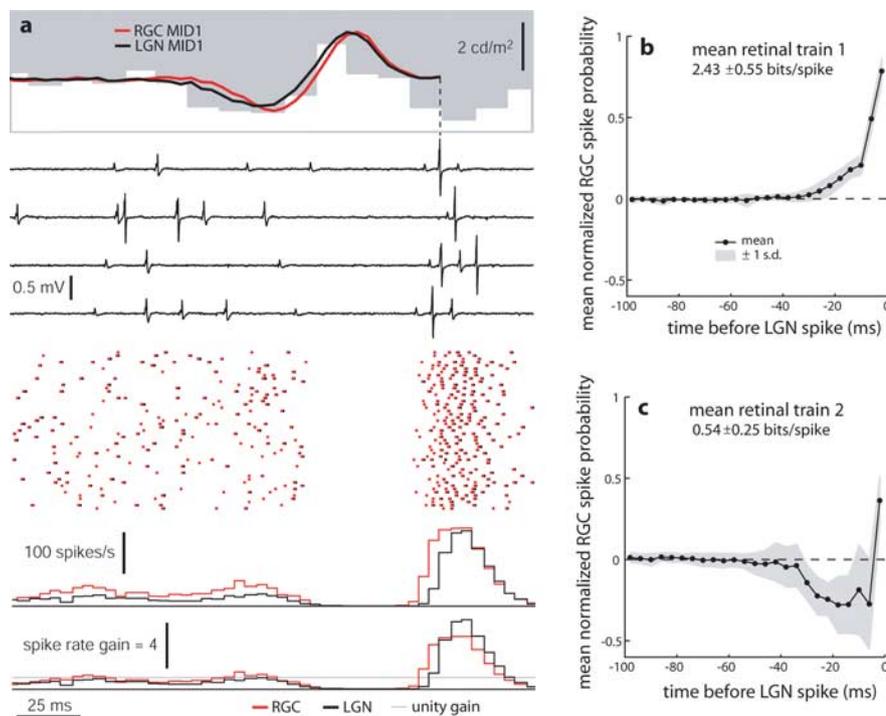


Figure 6. Short-term retinal spike summation improves LGN information transfer efficiency. *a*, Luminance steps, indicated in white, of a repeated stimulus that resembles the MID1 filters (top, aligned to a selected spike) lead to EPSPs that summate to drive LGN spikes, evident in raw traces (middle) as well as raster plots. Spike rates and corresponding spike rate gains (bottom) show that the LGN neuron (black) exhibits greater firing rate modulation than the ganglion cell (red). *b*, *c*, Different retinal spike trains convey different amounts of information. The population mean retinal train, showing increased EPSP probability during the 30 ms before an LGN spike, carries the most information (*b*), while the next most informative train shows inhibition over a longer period before an LGN spike (*c*). Shaded areas are ± 1 SD.

by extracellular recording. These results lead us to conclude that the filter transformations and the shift in gain functions realized by LGN neurons explain the lossless information transfer from the eye to cortex achieved by some cell pairs.

Temporal summation underlies spike information gain

What might be the biophysical mechanism for the precise selection of retinal spikes relayed to cortex? In these recordings, single isolated EPSPs were insufficient to bring an LGN neuron to threshold (Sincich et al., 2007). For an LGN spike to occur, at least two EPSPs had to arrive within ~ 30 ms of each other. For our example cell pair, representative recordings of EPSPs and

spikes driven by a repeated stimulus that closely matched the MID1 filters showed that solitary EPSPs reliably preceded each EPSP that generated an LGN spike (Fig. 6*a*). Linear models of EPSP summation can predict the responses of LGN neurons to a given retinal spike train arising from a spot stimulus (Carandini et al., 2007; Casti et al., 2008). These observations suggest that specific retinal spike sequences are likely to transmit information more effectively than others.

Quantitatively, one can determine what linear combinations of retinal spikes were most effective in eliciting spikes from the LGN neuron, using a strategy similar to the one followed for stimulus filters. Instead of quantifying information about the stimulus, we treated 200 ms of the EPSP train before each LGN spike as the vector to be analyzed. By this means, we found that the two most informative retinal input sequences were consistent for all LGN cells. The primary retinal train reflected the EPSP summation mechanism, with an ever-increasing probability beginning 35 ms before an LGN spike (Fig. 6*b*). The secondary retinal train, averaging 22% as much information as the first, showed a distinct inhibitory effect such that the probability of an LGN spike increased when retinal spikes were less likely during the prior 5–50 ms (Fig. 6*c*). This decrease in firing probability may represent the effect of several inhibitory sources, possibly working in concert: interneuron GABA currents (Blitz and Regehr, 2005), deactivation of T-type calcium current (Lu et al., 1992; Alexander et al., 2006), or spike after-hyperpolarizing potentials (Crunelli et al., 1987; Williams et al., 1996; Powers et al., 2005). When the spike train features were considered in combination, they increased information flow for the population by only 5% compared with the primary spike train feature. This small increase for combined trains suggests that the spike train features provide redundant contributions to coding (had they operated independently, their summed information would have been 4 times larger than the information in the combined spike train features).

Therefore information transmission occurs predominately through EPSP summation, and is only modestly improved with the addition of an inhibitory effect, whether originating from interneurons or intrinsic currents.

The combination of these two summation mechanisms was sufficient to re-encode retinal spikes with minimal loss of information. As an independent check on whether the spike train filters in Figure 6, *b* and *c*, adequately represent the firing behavior of LGN neurons, one can convolve the retinal spike trains with these filters and recompute the information carried in the resulting spike train. In this analysis, the relevant spike train features

and the nonlinearities were estimated from unrepeatable data, while the spike train transformation was applied to retinal spike trains in response to repeated stimulus segments (see Materials and Methods). We found that spikes in the transformed retinal spike trains were encoding stimuli just as efficiently as the actual LGN spikes (correlation coefficient $r = 0.99$, data in supplemental Table 1, available at www.jneurosci.org as supplemental material). This result confirms that the transformation of retinal spike trains occurs primarily through an EPSP summation mechanism, and is sufficient to achieve occasionally lossless information transmission.

Discussion

From early paired recordings of connected neurons, it has often been noted that the typical spike sequence changes from one neuron to the next, suggesting that a change in feature representation also occurs (Hubel and Wiesel, 1961; Bishop et al., 1962; Hamamoto et al., 1994; Benardete and Kaplan, 1997; Rukxenas et al., 2000). There has been considerable progress in elucidating how stimuli are represented in spike trains, especially in the visual system (Bialek et al., 1991; Meister and Berry, 1999). By comparison, there has been far less effort aimed at understanding how or why stimulus representations change from one neuron to the next. One reason is the technical difficulty of recording from two cells simultaneously for the length of time it takes to characterize firing behavior. Another reason is that synaptic convergence is quite high in most neural systems. Instances where a neuron receives input from only one or a few presynaptic cells are uncommon. Such an arrangement, though not necessary, makes the analysis of spike coding more tractable. The calyx of Held in the brainstem is an ideal example (von Gersdorff and Borst, 2002), but access to the presynaptic and postsynaptic cells is difficult, except *in vitro*. In neocortex, where cells are multiply innervated, some insights about coding have been obtained from connected cell pairs (Thomson et al., 2002), but the generally hyperpolarized state of neurons *in vitro* precludes expression of the normal activity pattern expected *in vivo*. Since neural coding is likely to be typical only *in vivo*, it is worthwhile to find an intact system where the inputs and outputs of a neuron can be readily recorded and realistic stimuli can be used. The macaque LGN comes close to fulfilling these criteria.

We have shown that the relevant stimulus features are more efficiently coded by LGN neurons than by retinal ganglion cells, based on the finding that each LGN spike conveys more information than is carried by each retinal spike. This boost in information per spike helps to compensate for a potentially large loss of information, as it has been recognized that many retinal EPSPs do not lead to LGN spikes (Hubel and Wiesel, 1961). In our population of neurons, mean EPSP efficacy was 53% during the unique stimulus sequences. In the presence of synaptic noise, one might have expected the information transfer efficacy to end up at or below 53%, as it did when retinal spikes were dropped through random failures in synaptic transmission (Fig. 2c). Instead, information transfer efficacy averaged 72% (Fig. 2d). Notably, EPSP efficacy never reached 100% in any cell, yet a third of the population achieved lossless information transfer rates. In these cases, the increase in information per spike was the maximum allowed by the laws of information transmission (Cover and Thomas, 1991). This result indicates that LGN neurons alter the incoming representations of the visual stimulus, and that a reliable mechanism determines which EPSPs elicit LGN spikes (Carandini et al., 2007; Casti et al., 2008).

The search for relevant stimulus dimensions revealed two dis-

tinct temporal filters for retinal ganglion cells as well as LGN neurons. Multiple filters have been reported in salamander retinal ganglion cells (Fairhall et al., 2006). However, the filters in the macaque LGN were not simply inherited from the retina. By recording the EPSPs as well as the LGN spikes, we could demonstrate that the retinal and LGN filters were genuinely different. The higher information content of LGN spikes suggests that the altered filters represent the statistics of the naturalistic stimuli better than ganglion cells. Although presynaptic and postsynaptic cells have not been recorded simultaneously in V1, it is clear that the receptive fields of both simple and complex cells are also best described by multiple filters (Touryan et al., 2002; Rust et al., 2005; Chen et al., 2007). Thus neural responses at every major stage of the early visual system—retina, V1, and now LGN—appear to be influenced by multiple stimulus combinations, with signal transformations at every step of the pathway. Spike-triggered averaging methods have proved insufficient for revealing the full complexity of neural responses. When used with non-Gaussian stimuli, these methods are inherently biased, leading to features that are shifted away from relevant dimensions by a finite vector that does not disappear even in the limit of infinite data (Paninski, 2003; Sharpee et al., 2004). This finite bias is large enough to obscure differences between the relevant stimulus features of retinal and LGN neurons (Fig. 3).

The increased information per spike means a greater modulation of the average firing rate in response to different stimuli relative to the average rate for all stimuli (Eq. 1). There are multiple ways to increase firing rate modulation. Previous *in vitro* studies showed that ganglion cell activity consists of separate “events” of high firing separated by epochs with no spikes (Berry et al., 1997; Fairhall et al., 2006). In this case, information per spike can be increased by reducing the duration of the firing events, or by dropping some events altogether. Given fewer events, when they do occur they signal stimuli from a smaller subset, reducing the uncertainty of which stimulus was presented and thus leading to more information. Alternatively, shortening the average duration of firing events allows one to distinguish a greater number of stimuli. Under our stimulus conditions, the number of firing events was similar for connected neural pairs ($p = 0.1$, Wilcoxon paired test), and the firing events (as defined by Berry et al., 1997) had similar duration between minima ($p = 0.3$, Wilcoxon paired test). Therefore, neither of these mechanisms explain the information capacity differences between ganglion cells and LGN neurons. Instead, the firing events *in vivo* appeared better “defined” by clear temporal sharpening of the firing peaks rather than the absence of firing (Figs. 1d), reflecting the importance of changes in spike rate gain as a means of increasing information capacity in single spikes.

To understand how the LGN neurons could achieve greater reliability in reporting different stimuli, we analyzed segments of the retinal spike trains that were most effective in eliciting LGN spikes. This analysis showed that inputs arriving within a ~30 ms period determine LGN firing. This is the same period during which nearly all EPSPs summate to reach spike threshold (Carandini et al., 2007; Sincich et al., 2007). Over this time scale, such summation is mediated mostly by NMDA current, with the effect being more pronounced at more depolarized membrane potentials (Sillito et al., 1990; Blitz and Regehr, 2003; Augustinaite and Heggelund, 2007). Very closely arriving EPSPs yield the most reliable LGN responses, and consequently if one computed retinal filters from such EPSPs they would be nearly identical to LGN filters. However, as the temporal jitter between EPSP pairs begins to increase, the first “priming” EPSP is shifted in time and would

begin to introduce noise into the retinal filters that would lower their information capacity. A second source of noise for the retinal spikes would be the EPSPs occurring at intervals longer than 30 ms and not associated with LGN spikes (several examples appear in Fig. 6*a*). These unsuccessful EPSPs occur throughout the stimulus ensemble, bringing irrelevant stimuli into the filter computation. The requirement for EPSP summation constrains the relevant stimuli to those that generate EPSP sequences which are more effective for reaching an LGN neuron's threshold, thereby refining the stimulus representations transmitted to cortex. Consequently, reliability is increased and LGN neurons can transmit more information with fewer spikes.

EPSP summation may not be the only way to raise information content per spike. In the cat inferior colliculus of the auditory system, it was found that more information was carried by single spikes when spike trains were sparser (Escabi et al., 2005). It was proposed that the cells with a lower firing rate had a higher spike threshold, and that this higher threshold could account for the increased information capacity (Goldman, 2004; Escabi et al., 2005). When the information in bits per spike was multiplied by the spike rate, this produced an information transmission rate that was always lower for spike trains with lower average firing rates. In contrast, when retinal ganglion cells or LGN neurons are considered alone, we found no relationship between spike rate and bit rate (RGC, $p > 0.5$; LGN, $p > 0.1$; Fisher Z transform test). As shown in Figure 5*c*, in one-third of the LGN cells, information rate did not decrease despite the decrease in the firing rate. Instead, the reencoding that occurs in LGN neurons allowed some of them to maintain their information transmission rates, despite a firing rate that was lower than retinal ganglion cells. This suggests that a difference in spike threshold is unlikely to be a sufficient mechanism for the preserved information transfer in the LGN.

Computations by neurons are usually cast in terms of an operation performed on multiple convergent inputs to yield an output. Our analysis of a major synaptic relay in the vertebrate visual system shows that even when neurons are driven by only one input they can perform a time-dependent operation having biological utility. The temporal receptive field is altered rather than faithfully relayed, information is well conserved, and the spike rate is reduced, implying a sparser representation at a lower metabolic cost (Laughlin, 2001; Olshausen and Field, 2004). If the information content per spike continues to increase along the visual pathways, it would help explain how higher cortical areas can represent complex stimuli with so few spikes.

References

- Adelman TL, Bialek W, Olberg RM (2003) The information content of receptive fields. *Neuron* 40:823–833.
- Agüera y Arcas B, Fairhall AL, Bialek W (2003) Computation in a single neuron: Hodgkin and Huxley revisited. *Neural Comput* 15:1715–1749.
- Alexander GM, Carden WB, Mu J, Kurukulasuriya NC, McCool BA, Nordskog BK, Friedman DP, Daunais JB, Grant KA, Godwin DW (2006) The native T-type calcium current in relay neurons of the primate thalamus. *Neuroscience* 141:453–461.
- Augustinaite S, Heggelund P (2007) Changes in firing pattern of lateral geniculate neurons caused by membrane potential dependent modulation of retinal input through NMDA receptors. *J Physiol* 582:297–315.
- Benardete EA, Kaplan E (1997) The receptive field of the primate P retinal ganglion cell. I: linear dynamics. *Vis Neurosci* 14:169–185.
- Berry MJ, Warland DK, Meister M (1997) The structure and precision of retinal spike trains. *Proc Natl Acad Sci U S A* 94:5411–5416.
- Bialek W, Rieke F, de Ruyter van Steveninck RR, Warland D (1991) Reading a neural code. *Science* 252:1854–1857.
- Bishop PO, Burke W, Davis R (1962) The interpretation of the extracellular response of single lateral geniculate cells. *J Physiol* 162:451–472.
- Blitz DM, Regehr WG (2003) Retinogeniculate synaptic properties controlling spike number and timing in relay neurons. *J Neurophysiol* 90:2438–2450.
- Blitz DM, Regehr WG (2005) Timing and specificity of feed-forward inhibition within the LGN. *Neuron* 45:917–928.
- Brenner N, Strong SP, Koberle R, Bialek W, de Ruyter van Steveninck RR (2000) Synergy in a neural code. *Neural Comput* 12:1531–1552.
- Carandini M, Horton JC, Sincich LC (2007) Thalamic filtering of retinal spike trains by postsynaptic summation. *J Vis* 7:20.1–20.11.
- Casti A, Hayot F, Xiao Y, Kaplan E (2008) A simple model of retina-LGN transmission. *J Comput Neurosci* 24:235–252.
- Chander D, Chichilnisky EJ (2001) Adaptation to temporal contrast in primate and salamander retina. *J Neurosci* 21:9904–9916.
- Chen X, Han F, Poo MM, Dan Y (2007) Excitatory and suppressive receptive field subunits in awake monkey primary visual cortex (V1). *Proc Natl Acad Sci U S A* 104:19120–19125.
- Chichilnisky EJ, Kalmar RS (2002) Functional asymmetries in ON and OFF ganglion cells of primate retina. *J Neurosci* 22:2737–2747.
- Cleland BG, Dubin MW, Levick WR (1971) Simultaneous recording of input and output of lateral geniculate neurones. *Nat New Biol* 231:191–192.
- Conley M, Fitzpatrick D (1989) Morphology of retinogeniculate axons in the macaque. *Vis Neurosci* 2:287–296.
- Cover TM, Thomas JA (1991) Elements of information theory. New York: Wiley.
- Croner LJ, Kaplan E (1995) Receptive fields of P and M ganglion cells across the primate retina. *Vision Res* 35:7–24.
- Crunelli V, Kelly JS, Leresche N, Pirchio M (1987) The ventral and dorsal lateral geniculate nucleus of the rat: intracellular recordings in vitro. *J Physiol* 384:587–601.
- Dan Y, Atick JJ, Reid RC (1996) Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *J Neurosci* 16:3351–3362.
- David SV, Gallant JL (2005) Predicting neuronal responses during natural vision. *Network* 16:239–260.
- de Boer R, Kuyper P (1968) Triggered correlation. *IEEE Trans Biomed Eng* 15:169–179.
- Dong DW, Atick JJ (1995) Statistics of natural time-varying images. *Netw Comput Neural Syst* 6:345–358.
- Efron B, Tibshirani RJ (1998) An introduction to the bootstrap. Boca Raton, FL: CRC.
- Escabi MA, Nassiri R, Miller LM, Schreiner CE, Read HL (2005) The contribution of spike threshold to acoustic feature selectivity, spike information content, and information throughput. *J Neurosci* 25:9524–9534.
- Fairhall AL, Burlingame CA, Narasimhan R, Harris RA, Puchalla JL, Berry MJ 2nd (2006) Selectivity for multiple stimulus features in retinal ganglion cells. *J Neurophysiol* 96:2724–2738.
- Goldman MS (2004) Enhancement of information transmission efficiency by synaptic failures. *Neural Comput* 16:1137–1162.
- Hamamoto J, Cheng H, Yoshida K, Smith EL 3rd, Chino YM (1994) Transfer characteristics of lateral geniculate nucleus X-neurons in the cat: effects of temporal frequency. *Exp Brain Res* 98:191–199.
- Hubel DH, Wiesel TN (1961) Integrative action in the cat's lateral geniculate body. *J Physiol* 155:385–398.
- Kaplan E, Shapley R (1984) The origin of the S (slow) potential in the mammalian lateral geniculate nucleus. *Exp Brain Res* 55:111–116.
- Laughlin SB (2001) Energy as a constraint on the coding and processing of sensory information. *Curr Opin Neurobiol* 11:475–480.
- Lu SM, Guido W, Sherman SM (1992) Effects of membrane voltage on receptive field properties of lateral geniculate neurons in the cat: contributions of the low-threshold Ca²⁺ conductance. *J Neurophysiol* 68:2185–2198.
- Machens CK, Wehr MS, Zador AM (2004) Linearity of cortical receptive fields measured with natural sounds. *J Neurosci* 24:1089–1100.
- Meister M, Berry MJ 2nd (1999) The neural code of the retina. *Neuron* 22:435–450.
- Olshausen BA, Field DJ (2004) Sparse coding of sensory inputs. *Curr Opin Neurobiol* 14:481–487.
- Paninski L (2003) Convergence properties of three spike-triggered analysis techniques. *Network* 14:437–464.
- Powers RK, Dai Y, Bell BM, Percival DB, Binder MD (2005) Contributions of the input signal and prior activation history to the discharge behaviour of rat motoneurons. *J Physiol* 562:707–724.

- Reid RC, Shapley RM (1992) Spatial structure of cone inputs to receptive fields in primate lateral geniculate nucleus. *Nature* 356:716–718.
- Reid RC, Shapley RM (2002) Space and time maps of cone photoreceptor signals in macaque lateral geniculate nucleus. *J Neurosci* 22:6158–6175.
- Reinagel P, Reid RC (2000) Temporal coding of visual information in the thalamus. *J Neurosci* 20:5392–5400.
- Reinagel P, Godwin D, Sherman SM, Koch C (1999) Encoding of visual information by LGN bursts. *J Neurophysiol* 81:2558–2569.
- Rieke F, Warland D, de Ruyter van Steveninck RR, Bialek W (1997) *Spikes: exploring the neural code*. Cambridge, MA: MIT.
- Ringach DL, Hawken MJ, Shapley R (2002) Receptive field structure of neurons in monkey primary visual cortex revealed by stimulation with natural image sequences. *J Vis* 2:12–24.
- Ruksenas O, Fjeld IT, Heggelund P (2000) Spatial summation and center-surround antagonism in the receptive field of single units in the dorsal lateral geniculate nucleus of cat: comparison with retinal input. *Vis Neurosci* 17:855–870.
- Rust NC, Schwartz O, Movshon JA, Simoncelli EP (2005) Spatiotemporal elements of macaque v1 receptive fields. *Neuron* 46:945–956.
- Sahani M, Linden JF (2003) How linear are auditory cortical responses? In: *Advances in neural information processing systems* (Becker S, Thrun S, Obermayer K, eds), pp 125–132. Cambridge, MA: MIT.
- Schwartz O, Pillow JW, Rust NC, Simoncelli EP (2006) Spike-triggered neural characterization. *J Vis* 6:484–507.
- Sen K, Theunissen FE, Doupe AJ (2001) Feature analysis of natural sounds in the songbird auditory forebrain. *J Neurophysiol* 86:1445–1458.
- Sharpee T, Rust NC, Bialek W (2004) Analyzing neural responses to natural signals: maximally informative dimensions. *Neural Comput* 16:223–250.
- Sharpee TO (2007) Comparison of information and variance maximization strategies for characterizing neural feature selectivity. *Stat Med* 26:4009–4031.
- Sharpee TO, Sugihara H, Kurgansky AV, Rebrik SP, Stryker MP, Miller KD (2006) Adaptive filtering enhances information transmission in visual cortex. *Nature* 439:936–942.
- Sillito AM, Murphy PC, Salt TE, Moody CI (1990) Dependence of retinogeniculate transmission in cat on NMDA receptors. *J Neurophysiol* 63:347–355.
- Simoncelli EP, Olshausen BA (2001) Natural image statistics and neural representation. *Annu Rev Neurosci* 24:1193–1216.
- Sincich LC, Adams DL, Economides JR, Horton JC (2007) Transmission of spike trains at the retinogeniculate synapse. *J Neurosci* 27:2683–2692.
- Smirnakis SM, Berry MJ, Warland DK, Bialek W, Meister M (1997) Adaptation of retinal processing to image contrast and spatial scale. *Nature* 386:69–73.
- Solomon SG, Peirce JW, Dhruv NT, Lennie P (2004) Profound contrast adaptation early in the visual pathway. *Neuron* 42:155–162.
- Strong SP, Koberle R, de Ruyter van Steveninck RR, Bialek W (1998) Entropy and information in neural spike trains. *Phys Rev Lett* 80:197–200.
- Theunissen FE, David SV, Singh NC, Hsu A, Vinje WE, Gallant JL (2001) Estimating spatio-temporal receptive fields of auditory and visual neurons from their responses to natural stimuli. *Network* 12:289–316.
- Thomson AM, Bannister AP, Mercer A, Morris OT (2002) Target and temporal pattern selection at neocortical synapses. *Philos Trans R Soc Lond B Biol Sci* 357:1781–1791.
- Touryan J, Lau B, Dan Y (2002) Isolation of relevant visual features from random stimuli for cortical complex cells. *J Neurosci* 22:10811–10818.
- Treves A (1995) Quantitative estimate of the information relayed by the Schaffer collaterals. *J Comput Neurosci* 2:259–272.
- Usrey WM, Reppas JB, Reid RC (1999) Specificity and strength of retinogeniculate connections. *J Neurophysiol* 82:3527–3540.
- van Hateren JH (1997) Processing of natural time series of intensities by the visual system of the blowfly. *Vision Res* 37:3407–3416.
- van Hateren JH, Rüttiger L, Sun H, Lee BB (2002) Processing of natural temporal stimuli by macaque retinal ganglion cells. *J Neurosci* 22:9945–9960.
- von Gersdorff H, Borst JG (2002) Short-term plasticity at the calyx of held. *Nat Rev Neurosci* 3:53–64.
- Weyand TG (2007) Retinogeniculate transmission in wakefulness. *J Neurophysiol* 98:769–785.
- Williams SR, Turner JP, Anderson CM, Crunelli V (1996) Electrophysiological and morphological properties of interneurons in the rat dorsal lateral geniculate nucleus in vitro. *J Physiol* 490:129–147.
- Wilson JR (1989) Synaptic organization of individual neurons in the macaque lateral geniculate nucleus. *J Neurosci* 9:2931–2953.