

The Statistical Computation Underlying Contrast Gain Control

Vincent Bonin, Valerio Mante, and Matteo Carandini

The Smith-Kettlewell Eye Research Institute, San Francisco, California 94115

In the early visual system, a contrast gain control mechanism sets the gain of responses based on the locally prevalent contrast. The measure of contrast used by this adaptation mechanism is commonly assumed to be the standard deviation of light intensities relative to the mean (root-mean-square contrast). A number of alternatives, however, are possible. For example, the measure of contrast might depend on the absolute deviations relative to the mean, or on the prevalence of the darkest or lightest intensities. We investigated the statistical computation underlying this measure of contrast in the cat's lateral geniculate nucleus, which relays signals from retina to cortex. Borrowing a method from psychophysics, we recorded responses to white noise stimuli whose distribution of intensities was precisely varied. We varied the standard deviation, skewness, and kurtosis of the distribution of intensities while keeping the mean luminance constant. We found that gain strongly depends on the standard deviation of the distribution. At constant standard deviation, moreover, gain is invariant to changes in skewness or kurtosis. These findings held for both ON and OFF cells, indicating that the measure of contrast is independent of the range of stimulus intensities signaled by the cells. These results confirm the long-held assumption that contrast gain control computes root-mean-square contrast. They also show that contrast gain control senses the full distribution of intensities and leaves unvaried the relative responses of the different cell types. The advantages to visual processing of this remarkably specific computation are not entirely known.

Key words: receptive field; adaptation; inhibition; retina; luminance; image statistics; texture

Introduction

In the early visual system, two fast-acting adaptation mechanisms control the gain of responses based on the locally prevalent image statistics, reducing gain as the stimuli become stronger, and increasing it as they become weaker. One is light adaptation, which operates in retina and adjusts gain based on the local mean intensity (Shapley and Enroth-Cugell, 1984). The other is contrast gain control, which operates in retina (Shapley and Victor, 1978, 1981; Chander and Chichilnisky, 2001; Baccus and Meister, 2002; Zghloul et al., 2005), is strengthened in thalamus and visual cortex (Kaplan et al., 1987; Sclar, 1987; Sclar et al., 1990; Cheng et al., 1995), and adjusts gain based on a measure of local stimulus contrast.

The measure of contrast that underlies contrast gain control is not entirely known. It is commonly thought to be the root-mean-square contrast, the local standard deviation of intensities relative to the mean intensity (Shapley and Victor, 1978, 1981; Bonin et al., 2005). Consistent with this hypothesis, gain is independent of the exact stimulus position within a local region (Shapley and Victor, 1978; Benardete and Kaplan, 1999) and decreases with

stimulus area (Shapley and Victor, 1981; Sclar et al., 1990; Bonin et al., 2005). However, contrast gain control could use approximate statistics, which do not require summing the squared intensity of each point in a patch of image. For example, Victor (1987) proposed that contrast gain control might sense the mean absolute deviations from the mean, not the squared deviations. Alternatively, contrast gain control might measure the prevalence of the darkest and/or of the lightest pixels in the image. This and other computations would not require processing of the whole distribution of light intensities. They would generally make contrast gain control dependent not only on the standard deviation of light intensities but also on higher moments of the distribution such as skewness (the asymmetry in the prevalence of particularly light or dark intensities) or kurtosis (the relative prevalence of particularly light or dark pixels).

In particular, there are at least two reasons to expect the measure of contrast to be sensitive to asymmetries in the distribution of light intensities. First, the early visual system employs separate pathways, ON and OFF, to encode light increments and decrements (Kuffler, 1953; Schiller, 1992), and the measure of contrast of each pathway may be biased by the range of intensities that is signaled by that pathway. Second, there is psychophysical evidence for a mechanism sensitive to higher moments of the light intensity distribution (Chubb et al., 1994, 2004). When viewing static texture stimuli, human subjects not only discriminate stimuli of different standard deviation (Fig. 1A, the root-mean-square contrast) but also stimuli of different skewness (Fig. 1B) or kurtosis (Fig. 1C). This performance can be explained by in-

Received Jan. 20, 2006; revised April 18, 2006; accepted May 8, 2006.

This work was supported by the James S. McDonnell Foundation 21st Century Research Award in Bridging Brain, Mind, and Behavior. We thank Séverine Durand for help with the experiments and Michael Landy for useful discussions.

Correspondence should be addressed to Vincent Bonin, The Smith-Kettlewell Eye Research Institute, 2318 Fillmore Street, San Francisco, CA 94115. E-mail: vincent@ski.org.

DOI:10.1523/JNEUROSCI.0284-06.2006

Copyright © 2006 Society for Neuroscience 0270-6474/06/266346-08\$15.00/0




Figure 1. White noise textures differing in standard deviation, skewness, and kurtosis. Each rectangle contains two squares, one for each value of the statistic being considered. The histograms on the top and bottom indicate the associated distribution of light intensities, i.e., the probability of having a pixel of a certain intensity as a function of intensity. **A**, Textures differing in standard deviation (root-mean-square contrast) and kurtosis. **B**, Textures differing only in skewness. **C**, Textures differing only in kurtosis. The methods to synthesize these textures were introduced by Chubb et al. (1994).

voking three perceptual mechanisms: one sensitive to mean intensity (e.g., light adaptation), one sensitive to standard deviation (e.g., contrast gain control), and one sensitive to higher moments of the light intensity distribution. A physiological correlate for the latter mechanism has not yet been identified. A simple explanation would be contrast gain control: for instance, if the measure of contrast were sensitive to skewness, the responses of the early visual system to textures of different skewness would be different, and the textures would be easy to discriminate.

Does contrast gain control measure root-mean-square contrast? Is it sensitive to higher moments of the light intensity distribution? We addressed these questions in the cat's lateral geniculate nucleus (LGN), which relays signals from retina to cortex. Borrowing the ingenious methods of Chubb et al. (1994, 2004), we characterized the computation underlying contrast gain control using white noise stimuli whose statistics were precisely controlled. We separately varied the standard deviation, the skewness, and the kurtosis and measured their individual effects on response gain.

Materials and Methods

We characterized the responses of 25 well isolated neurons recorded in lateral geniculate nucleus of three anesthetized, paralyzed cats. These neurons were held long enough (>2 h, commonly 4 h) to perform a series of more than six experiments.

Adult cats were anesthetized with ketamine (20 mg/kg) mixed with acepromazine (0.1 mg/kg) or xylazine (1 mg/kg). Anesthesia was maintained with a continuous intravenous infusion of pentothal (0.5–4 mg · kg⁻¹ · h⁻¹). Animals were paralyzed with pancuronium bromide (0.15 mg · kg⁻¹ · h⁻¹) and artificially respired with a mixture of O₂ and N₂O (typically 1:2). EEG, electrocardiogram, and end-tidal CO₂ were continuously monitored. Extracellular signals were recorded with Quartz-coated platinum/tungsten microelectrodes (Thomas Recording, Giessen, Germany), sampled at 12 kHz, and stored for off-line spike discrimination. A craniotomy was performed above the right LGN

(Horsley–Clarke, coordinates ~9 mm lateral and ~6 mm anterior). Electrodes were lowered vertically until visual responses were observed. The location of LGN was determined from the sequence of ocular dominance changes during penetration. Most cells (19 of 25) were located either in the first contralateral layer (presumably lamina A) or in the first ipsilateral layer (presumably lamina A₁).

Visual stimuli were displayed using the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) and presented monocularly on a calibrated monitor with mean luminance of 32 cd/m² and refresh rate of 124 Hz.

We classified cells into X and Y types using responses to large moving gratings of different spatial frequencies. Experiments included 14 logarithmically spaced frequencies. Cells were classified as Y type if the mean firing rate measured at high spatial frequencies was significantly higher than that expected from linear spatial summation (inferred by fitting responses to a difference of Gaussians followed by a rectification). All units presented here were of X type, which is consistent with the known laminar distribution of LGN cells (Wilson et al., 1976).

Stimuli were grids of uniform squares whose light intensities were drawn at random. Stimuli had 5–10 elements in width, covered the receptive field center and surround, were presented at a rate of 124 Hz, and lasted 20–30 s. Stimulus conditions were presented in randomized order and repeated 5–10 times. We varied the statistics of the stimulus across stimulus conditions.

We considered three statistics: standard deviation, skewness, and kurtosis. Let the random variable x_i denote the light intensity of any square in a stimulus. Standard deviation σ is the square root of the second central moment or variance $\sigma^2 = E(x_i - \bar{x})^2$, where $E()$ denotes the expected value operator, and $\bar{x}$ is mean light intensity. Skewness is the third central moment: $E(x_i - \bar{x})^3/\sigma^3$, and kurtosis is the fourth central moment $E(x_i - \bar{x})^4/\sigma^4$.

We used the method developed by Chubb et al. (1994) to synthesize the distribution from which the light intensities of the stimulus were drawn. Briefly, we devised a set of orthogonal basis functions to modulate the moments of the distribution. Let index $v \in \{0, 1, \dots, M-1\}$ denote one of the M different gray levels the stimuli can take. We computed the set of vectors $f_i(v) = v^i$, where $i \in \{0, 1, \dots, 4\}$. We then orthogonalized $f_i(v)$ using the Gram–Schmidt algorithm and normalized the result to an absolute maximum of $1/M$. We used the resulting basis functions $\lambda_i(v)$ to synthesize the distributions. Each function λ_i has distinct effects on the stimulus statistics. Translating the distribution along λ_2 changes its standard deviation and therefore by definition its kurtosis. Translations along λ_3 vary skewness but do not affect standard deviation or kurtosis. Translations along λ_4 vary kurtosis but do not change standard deviation or skewness.

Each experiment had seven conditions. In the control condition, gray levels were drawn from the uniform distribution U (with probability $1/M$). In the low and high standard deviation conditions intensities were drawn from $U - \kappa\lambda_2$ and $U + \kappa\lambda_2$. In the low and high skewness conditions the stimuli followed $U - \kappa\lambda_3$ and $U + \kappa\lambda_3$. In the low and high kurtosis conditions the stimuli followed $U - \kappa\lambda_4$ and $U + \kappa\lambda_4$. We used $\kappa = 0.9$ to ensure that all gray levels have nonzero probabilities. We used $M = 5$ gray levels to maximize the differences between the low and high conditions.

We modeled responses with a linear filter followed by a static nonlinearity (Chichilnisky, 2001). First, we estimated a separate filter for each stimulus by computing the average stimulus preceding a spike. We then estimated a separate nonlinearity for each stimulus by convolving the filter with the stimulus and scattering the resulting linear response against the measured response. We fixed the nonlinearity to the average measured response for a given linear response. At this stage, we had one linear filter and one nonlinear function for each stimulus. In a final step, we rescaled the linear filters obtained for the test condition so as to minimize the mean square error between the nonlinear functions obtained in the test conditions and that in the control condition.

We measured the amplitude of the linear filter for each stimulus. We considered the component of the filter that corresponds to the receptive field center and measured its amplitude. We defined amplitude as the standard deviation calculated across time. Other measures, such as peak-to-peak amplitude, yielded similar results.




Figure 2. Measuring LGN responses to the white noise textures. This figure illustrates responses to the control texture, which has uniform distribution. **A**, Some frames of the texture stimulus. **B**, Distribution of light intensities. **C**, Responses of an example LGN neuron (cell 51.3.1) to the texture stimulus. The raster plot indicates spike response in individual trials. The histogram shows firing rate response averaged across trials.

To estimate the true gain of the neurons, we performed simulations. We simulated the response of a fixed linear filter followed by a static nonlinearity and a Poisson generator. We set for each cell the linear filter and nonlinearity to the average calculated across stimuli. We ran our analysis on the simulated responses. Because the filter in the model is fixed, changes in amplitudes can be ascribed to biases in the estimation procedure. To estimate true gain, we normalized the changes in amplitudes observed in the neurons by data observed in the simulations.

Results

We estimated the gain of 25 LGN neurons in three anesthetized, paralyzed cats using white noise stimuli (Fig. 2). The stimuli consisted of grids of uniform squares whose light intensities were drawn at random (Fig. 2A) from a chosen distribution (Fig. 2B). The stimuli covered both the receptive field center and surround and elicited strong and reliable responses (Fig. 2C).

We summarized the observed responses with a classical model: a linear filter followed by a static nonlinearity (Fig. 3). The linear filter specifies the weights that the neuron applies when summing light intensities (Fig. 3A). The static nonlinearity converts the output of the linear filter into positive firing rates (Fig. 3B). We estimated the linear filter by computing the average stimulus that preceded a spike (Hunter and Korenberg, 1986; Sakai, 1992). We estimated the static nonlinearity by convolving the stimulus with the estimated filter and comparing the output of the filter with the observed firing rate (Chichilnisky, 2001) (Fig. 3B). As expected from previous studies (Dan et al., 1996; Chander and Chichilnisky, 2001; Baccus and Meister, 2002; Mante et al., 2005; Zghloul et al., 2005), at a given mean luminance and contrast, this simple model, a linear filter feeding into a static nonlinearity, captures the main features of the firing rate responses (Fig. 3D).

To compare linear filters across conditions, we considered the amplitude of the filter in the center of the receptive field. The filter consists of a collection of temporal weighting functions, one




Figure 3. A simple linear filter followed by a static nonlinearity describes the operation and the responses of the neuron. For the same neuron as in Figure 2, we here illustrate responses to the control texture, which has uniform distribution. **A**, Estimated linear filter. The dashed contour indicates the location at which the filter has maximal amplitude. **B**, Estimated static nonlinearity; measured mean firing rate for a given filter output. Error bars (barely visible) indicate ± 1 SE. **C**, Time course of filter component with maximal amplitude. Error bars (barely visible) indicate ± 1 standard deviation calculated across trials. **D**, Shaded areas indicate measured firing response. The solid lines indicate model predictions.

for each square in the stimulus. We selected the square placed on the receptive field center, which elicited the strongest response (Fig. 3A, dashed contours) and considered the amplitude of the associated weighting function (Fig. 3C). We asked how this amplitude is affected by the statistics of the stimulus.

We studied the effect of three statistics: standard deviation, skewness, and kurtosis. Each experiment included a control stimulus, in which intensities were drawn from a uniform distribution (Fig. 2B), and a set of test conditions, in which we precisely varied the standard deviation, skewness, and kurtosis of the distribution (Chubb et al., 1994) (Fig. 4A). We measured the effects of these variations by estimating a temporal weighting function for each condition (Fig. 4B).

We first describe the effects of these manipulations in stimulus statistics on the amplitude of the temporal weighting function and then, after correcting for biases introduced by our estimation method, interpret these effects in terms of neuronal gain.

Effects of standard deviation

Neural gain is well known to depend on stimulus standard deviation (Shapley and Victor, 1978, 1981; Victor, 1987; Baccus and Meister, 2002). To quantify this effect, we varied the standard deviation of our stimuli. In the control condition, the standard deviation of the stimulus was 23 cd/m², amounting to approximately one-third of the dynamic range of the display (64 cd/m², with a mean luminance of 32 cd/m²). In the test conditions, we held the mean constant and reduced standard deviation to 13.7 cd/m² (Fig. 4Ai) or increased it to 28.9 cd/m² (Fig. 4Aii).

As expected, varying standard deviation strongly affected the amplitude of the estimated weighting function. To assess the effects of standard deviation across neurons, we expressed the am-




Figure 4. Distribution of light intensities in the test conditions, and results for the example neuron. **A**, Distribution of light intensities in the six test conditions: low and high standard deviation (*i, ii*), low and high skewness (*iii, iv*), and low and high kurtosis (*v, vi*). **B**, Results for the example neuron from Figures 2 and 3. The solid curves show the temporal weighting functions measured for each of the test conditions. The dashed curve indicates the time course measured in the control condition (Fig. 3C).

plitudes of the weighting functions in the test conditions as multiples of the amplitude in the control condition. In the example neuron, dividing the standard deviation by 1.6 (to 14 cd/m²) multiplied the amplitude by a factor of 1.44 ± 0.04 (mean $\pm$ standard deviation) (Fig. 4*Bi*). Conversely, multiplying standard deviation by 1.3 (to 29 cd/m²) divided amplitude by 1.43 ± 0.05 (Fig. 4*Bii*).

We obtained comparable results in the population (Fig. 5*Ai, ii*). On average, multiplying standard deviation by 1.3 divided amplitude by 1.29 ± 0.02 (median, bootstrap estimates, $n = 25$) (Fig. 5*Aii*) and dividing standard deviation by 1.6 multiplied amplitude by 1.40 ± 0.03 (Fig. 5*Ai*).

Effects of skewness

Neural gain might also depend on the balance between the amounts of whiteness and blackness in the stimulus. This balance is unaffected by changes in standard deviation, which vary the prevalence of white and black pixels in equal amounts. To test this possibility, we varied the asymmetry in the intensity distribution,




Figure 5. Changes in response amplitude for the population of neurons ($n = 25$) and for simulated neurons in which gain is constant. **A**, Amplitude of responses in the six test conditions, relative to amplitude in the control condition. Open and closed histograms show gain of ON and OFF cells. Stars denote medians across cells. Amplitude is defined as the standard deviation of the temporal weighting function as estimated by spike-triggered averaging. **B**, Shaded areas show distribution of amplitude changes predicted by a linear–nonlinear Poisson model with fixed temporal weighting function. The dashed histograms are replotted from **A**.

or skewness, and measured the associated weighting functions. We did this without changing the stimulus mean and standard deviation. The control stimulus followed a uniform distribution of light intensities and was therefore not skewed (skewness, 0). In the test conditions, we skewed the stimulus toward either darker intensities (skewness, -0.4) (Fig. 4*Aiii*) or lighter intensities (skewness, 0.4) (Fig. 4*Aiv*).

Varying stimulus skewness has weak but significant effects on the estimated weighting function. For the example neuron, the weighting functions obtained when varying skewness resemble the one obtained in the control condition, but there are significant differences (Fig. 4*Biii, iv*, solid vs dashed curves). The weighting function was reduced by the skew toward dark intensities (Fig. 4*Biii*) but was unaffected by the skew toward lighter intensities (Fig. 4*Biv*).

Such weak effects were observed in most cells and depended on the relation between the stimulus and the polarity of the receptive field of the neuron (Fig. 5*Aiii, iv*). Skewing the stimulus toward the darker intensities typically reduced the amplitude of the weighting functions in ON cells by a factor of 0.94 ± 0.04

(median $\pm$ standard deviation, bootstrap, $n = 13$) but did not affect (1.00 ± 0.02 , $n = 12$) that of OFF cells (Fig. 5Aiii). Skewing the stimulus toward the lighter intensities increased the amplitude of the weighting function of ON cells by 1.06 ± 0.03 and decreased that of OFF cells by 0.88 ± 0.02 (Fig. 5Aiv).

Effects of kurtosis

Neural gain might also depend on the prevalence of the extreme intensity values in the image. Varying standard deviation or skewness does not isolate the effect of these extreme values: standard deviation determines both the extreme intensities and the intensities that are closer to the mean, and skewness determines the imbalance between dark and light intensities. To test the effect of intensities far away from the mean, we varied stimulus kurtosis without changing stimulus mean, standard deviation, and skewness. In the control condition, the stimulus had a kurtosis of 1.7 (for comparison, a Gaussian distribution has a kurtosis of 3). In the test conditions, we reduced stimulus kurtosis to 1.5 (Fig. 4Av) or increased it to 1.9 (Fig. 4Avi).

Varying kurtosis had weak effects on the estimated weighting functions. In the example neuron, for instance, decreasing kurtosis had little effect on the weighting function (Fig. 4Bv), whereas increasing kurtosis slightly reduced its amplitude (Fig. 4Bvi).

Similar results were obtained in the population (Fig. 5Av,vi). Across cells, decreasing kurtosis had no significant effect on the weighting functions (median normalized amplitude of 0.99 ± 0.02 and 1.02 ± 0.03 for ON and OFF cells) (Fig. 5Av). Increasing kurtosis, however, did somewhat reduce the amplitude of the weighting function (median normalized amplitudes of 0.93 ± 0.02 for both ON and OFF cells) (Fig. 5Avi).

Estimating true gain

We have seen strong effects of standard deviation and weak effects of skewness and kurtosis on the amplitude of the estimated temporal weighting function. Are some of these effects attributable to biases in our estimation method or do they reflect true changes in neuronal gain?

Spike-triggered averaging yields the true temporal weighting function of a neuron only under specific assumptions (Chichilnisky, 2001; Paninski, 2003). First, the responses of the neuron have to be well described by a linear receptive field followed by a static nonlinearity. Second, the white noise stimuli have to be drawn from a Gaussian distribution. When the light intensities of the stimulus are drawn from other distributions, instead, the method can show significant biases (Victor and Knight, 1979; Simoncelli et al., 2004).

To assess the biases in our estimation procedure, we simulated the spike responses of neurons with fixed gain and ran our analysis on these responses (Fig. 5B, shaded histograms). The model consisted of a fixed linear filter followed by a static nonlinearity and a Poisson generator. For each cell, we set the linear filter and the nonlinearity to the average filter and nonlinearity calculated across stimuli for that cell. We then calculated the responses of the model cells to our stimuli and estimated a temporal weighting function for each stimulus condition. Because the filter in the model is fixed, systematic changes in the temporal weighting functions measure the biases in the estimation procedure.

This analysis of the bias suggests that the large changes in amplitude observed when varying standard deviation reflect true gain changes, whereas the small changes observed when varying skewness and kurtosis do not reflect true gain changes (Fig. 5B, dashed vs shaded histograms). Both reducing and increasing




Figure 6. Changes in neuronal gain for the population of neurons ($n = 25$), the same as in Figure 5 except that open and closed histograms show for ON and OFF cells the ratios of observed amplitude changes over those predicted by the linear nonlinear Poisson model.

standard deviation reduced the amplitudes of the estimated weighting functions of the model neurons by 0.91 ± 0.01 and 0.93 ± 0.01 (median, $n = 25$) (Fig. 5Bii, shaded histograms). These effects, however, were small compared with the changes in estimated gain seen in the real neurons (Fig. 5Bii, dashed histograms). The effects of standard deviation, therefore, are far from being fully explained by biases in the estimation procedure. In contrast, the biases seen when varying skewness and kurtosis resemble the changes in amplitude observed in real neurons, suggesting that the latter do not reflect true changes in gain (Fig. 5Biii–vi). Just as in the real neurons, changing skewness or kurtosis with the model cells had small but significant effects on the amplitude of the estimated weighting functions, often with opposite effects in simulated ON cells and OFF cells. The distributions of amplitude for the data and for the simulations are very similar (Fig. 5Biii–vi), suggesting that the effects of skewness and kurtosis seen in the real data are entirely explained by biases in the estimation procedure.

To compensate for the biases in the estimation procedure and thereby estimate the true gain of the neurons, we normalized the changes in amplitude observed in the neural responses by the changes observed in the simulations with model neurons that have fixed gain (Fig. 6).

Changes in standard deviation elicited strong changes in gain (Fig. 6i,ii). Consistent with an almost perfect mechanism of gain control, the changes in gain almost entirely counteracted the changes in standard deviation: dividing standard deviation by 1.6 (Fig. 6i) multiplied gain by 1.54 ± 0.04 , and multiplying standard deviation by 1.3 divided amplitude by 1.20 ± 0.05 (Fig. 6ii). As expected from studies in retina (Chander and Chichilnisky, 2001; Zghloul et al., 2005) and in LGN (Bonin et al., 2005), these changes in gain were more pronounced in ON cells than in OFF cells. Increasing standard deviation by a factor of 2.1 (from the low to the high standard deviation condition) decreased the gain of ON cells by a factor 2.13 ± 0.13 (median, $n = 12$) (Fig. 6, open histograms) and that of OFF cells by 1.63 ± 0.07 ($n = 13$) (Fig. 6, closed histograms).

Neuronal gain, in contrast, was unaffected by changes in skewness (Fig. 6iii,iv) or in kurtosis (Fig. 6v,vi). For both ON and OFF cells, the estimates of gain do not differ significantly from 1 (medians, two-sided sign tests, $p > 0.05$). Thus, the changes in

Table 1. Summary of stimulus statistics and results

	Standard Deviation (cd/m ²)	Skewness (unitless)	Kurtosis (unitless)	Absolute deviation (cd/m ²)	Firing rate (spikes/s)		Gain (%)		Explained variance (%)	
					ON cells	OFF cells	ON cells	OFF cells	ON cells	OFF cells
Control	22.6	0.0	1.7	19.2	mean (SE)	mean (SE)	median (SD)	median (SD)	mean (SE)	mean (SE)
Tests	13.7	0.0	2.2	10.6	19.2 (4.2)	19.9 (3.1)	100 (2.8)	100 (2.0)	76.8	78.3
	28.9	0.0	1.2	27.8	17.8 (4.5)	16.4 (2.8)	162 (6)	147 (7)	79.1 (3.2)	79.5 (2.6)
	22.8	−0.4	1.7	19.3	19.2 (4.2)	21.1 (3.2)	77 (3)	88 (3)	74.4 (2.7)	75.6 (2.2)
	22.6	0.4	1.7	19.2	18.4 (4.3)	19.6 (3.1)	101 (2)	105 (2)	76.2 (3.0)	78.0 (2.4)
	22.8	0.0	1.5	21.1	19.4 (4.1)	19.3 (3.1)	101 (3)	99 (2)	78.0 (2.8)	76.6 (2.4)
	22.6	0.0	1.9	17.3	19.0 (4.2)	19.8 (3.2)	102 (3)	102 (2)	76.9 (2.6)	77.6 (2.3)
					18.9 (4.2)	19.5 (3.1)	101 (3)	102 (3)	74.7 (3.5)	75.1 (2.4)

amplitudes observed when varying skewness and kurtosis (Fig. 5*Aiii–vi*) entirely reflect biases in the estimation procedure (Fig. 5*Biii–vi*). These results are summarized in Table 1.

Model performance

The comparisons that we have made between the linear filters, and thus the gains, measured with the different stimulus statistics are meaningful only if the static nonlinearity at the output of the model remains constant across these manipulations (Fig. 7) (Chichilnisky, 2001; Baccus and Meister, 2002; Zaghloul et al., 2005). In our analysis, we imposed that the nonlinearity be exactly the same across conditions. Doing so is appropriate because, as illustrated for the example cell, the nonlinearities obtained in the test conditions (Fig. 7*A*, symbols) are indistinguishable from the one measured in the control condition (Fig. 7*A*, gray curves).

Similarly, the comparisons that we have made across stimulus conditions would be inappropriate if the simple model used to describe the responses (the linear filter followed by the static nonlinearity) were to work better for some stimulus conditions than for others. Indeed, the model is not perfect: it captures the main features of the responses but also exhibits deviations, often failing to capture the precise size of the firing events (Fig. 3*D*).

We addressed this concern in three ways. First, when estimating the static nonlinearity we plotted actual response as a function of predicted response in the various stimulus conditions (Figs. 3*B*, 7). The size of the error bars is constant across conditions (the error bars are in most cases invisible), suggesting that fit quality does not depend on the statistics of the stimulus. Second, we measured the percentage of stimulus-driven variance in the data that is explained by the model (Sahani and Linden, 2003; Machens et al., 2004). We found that this measure barely depends on the statistics of the stimulus (Table 1). Percentage of variance depended weakly on stimulus standard deviation, ranging from 79.3 ± 2.0% (mean ± SE, *n* = 25) at low standard deviation to 77.6 ± 1.7% and 75.4 ± 1.8% as standard deviation is increased. In the remaining conditions, explained variance also shows weak dependencies, with values ranging from 74.5 ± 2.0% to 77.3 ± 1.7%. Third, we computed the distribution of deviations from the model predictions and found it to be nearly independent of the statistics of the stimulus. We estimated the distribution of deviations for the test conditions and for the control condition and used quantile analysis to compare them (Fig. 7*B*). In all graphs, the quantiles of deviations in the test conditions resemble the quantiles in the control (the data points fall near the diagonal line). This identity indicates that model deviations for the test conditions and for control have nearly identical distributions.

Discussion

We have used white noise stimuli and a simple model of responses to measure the gain of neurons in lateral geniculate nu-

cleus. Our analysis revealed that gain is set by the standard deviation of light intensities and is independent of higher moments such as skewness and kurtosis.

Previous studies used stimuli whose statistics cannot be varied independently. These stimuli include sine waves (Shapley and Victor, 1978, 1981; Victor, 1987; Benardete et al., 1992; Bonin et




Figure 7. Model performance. *A*, The nonlinearity is approximately constant across stimulus conditions. Results for the example neuron of Figures 2, 3, and 4. Data points indicate the nonlinearities measured in the test conditions. The gray curve indicates the nonlinearity measured in the control condition (Fig. 3*B*). *B*, The quality of model predictions is similar across conditions. Results for the population of neurons (*n* = 25) are shown. Data points indicate quantiles of deviations in test conditions as a function of quantiles observed in the control condition. The dashed lines indicate unity relationship.

al., 2005; Zaghoul et al., 2005) and Gaussian noise (Chander and Chichilnisky, 2001; Kim and Rieke, 2001; Baccus and Meister, 2002; Zaghoul et al., 2005). The stimuli have an intensity distribution of fixed shape, so that varying their amplitude varies the standard deviation, but does not vary higher moments. Studies using more complex stimuli (Victor, 1987) did not explicitly test the role of the various intensity statistics in setting gain.

Our results directly confirm the long-standing hypothesis that the statistic driving contrast gain control is root-mean-square contrast (Shapley and Victor, 1978, 1981; Benardete et al., 1992; Benardete and Kaplan, 1999; Baccus and Meister, 2002; Bonin et al., 2005; Zaghoul et al., 2005). In particular, they argue against the proposal that contrast gain control might compute the absolute deviations of light intensities rather than the square deviations (Victor, 1987). Indeed, this measure of contrast is changed from the control value of 19.2 cd/m² to the values of 21.1 or 17.3 cd/m² in the stimuli in which kurtosis is decreased or increased (Table 1). If contrast gain control measured absolute deviations, gain should have decreased when kurtosis was decreased and increased when kurtosis was increased. Similarly, the results argue against the possibility that contrast gain control might measure the prevalence of the darkest and/or of the lightest pixels in the image. This and other alternative computations would have predicted a dependence of gain not only on standard deviation but also on skewness or kurtosis.

Our results were similar for ON and OFF cells; we found no bias for the positive and negative deviations from the mean that are signaled by these cells. Rather than adapting to the limited range of intensities that it signals, each cell adapts to the full range of stimulus intensities. Therefore, the adaptive mechanism maintains without corruption the relative responses of the different cell types. This is surprising because each pathway only has access to a subset of the stimulus features and because there is limited cross talk between the two pathways (Schiller, 1992; but see Zaghoul et al., 2003). To compute root-mean-square contrast without bias, the mechanism of contrast gain control might need to integrate signals between the ON and OFF pathways.

The advantages to visual processing of this remarkably specific computation are not known. Contrast gain control is thought to help map the distribution of contrasts in natural scenes onto a limited range of firing rates (Laughlin, 1981; Tadmor and Tolhurst, 2000; Schwartz and Simoncelli, 2001). Consistent with this hypothesis, and in agreement with recent studies in retina (Chander and Chichilnisky, 2001; Zaghoul et al., 2005) and LGN (Bonin et al., 2005), we found the effects of contrast gain control to be more pronounced in ON than in OFF cells. This asymmetry might reflect ecological constraints because the distribution of contrasts in the visual world is asymmetric (Ruderman, 1994; Tadmor and Tolhurst, 2000; Frazor and Geisler, 2006). Still, it is not entirely clear why the visual system would require a precise computation of standard deviation.

Our results remove an important obstacle on the road to predicting LGN responses to arbitrary, complex stimuli, including stimuli that would be encountered in nature. Although models based on a fixed linear filter can predict the gist of LGN responses to some complex stimuli (Dan et al., 1996), a fixed linear filter is unlikely to perform well in predicting responses to natural images. Indeed, these images contain local variations in mean luminance and in contrast that must strongly engage the gain control mechanisms (Mante et al., 2005; Frazor and Geisler, 2006). We have recently characterized these gain control mechanisms, and we have shown that incorporating them into a model of LGN responses allows one to predict responses to a variety of stimuli,

such as gratings and sums of gratings (Bonin et al., 2005; Mante et al., 2005). These previous results clarify the spatial extent of contrast gain control (Bonin et al., 2005) and the effect of contrast gain control on response amplitude (Bonin et al., 2005) and time course (Mante et al., 2005). They also clarify the spatial filtering properties of contrast gain control, which differ strikingly from those of the receptive field (Bonin et al., 2005). The measure of contrast used to set gain is influenced by a broad range of spatial frequencies, including very low frequencies to which LGN neurons barely respond. To obtain a model that can be applied to arbitrary images, however, one needs to know how to compute the measure of contrast that drives the contrast gain control mechanism. The results obtained here indicate that this quantity is proportional to the local standard deviation of light intensities.

Our finding that the processing in LGN is invariant to stimulus skewness indicates that the psychophysical results of Chubb et al. (1994, 2004) are not attributable to contrast gain control in the early visual system. As mentioned in the Introduction, from discrimination tests such as the one illustrated in Figure 1, these authors found evidence for three perceptual mechanisms, one sensitive to the mean luminance, one sensitive to the standard deviation, and one sensitive to asymmetries in the distribution of light intensities (and in particular to the prevalence of the darkest intensities). Our results indicate that contrast gain control in the early visual system does not contribute to the third mechanism. A possibility is that a neural correlate for such a mechanism is not visible at the level of single cell responses. Another possibility is that it is only visible in higher visual areas.

A limitation of our work is that we have concentrated on contrast gain control, and therefore we have not considered the effects of varying mean luminance. The effects of light adaptation and contrast gain control, however, are functionally independent (Mante et al., 2005); thus, the two mechanisms can be safely studied separately. Our results, therefore, are likely to hold for a range of mean light levels.

Another limitation of our work is that we have not considered how the statistics of the stimulus influence the time course of the responses. Contrast gain control not only affects the gain of the responses but also their temporal dynamics (Shapley and Victor, 1978, 1981; Mante et al., 2005). We have looked for changes of integration time accompanying the changes in gain but found no significant effect. A possible explanation is that our estimates of integration time are too coarse to reveal these variations (Pillow and Simoncelli, 2003). Another possibility is that the changes in dynamics occur at lower contrasts than could be tested with our stimuli.

Moreover, we have not considered the dynamics of contrast gain control, i.e., how a change in contrast at a given time can influence gain at a later time. We cannot study these dynamics here because in our stimuli, contrast was fixed through time. In particular, we cannot distinguish fast components of contrast gain control (Victor, 1987; Baccus and Meister, 2002; Zaghoul et al., 2005) from slower components of contrast adaptation (Chander and Chichilnisky, 2001; Kim and Rieke, 2001; Baccus and Meister, 2002; Solomon et al., 2004).

Finally, our results leave open an important question: whether the gain of LGN neurons depends on the phase spectrum of the stimulus. Studies in retina (Smirnakis et al., 1997) and primary visual cortex (Mechler et al., 2002; Felsen et al., 2005) suggest the existence of nonlinear mechanisms sensitive to structure in the spatial phase spectrum, e.g., as associated with edges or more complex features. Our white noise stimuli are not ideal to test for such mechanisms because they have a flat amplitude spectrum

and a random phase spectrum and therefore do not allow investigation of the effects of second-order spatial statistics. The amplitude spectrum solely depends on second-order statistics such as the standard deviation. The phase spectrum, instead, depends on higher-order statistics such as skewness and kurtosis. By changing skewness or kurtosis while holding standard deviation constant, we have varied the phase spectrum from one random distribution to another, while holding the amplitude spectrum constant. Although response gain was not affected by these variations, it might be affected by manipulations that introduce spatial structure, i.e., a nonrandom phase spectrum. This issue remains open and thus requires further investigation.

References

- Baccus SA, Meister M (2002) Fast and slow contrast adaptation in retinal circuitry. *Neuron* 36:909–919.
- Benardete EA, Kaplan E (1999) The dynamics of primate M retinal ganglion cells. *Vis Neurosci* 16:355–368.
- Benardete EA, Kaplan E, Knight BW (1992) Contrast gain in the primate retina: P cells are not X-like, some M cells are. *Vis Neurosci* 8:483–486.
- Bonin V, Mante V, Carandini M (2005) The suppressive field of neurons in lateral geniculate nucleus. *J Neurosci* 25:10844–10856.
- Brainard DH (1997) The psychophysics toolbox. *Spat Vis* 10:433–436.
- Chander D, Chichilnisky EJ (2001) Adaptation to temporal contrast in primate and salamander retina. *J Neurosci* 21:9904–9916.
- Cheng H, Chino YM, Smith III EL, Hamamoto J, Yoshida K (1995) Transfer characteristics of lateral geniculate nucleus X neurons in the cat: effects of spatial frequency and contrast. *J Neurophysiol* 74:2548–2557.
- Chichilnisky EJ (2001) A simple white noise analysis of neuronal light responses. *Network* 12:199–213.
- Chubb C, Econopoulou J, Landy MS (1994) Histogram contrast analysis and the visual segregation of IID textures. *J Opt Soc Am A Opt Image Sci Vis* 11:2350–2374.
- Chubb C, Landy MS, Econopoulou J (2004) A visual mechanism tuned to black. *Vision Res* 44:3223–3232.
- 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.
- Felsen G, Touryan J, Han F, Dan Y (2005) Cortical sensitivity to visual features in natural scenes. *PLoS Biol* 3:e342.
- Frazor RA, Geisler WS (2006) Local luminance and contrast in natural images. *Vision Res* 46:1585–1598.
- Hunter IW, Korenberg MJ (1986) The identification of nonlinear biological systems: Wiener and Hammerstein cascade models. *Biol Cybern* 55:135–144.
- Kaplan E, Purpura K, Shapley R (1987) Contrast affects the transmission of visual information through the mammalian lateral geniculate nucleus. *J Physiol (Lond)* 391:267–288.
- Kim KJ, Rieke F (2001) Temporal contrast adaptation in the input and output signals of salamander retinal ganglion cells. *J Neurosci* 21:287–299.
- Kuffler SW (1953) Discharge patterns and functional organization of mammalian retina. *J Neurophysiol* 16:37–68.
- Laughlin S (1981) A simple coding procedure enhances a neuron's information capacity. *Z Naturforsch [C]* 36:910–912.
- Machens CK, Wehr MS, Zador AM (2004) Linearity of cortical receptive fields measured with natural sounds. *J Neurosci* 24:1089–1100.
- Mante V, Frazor RA, Bonin V, Geisler WS, Carandini M (2005) Independence of luminance and contrast in natural scenes and in the early visual system. *Nat Neurosci* 8:1690–1697.
- Mechler F, Reich DS, Victor JD (2002) Detection and discrimination of relative spatial phase by V1 neurons. *J Neurosci* 22:6129–6157.
- Paninski L (2003) Convergence properties of three spike-triggered analysis techniques. *Network* 14:437–464.
- Pelli DG (1997) The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat Vis* 10:437–442.
- Pillow JW, Simoncelli EP (2003) Biases in white noise analysis due to non-Poisson spike generation. *Neurocomputing* 52–54:109–115.
- Ruderman DL (1994) The statistics of natural images. *Netw Comput Neural Sys* 5:517–548.
- Sahani M, Linden JF (2003) How linear are auditory cortical responses? In: *Advances in neural information processing systems 15* (Becker S, Thrun S, Obermayer K, eds), pp 109–116. Cambridge, MA: MIT.
- Sakai HM (1992) White-noise analysis in neurophysiology. *Physiol Rev* 72:491–505.
- Schiller PH (1992) The ON and OFF channels of the visual system. *Trends Neurosci* 15:86–92.
- Schwartz O, Simoncelli EP (2001) Natural signal statistics and sensory gain control. *Nat Neurosci* 4:819–825.
- Sclar G (1987) Expression of “retinal” contrast gain control by neurons of the cat's lateral geniculate nucleus. *Exp Brain Res* 66:589–596.
- Sclar G, Maunsell JHR, Lennie P (1990) Coding of image contrast in central visual pathways of the macaque monkey. *Vision Res* 30:1–10.
- Shapley RM, Enroth-Cugell C (1984) Visual adaptation and retinal gain controls. *Prog Retinal Res* 3:263–346.
- Shapley RM, Victor JD (1978) The effect of contrast on the transfer properties of cat retinal ganglion cells. *J Physiol (Lond)* 285:275–298.
- Shapley RM, Victor J (1981) How the contrast gain modifies the frequency responses of cat retinal ganglion cells. *J Physiol (Lond)* 318:161–179.
- Simoncelli EP, Pillow JW, Paninski L, Schwartz O (2004) *Characterization of neural responses with stochastic stimuli*. In: *The cognitive neurosciences*, Ed 3 (Gazzaniga MS, ed) Cambridge, MA: MIT.
- 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.
- Tadmor Y, Tolhurst DJ (2000) Calculating the contrasts that retinal ganglion cells and LGN neurones encounter in natural scenes. *Vision Res* 40:3145–3157.
- Victor J (1987) The dynamics of the cat retinal X cell centre. *J Physiol (Lond)* 386:219–246.
- Victor JD, Knight BW (1979) Nonlinear analysis with an arbitrary stimulus ensemble. *Q Appl Math* 37:113–136.
- Wilson PD, Rowe MH, Stone J (1976) Properties of relay cells in the cat's lateral geniculate nucleus: a comparison of W-cells with X- and Y-cells. *J Neurophysiol* 39:1193–1209.
- Zaghloul KA, Boahen K, Demb JB (2003) Different circuits for ON and OFF retinal ganglion cells cause different contrast sensitivities. *J Neurosci* 23:2645–2654.
- Zaghloul KA, Boahen K, Demb JB (2005) Contrast adaptation in subthreshold and spiking responses of mammalian Y-type retinal ganglion cells. *J Neurosci* 25:860–868.