

Independence of luminance and contrast in natural scenes and in the early visual system

Valerio Mante¹, Robert A Frazor^{1,2}, Vincent Bonin¹, Wilson S Geisler² & Matteo Carandini¹

The early visual system is endowed with adaptive mechanisms that rapidly adjust gain and integration time based on the local luminance (mean intensity) and contrast (standard deviation of intensity relative to the mean). Here we show that these mechanisms are matched to the statistics of the environment. First, we measured the joint distribution of luminance and contrast in patches selected from natural images and found that luminance and contrast were statistically independent of each other. This independence did not hold for artificial images with matched spectral characteristics. Second, we characterized the effects of the adaptive mechanisms in lateral geniculate nucleus (LGN), the direct recipient of retinal outputs. We found that luminance gain control had the same effect at all contrasts and that contrast gain control had the same effect at all mean luminances. Thus, the adaptive mechanisms for luminance and contrast operate independently, reflecting the very independence encountered in natural images.

In the early visual system, two rapid adaptive mechanisms control the gain of neural responses¹: luminance gain control and contrast gain control. Luminance gain control (also known as light adaptation) occurs largely in the retina. It adjusts sensitivity to match the locally prevalent luminance (light intensity). Rather than responding linearly with luminance, which can potentially vary over an extremely wide range, the retina effectively divides luminance by the local mean luminance^{1–3}. Contrast gain control begins in the retina^{1,4–7} and is strengthened at subsequent stages of the visual system^{8,9}. It effectively divides¹⁰ the responses by a measure that grows with the locally prevalent root-mean-square (r.m.s.) contrast, the standard deviation of the stimulus luminance divided by the mean luminance. This division provides a degree of contrast invariance: rather than depending linearly on contrast, responses are reduced in those locations of the visual field where contrast is high and increased where contrast is low.

At least some components of luminance and contrast gain control should operate rapidly, because the eyes typically fixate a given location for only 200–300 ms, and eye movements will often bring the receptive fields of neurons in the early visual system over image patches that differ in luminance and contrast. What ranges of luminance and contrast are typically encountered in natural scenes? To what extent do luminance and contrast vary together? We addressed these questions by measuring luminance and contrast for image patches selected from calibrated natural images. We found that the typical ranges are substantial and that luminance and contrast are largely statistically independent. We then turned to the gain control mechanisms in the early visual system and found that their operation reflects the independence encountered in natural images. This finding indicates a close match between the statistical properties of natural scenes and the processing of luminance and contrast in visual systems^{11–15}. Moreover,

it directly supports the hypothesis^{1,2,16,17} that contrast is a fundamental independent variable encoded by the early visual system.

RESULTS

Luminance and contrast in natural scenes

We studied the changes in luminance and contrast encountered during a simulated saccadic inspection of a natural scene and found that these changes cover more than an order of magnitude (Fig. 1). We analyzed 300 calibrated natural images¹⁸ and simulated scan paths based on the measured statistics of saccades¹⁹. We considered image patches such as would be covered by a receptive field of a neuron in the early visual system (we repeated this analysis for different patch sizes). For each patch, we measured local luminance (the mean luminance in the patch) and local contrast (the r.m.s. contrast in the patch). The typical distance between fixations was sufficient to cause large changes in the local luminance and local contrast that fall within the receptive field of visual neurons (Fig. 1b). An analysis of the population of images confirmed that within an image, local luminance and local contrast typically varied by more than a factor of 10 (Fig. 2a–d, marginal distributions).

We found local luminance and contrast to be nearly statistically independent. Indeed, the joint distributions of luminance and contrast were approximately separable (Fig. 2b). Independence was confirmed by examination of the conditional probabilities: in the ranges of luminance and contrast that were most likely, all contrasts were roughly equally probable given a luminance (Fig. 2c), and all luminances were roughly equally probable given a contrast (Fig. 2d). This substantial degree of independence was reflected in the correlation between luminance and contrast, which was a low -0.2 . Independence held regardless of patch size: for patches ranging from 0.1° to 1.0° (full-width

¹The Smith-Kettlewell Eye Research Institute, San Francisco, California 94115, USA. ²Department of Psychology and Center for Perceptual Systems, University of Texas, Austin, Texas 78712, USA. Correspondence should be addressed to M.C. (matteo@ski.org).

Received 28 July; accepted 6 September; published online 13 November 2005; doi:10.1038/nn1556




Figure 1 Luminance and contrast in a natural scene. **(a)** A sequence of fixations in a natural scene. The crosses indicate fixation locations and the circles represent the corresponding locations of an arbitrary receptive field (diameter: 1°). **(b)** Enlargements of the image patches falling within the receptive field as a function of their r.m.s. contrast (ordinate) and average luminance (abscissa).

at half-height), the within-image correlations were always low, ranging from $-0.22 (\pm 0.02, \text{s.e.m.}, n = 300)$ to $-0.21 (\pm 0.02)$.

The mild negative correlation between luminance and contrast was due to the portions of the images taken up by the sky, where local luminance was high but local contrast was, on average, very low. We measured local luminance and local contrast for some of the obvious constituents of natural images: sky, foliage, ground and backlit foliage. From each image, we selected, by hand, rectangular regions containing these constituents, while excluding ambiguous regions. Approximate statistical independence was found for each of these constituents—the correlations were 0.0, 0.2, -0.1 and -0.2 for sky, foliage, ground and backlit foliage. This analysis of the constituents also showed that the

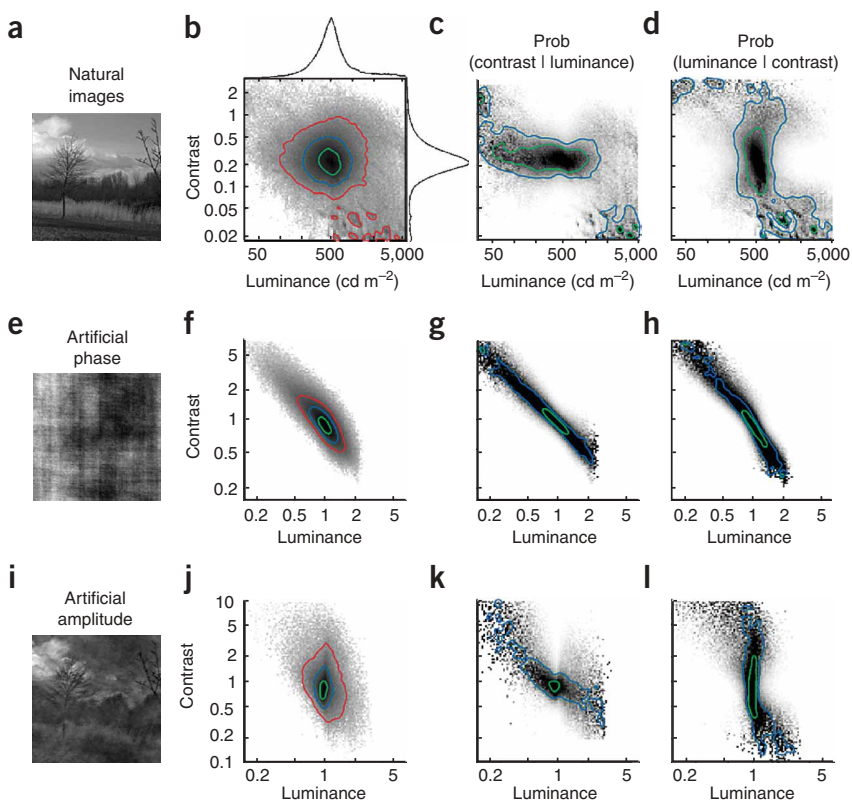
regions of high luminance and low contrast (bottom right of **Fig. 2b**) were indeed due to the sky.

To gain more insight into the observed independence of local luminance and contrast, we evaluated, separately, the effects of the phase and amplitude spectra of the natural images on the joint distribution of luminance and contrast. We assessed the role of phase by measuring local luminance and contrast for artificial-phase images (**Fig. 2e–h**). These images had the same amplitude spectrum (that is, the same autocorrelation function) as the original images, but they had a random phase spectrum²⁰. They were, approximately, samples of Gaussian noise with an amplitude spectrum that fell as $1/f$, a distribution that is typically assumed to be representative of natural images²¹. In these artificial-phase images, luminance and contrast were far from independent (**Fig. 2f**); the average within-image correlation was -0.77 ± 0.01 . The opposite was true in artificial-amplitude images (**Fig. 2i–l**). These images preserved the phases of natural images, but had an amplitude spectrum that decreased with frequency as $1/f^n$ (as fitted to the amplitude versus frequency curve of each image, $n = 1.25 \pm 0.02$). In these images, luminance and contrast were much more independent (**Fig. 2j**); the average within-image correlation was -0.05 ± 0.02 . A look at the conditional probabilities for luminance and contrast confirmed their marked dependence in artificial-phase images (**Fig. 2g,h**) and a much higher degree of independence in artificial-amplitude images (**Fig. 2k,l**). Thus, statistical independence of luminance and contrast was not trivial ($1/f$ noise does not have this property), but was dependent on the typical phase structure of natural images (see Discussion).

Luminance gain control and contrast gain control

Thanks to luminance and contrast gain control, sudden changes in luminance or contrast that occurred between fixations had a reduced

Figure 2 Statistics of local luminance and contrast in natural images. **(a)** A natural image (same as in **Fig. 1a**). **(b)** Joint distribution of luminance and contrast as sampled from all 300 images. These distributions represent the variation of luminance and contrast within a typical image: specifically, we first computed the overall average luminance and contrast across images, and then rescaled each image so that its average luminance and contrast would match the overall average. The contours delineate the regions containing 90% (red), 65% (blue) and 40% (green) of the observations. The curves on the sides of the joint distribution indicate the marginal distributions of luminance and contrast. **(c)** Conditional probability of observing a certain contrast given a specified luminance. This distribution is obtained by normalizing vertical slices of the joint distribution in **b**. **(d)** Conditional probability of observing a certain luminance given a specified contrast. This distribution is obtained by normalizing horizontal slices of the distribution in **b**. **(e)** An artificial-phase image. This image has the same amplitude spectrum as the image in **a**, but a random phase spectrum. **(f–h)** Joint distribution and conditional probability distributions for the 300 artificial-phase images. Format as in **b–d**. **(i)** An artificial-amplitude image. This image has the same phase spectrum as the image in **a**, but the amplitude at each frequency is given by the $1/f^n$ spectrum that best fits the spectrum of the original image. **(j–l)** Joint distribution and conditional probability distributions for the 300 artificial-amplitude images. Format as in **b–d**.



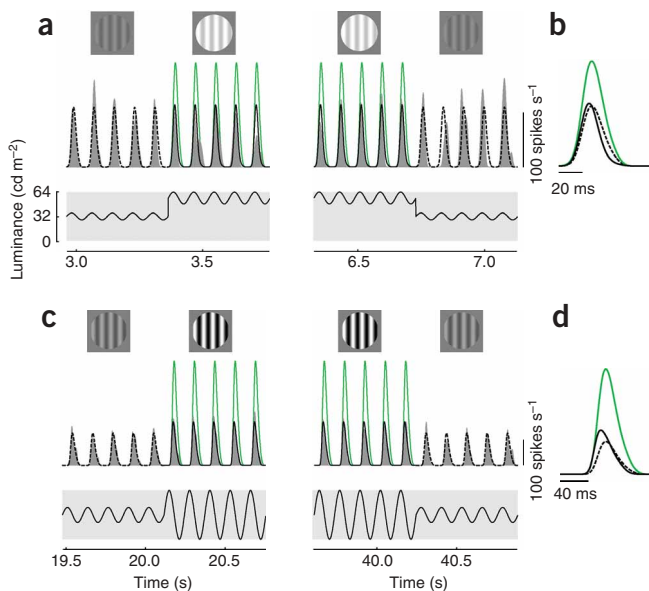


Figure 3 Effect and time course of gain control mechanisms in LGN. **(a)** Response of an LGN neuron to a drifting grating of constant contrast (14%), whose luminance steps from 32 cd m^{-2} to 56 cd m^{-2} (left) and back to 32 cd m^{-2} (right). Spatial frequency and temporal frequency (12.5 Hz) are optimal for this neuron. The temporal profile of the stimulus is shown below the responses. Histograms (gray) were obtained by convolving the spike trains with a Gaussian window ($\sigma = 5 \text{ ms}$), and averaging over three stimulus presentations. From the histograms, we computed the average response to a cycle of the stimulus before (dashed) and after (black) the step in luminance. The linear prediction (green) was obtained by scaling the response before the step (dashed) by the ratio of the two luminances. **(b)** Comparison of average responses to low luminance (dashed) and high luminance (black), and of the response expected in the absence of gain control (green). **(c)** Response of an LGN neuron to a drifting grating of constant luminance (32 cd m^{-2}) whose contrast steps from 31% to 100% (left) and back to 31% (right). Spatial frequency and temporal frequency (7.8 Hz) are optimal for this neuron. Histograms (gray) are the average over five stimulus presentations. The linear prediction (green) was obtained by scaling the response before the step (dashed) by the ratio of the two contrasts. **(d)** Comparison of average responses to low contrast (dashed) and high contrast (black), and of the response expected in the absence of gain control (green).

impact on the responses of the early visual system (Fig. 3). We recorded the responses of neurons in the lateral geniculate nucleus (LGN), which receives the output of the retina and provides input to the visual cortex. The recordings were performed extracellularly in anesthetized, paralyzed cats. LGN responses were barely affected by sudden steps in luminance (Fig. 3a) and were weakly affected by changes in contrast (Fig. 3c). The measured responses were much smaller and occurred faster than the high-luminance responses predicted by low-luminance measurements of the receptive field (Fig. 3b) or the high-contrast responses predicted by low-contrast measurements of the receptive field (Fig. 3d). These reductions in gain and the changes in dynamics occurred well within a cycle of the drifting grating (80 ms in Fig. 3a, 128 ms in Fig. 3c), confirming that the gain control mechanisms operate very quickly, in less than 100 ms ^{1,5,6,22–26}.

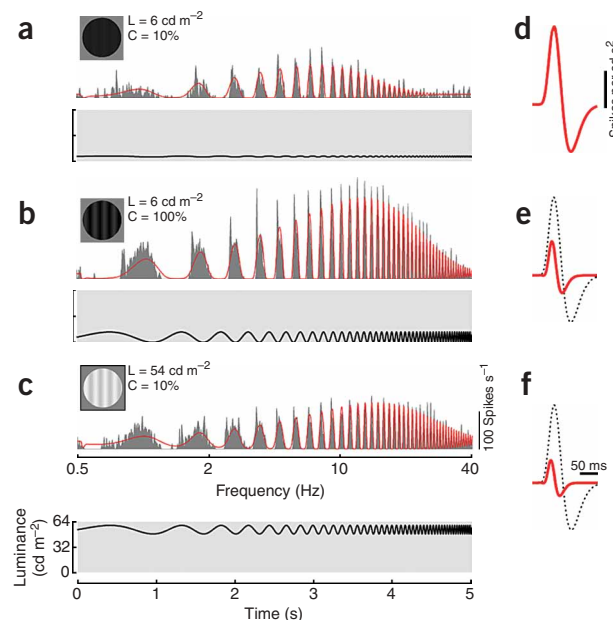
Do the mechanisms of gain control for luminance and contrast reflect the independence of luminance and contrast seen in natural images? To react appropriately to the changes in luminance and contrast, the corresponding gain control mechanisms should be functionally independent. In other words, within the range of luminances encountered during natural viewing, luminance gain control should have the same effects at all contrasts, and contrast gain control should have the same effects at all mean luminances. Instead, if the gain control mechanisms were appropriate for statistics other than those in the natural environment—for example, for those of $1/f$ noise—one would expect that contrast gain control would be biased by local luminance or that luminance gain control would be biased by local

contrast. In other words, one would expect the visual system to exploit the redundancy implicit in any lack of independence.

Independence of gain control mechanisms

To test for independence, we characterized the effects of luminance and contrast gain control in the LGN. We recorded responses to drifting gratings (Fig. 4) with mean luminance ($6\text{--}56 \text{ cd m}^{-2}$) and contrast (10–100% Michelson contrast; 0.07–0.7 r.m.s. contrast) covering a range extending over a factor of 10, similar to the excursion seen in patches of natural images (Fig. 2b). To fully quantify the effects of gain control on both the amplitude and the dynamics of the responses¹, we measured responses to a range of frequencies by increasing the drift rate exponentially with time from 0.5 Hz to 40 Hz in 5 s and back to 0 in the subsequent 5 s (Fig. 4a–c, and Supplementary Fig. 1 online). The responses to these stimuli can be read as tuning functions for stimulus temporal frequency. As expected¹, the preferred temporal

Figure 4 Characterizing LGN responses at various luminances and contrasts. **(a–c)** Responses of an LGN neuron (X-type, on-center) to temporal frequency sweeps at **(a)** low luminance and low contrast ($L = 6 \text{ cd m}^{-2}$, $C = 10\%$, Michelson contrast), **(b)** low luminance and high contrast ($L = 6 \text{ cd m}^{-2}$, $C = 100\%$) and **(c)** high luminance and high contrast ($L = 54 \text{ cd m}^{-2}$, $C = 10\%$). Histograms (gray) were obtained by averaging over ten stimulus presentations. Red curves are descriptions of the responses by the descriptive model (Fig. 5a). Stimuli were sinusoidal gratings at optimal spatial frequency (icons). The temporal profile of the stimuli is shown under the responses; drift rate increased exponentially with time, from 0.5 Hz to 40 Hz in 5 s, and back (not shown). **(d–f)** Impulse responses used for the predictions in **a–c**. The impulse response is smaller and faster at the higher contrast **(e)** or luminance **(f)** than at low luminance and contrast **(d)**, and dotted curves.



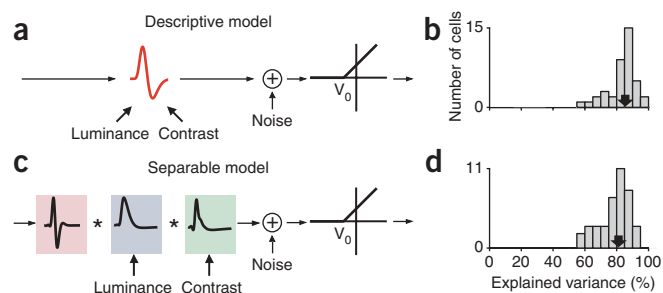
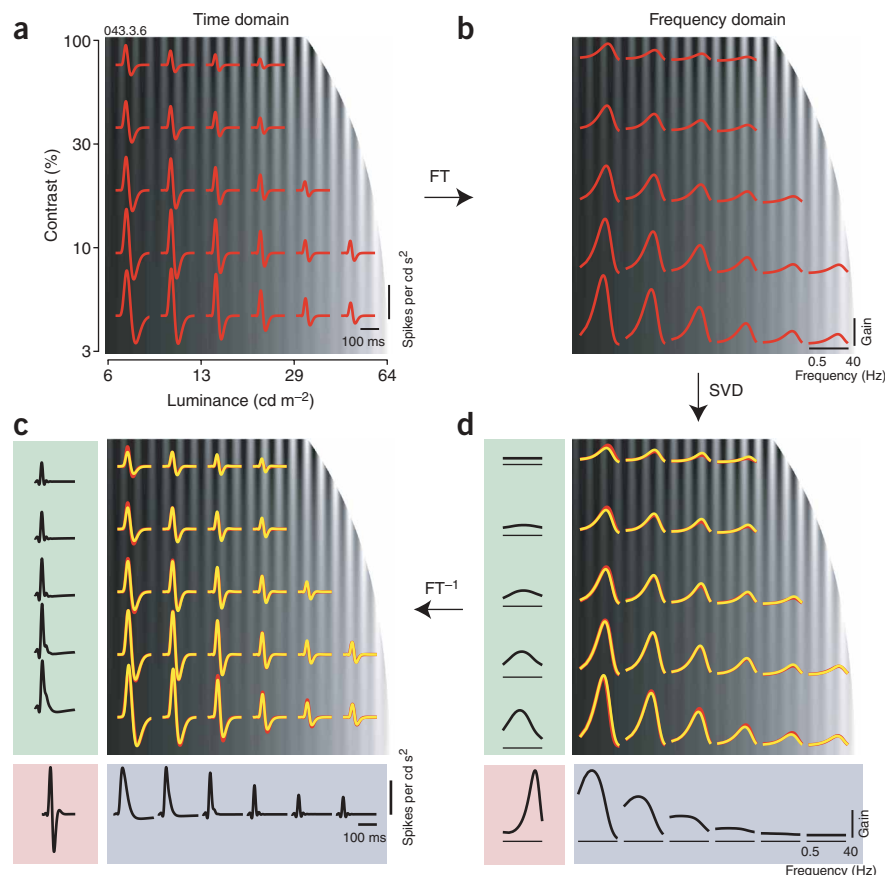


Figure 5 The two models used to describe LGN responses, and a measure of their performance. **(a)** The descriptive model: stimulus luminance is integrated by the spatial receptive field, filtered by the impulse response (red), added to Gaussian noise and rectified. **(b)** Quality of predictions for the descriptive model, measured by the percentage of stimulus-driven response variance explained by the model. $N = 40$. The median is 85% (arrow). **(c)** The separable model. The impulse response is the convolution of three filters: a fixed filter (pink), a luminance gain filter (blue) and a contrast gain filter (green). **(d)** Quality of predictions for the separable model. The median explained variance is 81% (arrow).

frequency was substantially higher at higher contrast (**Fig. 4b**) or at higher luminance (**Fig. 4c**) than at lower luminance and contrast (**Fig. 4a**).

To summarize the responses in each stimulus condition (that is, each fixed mean luminance and contrast), we fitted them with a descriptive model. In this model (**Fig. 5a**), stimulus luminance was filtered by a linear receptive field, whose output was rectified to yield positive firing rates. The impulse response (the temporal profile of the receptive field) was estimated independently for each stimulus condition (**Fig. 4d–f**). The descriptive model captured response amplitude and phase over the entire range of tested temporal frequencies (**Fig. 4a–c** and **Supplementary Fig. 2** online). For over half of the cells in our population ($N = 40$), it accounted for more than 85% of the stimulus-driven variance in the responses (**Fig. 5b**).

Figure 6 Independence of the effects of luminance gain control and contrast gain control. **(a)** Impulse responses of the cell in **Figure 4**, measured by fitting the descriptive model for all combinations of mean luminance and contrast. **(b)** The amplitude of the transfer functions corresponding to those impulse responses, as a function of frequency. **(c)** Impulse responses predicted by the separable model (yellow), compared to those predicted by the descriptive model (red, replotted from **a** for comparison). The latter are barely visible in the superposition, indicating that the predicted impulse responses are extremely similar. Each impulse response (yellow) is the convolution of the fixed filter (pink) with the luminance gain filter in the appropriate column (blue) and a contrast gain filter in the appropriate row (green). **(d)** The amplitude transfer functions corresponding to those impulse responses, as a function of frequency. The arrows between the panels indicate the sequence of operations: Fourier transform (FT, **a** to **b**), singular value decomposition (SVD, **b** to **d**), and inverse Fourier transform (FT⁻¹, **d** to **c**).



Differences between predicted and measured responses were due mostly to short transients occurring at the onset of the rising phase of a cycle. These transients corresponded to bursts of action potentials²⁷, which the model was not designed to produce.

When either mean luminance or contrast was changed, the gain control mechanisms had potent effects on the impulse response^{4,28} (**Fig. 4d–f**). One effect of increasing mean luminance or contrast was on the impulse response's amplitude, which decreased markedly. For example, the peak-to-peak amplitude was reduced by 59% by increasing the contrast (**Fig. 4e**) and by 71% by increasing the luminance (**Fig. 4d**). This decrease countered the increase in signal strength, reducing the dependence of the responses on mean luminance and contrast and adjusting the cell's dynamic range to the prevalent stimulus conditions. The other effect of increasing mean luminance or contrast was on the impulse response's time course, which became more transient. For example, the duration of the impulse response was reduced from 94 ms at low luminance and contrast (**Fig. 4d**) to 55 ms at the higher contrast (**Fig. 4e**) and to 48 ms at the higher luminance (**Fig. 4f**). This reduction modified the temporal frequency tuning of the responses, increasing the preferred temporal frequency as mean luminance and contrast increased (**Fig. 4a–c**).

Our method for measuring impulse responses yielded robust measures, which were not affected by slow contrast-adaptation mechanisms^{6,29–32}. The responses in the first 5 s, when frequency ramped up from 0.5 Hz to 40 Hz, were essentially identical to those in the subsequent 5 s, when frequency ramped back down (**Supplementary Fig. 1**). Moreover, responses commonly remained constant over 20 s of stimulation with a drifting grating of constant frequency (**Fig. 3c**).

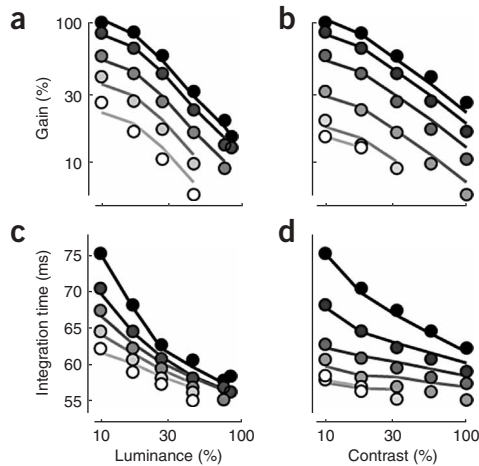


Figure 7 Summary of the effects of luminance and contrast, and predictions of the separable model. Data points indicate the overall gain (top) and integration time (bottom) for the data of **Figure 6**. Lines show the predictions of the separable model. **(a)** Overall gain as a function of mean luminance, for different contrasts (black to white: 10% to 100%). **(b)** Gain replotted as a function of contrast, for different mean luminances (black to white: 10% to 84%). **(c)** Overall integration time as a function of luminance, for different contrasts (as in **a**). **(d)** Integration time as a function of contrast, for different luminances (as in **b**).

Thus, slow contrast adaptation mechanisms operating over time courses of seconds had little (if any) role in these responses.

To test whether luminance gain control and contrast gain control operate independently of each other, we asked whether these responses could be explained by a separable model (**Fig. 5c,d**). This separable model is a special case of the descriptive model described above (**Fig. 5a**). In the model, the impulse response is described by a fixed filter followed by two variable filters: one for luminance gain control, which depends only on mean luminance, and one for contrast gain control, which depends only on contrast. This model makes a strong prediction: the effects on the impulse response of changing luminance should be the same for all contrasts, and the effects of changing contrast should be the same for all luminances.

We tested this prediction on measurements made at a variety of luminances and contrasts (**Fig. 6**). To estimate the filters in the separable model, we started from the impulse responses measured with the descriptive model (**Fig. 6a**), and we applied a series of simple mathematical operations (**Fig 6b–d**; see also **Supplementary Methods** online). The first operation was the Fourier transform (**Fig. 6b**), which expressed each impulse response as a transfer function in the frequency domain; the advantage of this representation is that the transfer function of a series of filters (the three filters in the separable model, **Fig. 5c**) is the product of the transfer functions of the filters. According to the separable model, the matrix of transfer functions (**Fig. 6b**) should be separable: each transfer function should be the product of three transfer functions—a fixed one, one that depends only on luminance and one that depends only on contrast. The best estimates for the three transfer functions were obtained through singular value decomposition (**Fig. 6d**). This operation yielded (i) a fixed transfer function (pink), (ii) transfer functions that depended only on luminance (blue) and (iii) transfer functions that depended only on contrast (green). Finally, we performed an inverse Fourier transform to convert the transfer functions back into impulse responses in the time domain. The predicted impulse responses were very similar

to the measured ones (**Fig. 6c**). For each combination of luminance and contrast, the model predicted that the impulse response was the convolution of the fixed filter with the corresponding snapshots of the filters for luminance gain control and contrast gain control (that is, the filters in the appropriate column and row). As expected, these filters become larger and slower as luminance or contrast is decreased (**Fig. 6c**).

The separable model provided excellent fits to the data. First, it predicted impulse responses that were barely distinguishable from those estimated by the descriptive model (**Fig. 6c**). This was the case not only for the on-center, X-type cell (**Fig. 6**) but also for off-center cells and Y-type cells (**Supplementary Fig. 3** online). Second, it predicted the firing rate responses almost as well as did the descriptive model (compare **Fig. 5b** and **5d**). The percentage of stimulus-driven variance explained by the two models was comparable, with a median across cells of 81% for the separable model versus 85% for the descriptive model. In fact, we chose the example cell (shown in **Fig. 6**) because the quality of the fits was the same as the median values, 81% and 85%.

This performance was notable, given that the separable model has many fewer degrees of freedom than the descriptive model. To predict the responses, the descriptive model requires 25 filters (one impulse response for each combination of mean luminance and contrast), whereas the separable model requires only 10 filters (the fixed filter plus the snapshots of the variable filters for five luminances and four contrasts; see **Methods**).

Another way to gauge the quality of the separable model was to consider what would happen if the luminance and contrast gain mechanisms were instead matched to the statistics of $1/f$ noise, for which there is an inverse relationship between luminance and contrast (**Fig. 2f**). If this were the case, the full matrix of responses should be explained by only one gain control mechanism, which could operate based on luminance alone or contrast alone. We tested this hypothesis by trying to predict the full set of responses with a one-dimensional subset of impulse responses. We used the impulse responses estimated at combinations of luminance and contrast lying close to a line with the slope observed in a $1/f$ world (**Fig. 2f**). For each luminance, the impulse response nearest to the line was used to predict the responses obtained at all contrasts. This method yielded poor fits, explaining only 35% (median) of the stimulus-driven variance of the responses, implying that our analysis was sensitive enough to reject plausible alternatives to the independence assumption.

Finally, an intuitive way to summarize the effects of gain control—and to gauge the performance of the separable model—is to consider overall measures of gain and integration time (**Fig. 7**). As a measure of overall gain, we took the average of the amplitude of the transfer function between 0.5 Hz and 15 Hz (at higher frequencies gain is barely affected by changes in luminance and contrast). This overall gain was plotted as a function of luminance (**Fig. 7a**). The slope of the curves was close to -1 in logarithmic axes, indicating that overall gain was inversely proportional to luminance¹. On the other hand, overall gain decreased more modestly with contrast (**Fig. 7b**, slope of curves is shallower than -1). Indeed, at low temporal frequencies, LGN responses were largely independent of mean luminance (**Fig. 3a,b**), whereas they did grow with contrast (**Fig 3c,d**). As an overall measure of integration time, we took the slope of the best-fitting line relating the phase of the transfer function to frequency, weighted by the amplitude at each frequency³³. As expected, integration time decreased with luminance (**Fig. 7c**) and with contrast (**Fig. 7d**). All of these effects were very well captured by the separable model (**Fig. 7**), confirming that the effects of luminance and contrast gain control are independent of each other.

DISCUSSION

We have demonstrated that luminance and contrast are largely independent in natural images, and that the mechanisms of luminance gain control and contrast gain control in the early visual system reflect this independence.

Luminance and contrast in natural scenes

The independence of luminance and contrast in natural images is not easily explained. According to the familiar view of image formation, the luminance that reaches the eye is the product of the reflectance of surfaces and an independent illuminant that changes slowly across the scene¹. We simulated this layout and modeled the reflectance of a scene as a sample of $1/f$ noise and the illuminant as an independent sample of low-pass filtered noise (for example, $1/f^n$ noise, $n > 1$). In the resulting images, local luminance and local contrast invariably showed a strong negative correlation. The correlation was weaker than for a uniform illuminant, but still substantially more negative than for natural images; thus, independence of the illuminance and reflectance functions cannot be the whole explanation.

A factor contributing to the statistical independence is that regions of a scene that have high luminance are likely to have more directional illumination, and hence higher-contrast shadows and shading. This factor contributes a positive correlation between luminance and contrast that partially cancels the expected negative correlation. Another possible factor is that the distribution of reflectance in natural environments might be skewed towards the higher values, instead of being symmetric (as in the case of $1/f$ noise).

The analysis of natural images reported here does not take into account the longer-term changes in local luminance owing to the day-night cycle or to switching of environments (for example, moving from an open field to under a forest canopy). It seems likely that these longer term changes are also uncorrelated with local contrast. Nonetheless, our results are of most relevance to the more rapid components of luminance and contrast gain control that compensate for the changes in receptive field stimulation resulting from movements of eye, head, and body. Under these circumstances, we have shown that the visual system will mostly encounter changes in luminance of about one order of magnitude (but occasionally larger—see Fig. 2b), substantially less than the many orders of magnitude over which the eye is able to operate³.

Independence of gain control mechanisms

To study the effect of luminance and contrast on responses of the early visual system, we turned to simple stimuli, the drifting gratings that are classically used to study vision. Drifting gratings of optimal spatial frequency present many advantages for our purposes. First, drifting gratings are defined by few parameters, which explicitly include luminance and contrast. Second, drifting gratings drive neurons in the early visual system very well, so that we could test low luminances and low contrasts that would otherwise elicit very few spikes. Third, drifting gratings elicit, in the receptive field, responses that oscillate at the frequency of drift, so that it is straightforward to predict the responses of the receptive field and see how these responses are affected by gain control. In principle, it would have been desirable to use natural stimuli, but describing LGN responses to such stimuli requires a rather involved model that goes well beyond the simple temporal aspects discussed here.

The independence of gain-control mechanisms for luminance and contrast that we have found may not seem surprising at first, given the common assumptions that the output of luminance gain control (i) removes all effects of mean luminance and (ii) is

completely independent of contrast. However, these assumptions are not accurate. First, the output of the luminance gain control mechanism does depend on mean luminance: it emphasizes the low temporal frequencies at low mean luminance and the high temporal frequencies at high mean luminance (Fig. 4). To interpret these responses and estimate stimulus contrast, a subsequent contrast gain-control stage would have to know the mean luminance. Second, the output of mechanisms performing luminance gain control does, to some extent, depend on contrast³⁴: for a high-contrast stimulus, luminance gain can vary over time, because mean luminance is computed locally and rapidly¹.

One limitation of our study is that we attributed all the effects of contrast gain control to changes in the impulse response of the neurons. In principle, contrast gain control might also affect the spatial receptive field³⁵ or the resting potential^{6,32}. In practice, however, a model in which contrast gain control leaves constant both the spatial receptive field and the resting potential provides an excellent fit to the responses of LGN neurons to stimuli of different contrasts, sizes and spatial frequencies¹⁰. In ganglion cells, moreover, resting potential seems to be affected only very slightly by changes in contrast²⁶.

Similarly, we have not quantified the effects of luminance gain control on the spatial receptive field and on the resting potential. The surround of retinal ganglion cells becomes relatively weaker at lower luminances^{36,37}, but this weakening takes place only at the lowest levels of luminance and is barely noticeable when one compares responses within a limited range of luminances, as we do here. Mean luminance might affect the resting potential, but this effect is small and variable, at least as gauged from the resting firing rate³⁶. Indeed, in our models we have left the resting potential free to vary with mean luminance, but the resulting estimates did not vary by much and did not depend on luminance in an orderly fashion.

Conclusions

Our results add to the growing body of evidence for a close match between the statistical properties of natural scenes and the processing of contrast in visual systems. For example, there appears to be a close correspondence between the range of local contrasts in natural images and the dynamic range of single neuron responses in the eye^{11,12} and in the LGN¹³. Similarly, the statistics of natural images, together with the observation that signal strength, compared to noise strength, is smaller at low contrasts, can predict how the shape of the impulse response changes across contrasts¹⁴. There is also computational evidence that, for natural images, rapid local contrast adaptation enhances faint contours¹² and increases statistical independence (reduces the redundancy) in the responses of orientation-selective and spatial frequency-selective neural populations in visual cortex¹⁵. Ultimately, however, a naturalistic explanation of the computational advantages of contrast gain control will have to account for the large differences seen across species and across the different retinogeniculate streams such as the M and P pathways in primates³⁸.

In summary, we have shown that there is little dependence between the local luminance and the local contrast in natural images and that this independence is elegantly reflected in the neural responses of the LGN. Our results provide direct support for the hypothesis^{1,2,16,17} that contrast is a fundamental independent variable encoded by the early visual system. They strengthen and validate a large body of neurophysiological, psychophysical and theoretical research that has implicitly assumed that luminance and contrast gain control are functionally independent, and provide a new example of how measuring the statistics of natural environments can provide insight into sensory systems.

METHODS

Measurement of local luminance and contrast. From a publicly available image bank¹⁸, we selected 300 images (12-bit, gray scale) that did not include animals or artificial objects. For each image, we considered a sequence of eye movements by sampling from eye movement distributions measured separately¹⁹ (sampling from a uniform distribution gave similar results), with the constraint that no two samples could be closer than half the diameter of the patch.

The local luminance of a patch was defined as

$$L = \sum_{i=1}^N w_i L_i$$

where N is the total number of pixels in the patch, L_i is the luminance of the i^{th} pixel and w_i is the weight from a windowing function.

The local contrast of a patch was defined as

$$C = \sqrt{\frac{\sum_{i=1}^N w_i (L_i - L)^2}{L^2}}$$

The weighting function was a circularly symmetric raised cosine

$$a_i = \cos\left(\frac{2\pi}{d} \sqrt{(x_i - x_c)^2 + (y_i - y_c)^2}\right) + 1$$

where d is the patch diameter, (x_i, y_i) is the location of the i^{th} pixel in the patch and (x_c, y_c) is the location of the center of the patch. The weights were normalized to sum to 1: $w_i = a_i / \sum_{i=1}^N a_i$.

The results described here are for a patch diameter of $d = 64$ pixels, which corresponds to approximately 1° of visual angle¹⁸; however, we obtained very similar results (that is, small correlations) for a wide range of patch sizes.

To determine the variation of local luminance and contrast within the typical single image (or constituent sub-image), we first computed the overall average luminance and contrast across all images (sub-images) and then rescaled each image (sub-image) so that its own average luminance and contrast would match the overall average. Correlations of luminance and contrast were measured on the individual images and then averaged across images.

Physiological recordings. Methods for recording from single neurons in anesthetized cats have been described elsewhere³⁹. The Animal Care and Use Committee of the Smith-Kettlewell Eye Research Institute approved all procedures.

Extracellular signals were recorded with quartz-coated platinum-tungsten microelectrodes (Thomas Recording). Firing rates were obtained by convolving spike trains with a Gaussian window (s.d.: 5 ms).

Stimuli were drifting gratings presented monocularly on a CRT-screen (refresh rate: 125 Hz). Gratings had optimal spatial frequency and position. For neurons that were strongly suppressed by large gratings, stimulus size was set to the optimal value, as measured with gratings at 50% contrast and 32 cd m⁻²; for the other neurons, the stimuli covered the entire receptive field. Combinations of mean luminance (4–6 values between 6 and 56 cd m⁻²) and contrast (3–5 values between 10% and 100%) were presented in a randomized order (12–25 stimuli repeated 6–12 times). The appearance of a grating was preceded by 2–2.5 s of uniform screen at the mean luminance of the stimulus.

We classified neurons as on-center (26 cells) or off-center (14 cells) by mapping the receptive field with rapid sequences of flashed gratings. We classified neurons as X-type (34 cells) or Y-type (6 cells) based on standard criteria⁴⁰.

Models. For each combination of mean luminance L and contrast C , we estimated the impulse response $f_{L,C}(t)$ of a neuron by fitting a descriptive model (Fig. 5a).

The first stage of the model is the convolution between the linear receptive field $h(x,y,t)$ and the stimulus $s(x,y,t)$, $v_{\text{drive}}(t) = [h * s](x_0, y_0, t)$, where $v_{\text{drive}}(t)$ is the stimulus-driven membrane potential relative to rest, and (x_0, y_0) are coordinates of the receptive field center. The stimulus $s(x,y,t)$ is the luminance

distribution obtained by subtracting the mean from the luminance shown on the screen. The receptive field $h(x,y,t)$ has center-surround organization: $h(x,y,t) = G_c(x,y) \times f(t) - G_s(x,y) \times f(t - \delta)$, where G_c and G_s are Gaussian spatial profiles for center and surround, δ is the delay between center and surround, and $f(t)$ is a difference-of-gammas temporal impulse response (see **Supplementary Methods**). The latter is identical for center and surround. The parameters of G_c and G_s , and the delay δ were fixed for a given neuron.

The second and third stages of the model add Gaussian noise $n(t)$ with fixed variance σ^2 to the visual response $v_{\text{drive}}(t)$, and rectify the result to yield a firing rate $r(t) = \lfloor v_0(L) + v_{\text{drive}}(t) + n(t) \rfloor$, where $\lfloor \cdot \rfloor$ indicates rectification and v_0 is the difference between the spiking threshold and the resting potential. The variance of the Gaussian noise was fixed for a given cell, whereas the resting potential was allowed to vary with mean luminance L to account for changes in spontaneous firing rate seen at different mean luminances. The sequence of fits used to estimate model parameters is described in **Supplementary Methods**.

In the separable model (Fig. 5c), the impulse response $f_{L,C}$ is the convolution of three filters: $f_{L,C}(t) \approx [f_0 * f_L * f_C](t)$, where f_0 is fixed, f_L depends only on mean luminance (it describes the effects of luminance gain control) and f_C depends only on contrast (it describes the effects of contrast gain control).

We estimate the three filters from the impulse responses obtained with the descriptive model. This procedure is described in **Supplementary Methods**.

Quality of predictions. To quantify how well the model predictions r_i capture the measured responses s_i (both consisting of $t = 1, \dots, M$ samples) we estimated the fraction of stimulus-driven variance in the responses accounted for by the model^{41,42}:

$$\beta = \frac{\sigma_s^2 - \sigma_e^2}{\sigma_s^2 - \sigma_n^2}$$

where

$$\sigma_s^2 = \left\langle \frac{1}{M} \sum_i s_i^2 \right\rangle$$

is the power in the response,

$$\sigma_e^2 = \left\langle \frac{1}{M} \sum_i (r_i - s_i)^2 \right\rangle$$

is the mean square distance between data and model and

$$\sigma_n^2 = \frac{d}{d-1} \left[\left\langle \frac{1}{M} \sum_i s_i^2 \right\rangle - \frac{1}{M} \sum_i \langle s_i^2 \rangle \right]$$

is an estimate of the variance in s_i that is due to noise. Angular brackets indicate the average over d presentations of the same stimulus. The mean of the responses s_i was removed before these computations.

The quantity β is an intuitive measure of fit quality, similar but superior to the commonly used ‘percentage of the variance’. A perfect model—that is, a model that is a perfect description of a system—could never account for 100% of the variance, because some of that variance is due to noise. By estimating the fraction of the variance that is due to the noise, the quantity σ_n^2 , and subtracting it from the denominator in the definition of β , one eliminates this contribution. Rather reasonably, then, a perfect model is then one that yields $\beta = 1$, or 100%. Values $\beta > 1$ indicate overfitting, and $\beta = 0$ as usual indicates that a constant value would have been better than the proposed model. The assumptions behind this estimation are minimal⁴¹: the noise should have zero mean and a non-infinite variance, and should be independent between trials. The estimator σ_n^2 is unbiased, and holds even if the variance or other property of the noise depends on the signal strength (as is common with neural signals).

Note: Supplementary information is available on the Nature Neuroscience website.

ACKNOWLEDGMENTS

We thank D. Ringach for suggesting the singular value decomposition method and J. Victor for helpful comments. The study of natural images was performed in W.S.G.’s laboratory, supported by National Eye Institute grant R01EY11747. The study of physiological responses was performed in M.C.’s laboratory,

supported by the James S. McDonnell Foundation 21st Century Research Award in Bridging Brain, Mind and Behavior.

COMPETING INTERESTS STATEMENT

The authors declare that they have no competing financial interests.

Published online at <http://www.nature.com/natureneuroscience/>
Reprints and permissions information is available online at <http://npg.nature.com/reprintsandpermissions/>

1. Shapley, R.M. & Enroth-Cugell, C. Visual adaptation and retinal gain controls. in *Progress in Retinal Research* Vol. 3 (eds. Osborne, N. & Chader, G.) 263–346 (Pergamon, London, 1984).
2. Troy, J.B. & Enroth-Cugell, C. X and Y ganglion cells inform the cat's brain about contrast in the retinal image. *Exp. Brain Res.* **93**, 383–390 (1993).
3. Rodieck, R.W. *The First Steps in Seeing* (Sinauer, Sunderland, Massachusetts, 1998).
4. Shapley, R.M. & Victor, J.D. The effect of contrast on the transfer properties of cat retinal ganglion cells. *J. Physiol. (Lond.)* **285**, 275–298 (1978).
5. Victor, J. The dynamics of the cat retinal X cell centre. *J. Physiol. (Lond.)* **386**, 219–246 (1987).
6. Baccus, S.A. & Meister, M. Fast and slow contrast adaptation in retinal circuitry. *Neuron* **36**, 909–919 (2002).
7. Demb, J.B. Multiple mechanisms for contrast adaptation in the retina. *Neuron* **36**, 781–783 (2002).
8. Kaplan, E., Purpura, K. & Shapley, R. Contrast affects the transmission of visual information through the mammalian lateral geniculate nucleus. *J. Physiol. (Lond.)* **391**, 267–288 (1987).
9. Sclar, G., Maunsell, J.H. & Lennie, P. Coding of image contrast in central visual pathways of the macaque monkey. *Vision Res.* **30**, 1–10 (1990).
10. Bonin, V., Mante, V. & Carandini, M. The suppressive field of neurons in lateral geniculate. *J. Neurosci.* (in the press).
11. Laughlin, S. A simple coding procedure enhances a neuron's information capacity. *Z. Naturforsch. [C]* **36**, 910–912 (1981).
12. Ruderman, D.L. The statistics of natural images. *Network: Comput. Neural Syst.* **5**, 517–548 (1994).
13. Tadmor, Y. & Tolhurst, D.J. Calculating the contrasts that retinal ganglion cells and LGN neurones encounter in natural scenes. *Vision Res.* **40**, 3145–3157 (2000).
14. Van Hateren, J.H. Spatiotemporal contrast sensitivity of early vision. *Vision Res.* **33**, 257–267 (1993).
15. Schwartz, O. & Simoncelli, E.P. Natural signal statistics and sensory gain control. *Nat. Neurosci.* **4**, 819–825 (2001).
16. Shapley, R. The importance of contrast for the activity of single neurons, the VEP and perception. *Vision Res.* **26**, 45–61 (1986).
17. Shapley, R.M. & Man-Kit Lam, D. (eds.). *Contrast Sensitivity* (Bradford Books, 1993).
18. van Hateren, J.H. & van der Schaaf, A. Independent component filters of natural images compared with simple cells in primary visual cortex. *Proc. Biol. Sci.* **265**, 359–366 (1998).
19. Najemnik, J. & Geisler, W.S. Optimal eye movement strategies in visual search. *Nature* **434**, 387–391 (2005).
20. Oppenheim, A.V. & Lim, J.S. The importance of phase in signals. *Proc. IEEE.* **69**, 529–541 (1981).
21. Field, D.J. Relations between the statistics of natural images and the response properties of cortical cells. *J. Opt. Soc. Am. A* **4**, 2379–2394 (1987).
22. Saito, H. & Fukada, Y. Gain control mechanisms in X- and Y-type retinal ganglion cells of the cat. *Vision Res.* **26**, 391–408 (1986).
23. Lankheet, M.J., Van Wezel, R.J., Prickaerts, J.H. & van de Grind, W.A. The dynamics of light adaptation in cat horizontal cell responses. *Vision Res.* **33**, 1153–1171 (1993).
24. Yeh, T., Lee, B.B. & Kremers, J. The time course of adaptation in macaque retinal ganglion cells. *Vision Res.* **36**, 913–931 (1996).
25. Lee, B.B., Dacey, D.M., Smith, V.C. & Pokorny, J. Dynamics of sensitivity regulation in primate outer retina: the horizontal cell network. *J. Vis.* **3**, 513–526 (2003).
26. Zaghoul, K.A., Boahen, K. & Demb, J.B. Contrast adaptation in subthreshold and spiking responses of mammalian Y-type retinal ganglion cells. *J. Neurosci.* **25**, 860–868 (2005).
27. Sherman, S.M. Tonic and burst firing: dual modes of thalamocortical relay. *Trends Neurosci.* **24**, 122–126 (2001).
28. Tranchina, D., Gordon, J. & Shapley, R.M. Retinal light adaptation - evidence for a feedback mechanism. *Nature* **310**, 314–316 (1984).
29. Smirnakis, S.M., Berry, M.J., Warland, D.K., Bialek, W. & Meister, M. Adaptation of retinal processing to image contrast and spatial scale. *Nature* **386**, 69–73 (1997).
30. Chander, D. & Chichilnisky, E.J. Adaptation to temporal contrast in primate and salamander retina. *J. Neurosci.* **21**, 9904–9916 (2001).
31. Brown, S.P. & Masland, R.H. Spatial scale and cellular substrate of contrast adaptation by retinal ganglion cells. *Nat. Neurosci.* **4**, 44–51 (2001).
32. Solomon, S.G., Peirce, J.W., Dhruv, N.T. & Lennie, P. Profound contrast adaptation early in the visual pathway. *Neuron* **42**, 155–162 (2004).
33. Reid, R.C., Victor, J.D. & Shapley, R.M. Broadband temporal stimuli decrease the integration time of neuron in cat striate cortex. *Vis. Neurosci.* **9**, 39–45 (1992).
34. Lankheet, M.J., van Wezel, R.J. & van de Grind, W.A. Light adaptation and frequency transfer properties of cat horizontal cells. *Vision Res.* **31**, 1129–1142 (1991).
35. Nolt, M.J., Kumbhani, R.D. & Palmer, L.A. Contrast-dependent spatial summation in the lateral geniculate nucleus and retina of the cat. *J. Neurophysiol.* **92**, 1708–1717 (2004).
36. Derrington, A.M. & Lennie, P. The influence of temporal frequency and adaptation level on receptive field organization of retinal ganglion cells in cat. *J. Physiol. (Lond.)* **333**, 343–366 (1982).
37. Troy, J.B., Bohnsack, D.L. & Diller, L.C. Spatial properties of the cat X-cell receptive field as a function of mean light level. *Vis. Neurosci.* **16**, 1089–1104 (1999).
38. Lennie, P. Parallel visual pathways: a review. *Vision Res.* **20**, 561–594 (1980).
39. Freeman, T.C., Durand, S., Kiper, D.C. & Carandini, M. Suppression without inhibition in visual cortex. *Neuron* **35**, 759–771 (2002).
40. Hochstein, S. & Shapley, R.M. Quantitative analysis of retinal ganglion cell classifications. *J. Physiol. (Lond.)* **262**, 237–264 (1976).
41. Sahani, M. & Linden, J.F. in *Advances in Neural Information Processing Systems* (eds. Becker, S., Thrun, S. & Obermayer, K.) 125–132 (MIT Press, Cambridge, Massachusetts, 2003).
42. Machens, C.K., Wehr, M.S. & Zador, A.M. Linearity of cortical receptive fields measured with natural sounds. *J. Neurosci.* **24**, 1089–1100 (2004).