



Local luminance and contrast in natural images

Robert A. Frazor¹, Wilson S. Geisler^{*}

Department of Psychology and Center for Perceptual Systems, University of Texas at Austin, Austin, TX 78712, USA

Received 28 February 2005

Abstract

Within natural images there is substantial spatial variation in both local contrast and local luminance. Understanding the statistics of these variations is important for understanding the dynamics of receptive field stimulation that occur under natural viewing conditions and for understanding the requirements for effective luminance and contrast gain control. Local luminance and contrast were measured in a large set of calibrated 12-bit gray-scale natural images, for a number of analysis patch sizes. For each image and patch size we measured the range of contrast, the range of luminance, the correlation in contrast and luminance as a function of the distance between patches, and the correlation between contrast and luminance within patches. The same analyses were also performed on hand segmented regions containing only “sky”, “ground”, “foliage”, or “backlit foliage”. Within the typical image, the 95% range (2.5–97.5 percentile) for both local luminance and local contrast is somewhat greater than a factor of 10. The correlation in contrast and the correlation in luminance diminish rapidly with distance, and the typical correlation between luminance and contrast within patches is small (e.g., -0.2 compared to -0.8 for $1/f$ noise). We show that eye movements are frequently large enough that there will be little correlation in the contrast or luminance on a receptive field from one fixation to the next, and thus rapid contrast and luminance gain control are essential. The low correlation between local luminance and contrast implies that efficient contrast gain control mechanisms can operate largely independently of luminance gain control mechanisms.

© 2005 Elsevier Ltd. All rights reserved.

1. Introduction

When we explore a natural environment with our eyes, the local contrast and the local luminance that fall within the receptive field of a given visual neuron change from one fixation to the next. Further, the eyes typically fixate a given location for only 200–300 ms, and hence these changes in contrast and luminance typically occur at a rapid pace. For example, Fig. 1 illustrates the changes in contrast and luminance that would be expected during saccadic inspection of a natural scene. The “plus” signs represent a sequence of fixation locations, and the “circles” represent the corresponding sequence of locations of an arbitrary receptive field of 1 deg diameter. Enlargements of the image patches falling within the receptive field are

shown around the outside of the scene. Each of these image patches is labeled with the point in time it fell within the receptive field, with its luminance, and with its root-mean-squared (RMS) contrast; time proceeds in clockwise fashion around the figure beginning at the top. As can be seen, the contrast and luminance change from fixation to fixation.

Presumably the statistical characteristics of these variations in local contrast and luminance have had a substantial influence, through natural selection, on the design of the contrast and luminance gain control mechanisms in the visual system. Therefore, appropriate analyses of the statistical properties of natural images may be of considerable value for understanding and predicting the functional behavior of contrast and luminance gain control.

There is much circumstantial evidence for a tight linkage between the statistics of natural scenes and the design of the visual system (Atick & Redlich, 1992; Bell & Sejnowski, 1997; Field, 1987; Geisler, Perry, Super, & Gallogly, 2001; Laughlin, 1981; Olshausen & Field, 1997; Purves & Lotto,

^{*} Corresponding author. Tel.: +1 512 471 5380; fax: +1 512 471 7356.
E-mail address: geisler@psy.utexas.edu (W.S. Geisler).

¹ Present address: Smith Kettlewell Eye Research Institute, San Francisco, CA, USA.

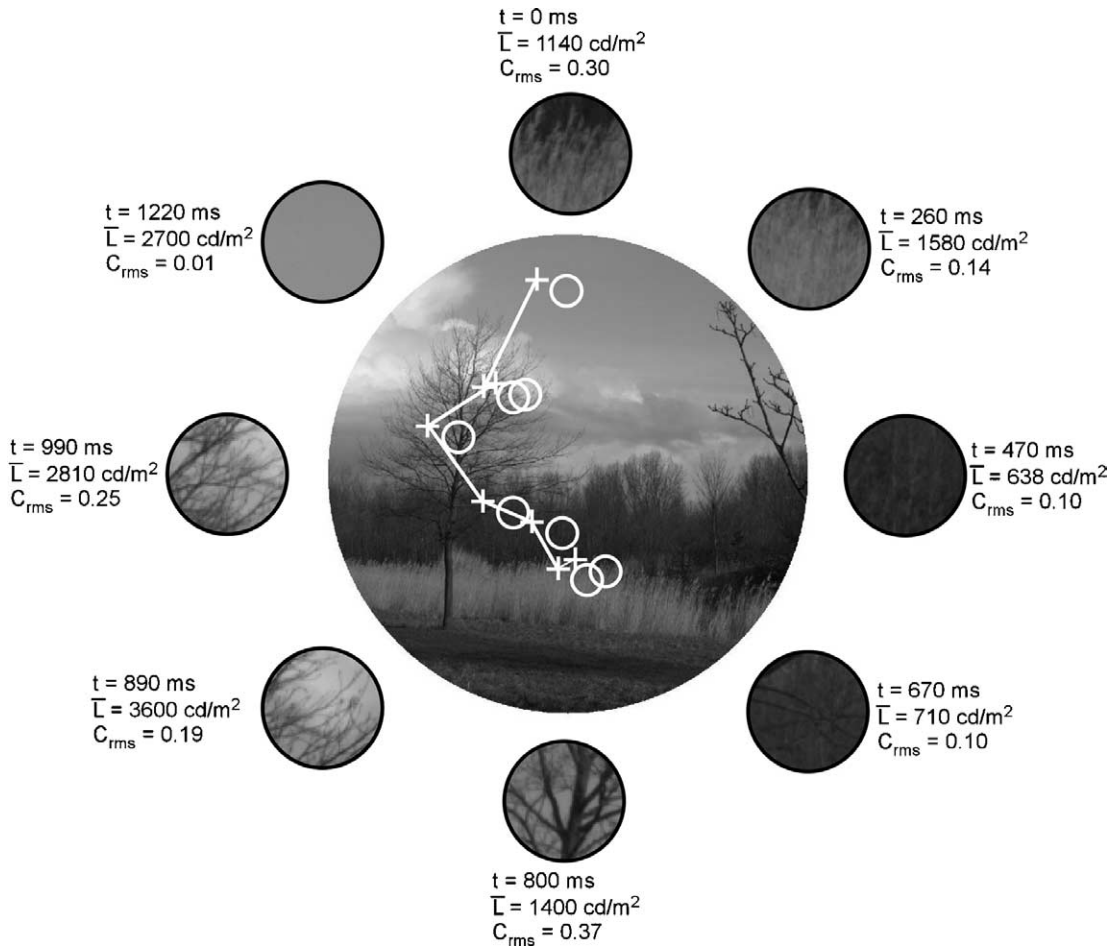


Fig. 1. Demonstration of the variation in contrast and luminance that might fall on a receptive field during a sequence of eye fixations. The plus signs show a random sequence of fixations created by sampling from eye movement histograms measured by Najemnik and Geisler (2005). Specifically, the successive eye positions were obtained by randomly sampling from the histogram of distances between fixations, and the length of time the eye stayed at a given position was obtained by randomly sampling from the histogram of fixation durations. The circles show a receptive field (1 deg in diameter) at an arbitrary location relative to the fixation point.

2003; Ruderman, 1994; Tolhurst, Tadmor, & Chao, 1992; van Hateren, 1992; van Hateren & van der Schaaf, 1998; for reviews see Simoncelli & Olshausen, 2001; Geisler & Diehl, 2002). Luminance and contrast are fundamental stimulus dimensions, and hence their statistics have received considerable attention. A number of studies have been concerned with measuring the distribution of local contrast in natural images and comparing it to the shape of contrast response functions in the eye (Laughlin, 1981; Ruderman, 1994), lateral geniculate nucleus (Tadmor & Tolhurst, 2000), and primary visual cortex (Brady & Field, 2000; Clatworthy, Chirimuuta, Lauritzen, & Tolhurst, 2003). Other studies have been concerned with characterizing the distributions of contrast in different environments (Balboa & Grzywacz, 2003) or the distribution of contrast at the center of gaze (Reinagel & Zador, 1999).

Although the present study is of some relevance to these issues (see Section 4) our primary aim was to obtain a better understanding of the statistical properties of receptive field stimulation during typical saccadic inspection, and hence to obtain a better understanding of the functional

requirements for effective luminance and contrast gain control. Specifically, we measured the variation and covariation of local contrast and luminance in natural images as a function of analysis patch size, distance between patches, and general type of image region (“foliage”, “ground”, “sky”, etc.). A subset of the measurements reported here are described in Mante, Bonin, Frazor, Geisler, and Carandini (2005).

2. Methods

Local luminance and contrast were measured in a set of calibrated natural images. The image set consisted of 300 “rural” images (i.e., minimum of manmade objects or animals) and 100 “urban” images (i.e., taken within a city environment) from a publicly available image database (van Hateren & van der Schaaf, 1998; the images may be obtained at <http://hlab.phys.rug.nl/archive.html>). The images were obtained with a Kodak DCS420 digital camera and were calibrated to result in approximately 12-bit values that are linear with respect to luminance. Complete details of the calibration procedures are given elsewhere (van Hateren & van der Schaaf, 1998). Scale factors, provided at the publicly available web site, were then used to convert the images from linear pixel values to linear luminance values (although, as noted by van Hateren and van der Schaaf,

the spectral sensitivity of the camera is not identical to human photopic spectral sensitivity). The 1536 by 1024 images were cropped to the center 1024 by 1024 pixels.

Images for this study were selected from the full set based upon a number of criteria. Images that were blurry or had a very narrow depth of field were not considered. Rural images were required to not contain humans, animals or manmade objects (including paved roads), unless they were small or at a great distance so that they occupied a small percentage of the whole image. Certain other images were removed from consideration because of their uniqueness (e.g., images dominated by a single tree trunk, or dominated by large bodies of water with specular reflections). Because the majority of the remaining images in the database are dominated by foliage, and because we were interested in what contributions to local contrast and luminance are made by various physical constituents of a scene, we qualitatively divided the remaining images into subsets based upon what kinds of physical constituents were in the image (i.e., foliage, ground, sky, or their various combinations). Fixed numbers of images from each of the subsets were then randomly selected.

The above method of selecting images was not perfectly objective, and may not be representative of the frequency with which the various kinds of image region are encountered in the environment. However, it did provide a wide variety of natural images that allowed us to evaluate how the various kinds of physical constituents contribute to the distributions of local luminance and contrast.

Local luminance and contrast were measured in image patches formed by windowing with a circularly symmetric raised cosine weighting function:

$$w_i = 0.5 \left(\cos \left(\frac{\pi}{p} \sqrt{(x_i - x_c)^2 + (y_i - y_c)^2} \right) + 1 \right), \quad (1)$$

where p is the patch radius, (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. Four different analysis patch radii were used (8, 16, 32, and 64 pixels). van Hateren and van der Schaaf (1998) report that each pixel corresponds to approximately 1 min of arc, thus the image patches have diameters of approximately 0.26, 0.54, 1.06, and 2.14 deg, respectively.

For each image and image patch size, image patch locations were selected by random sampling from an image, with the restriction that the center-to-center spacing between all selected patches exceeded the patch radius. The process of image patch selection from a given image continued until the restriction on the center-to-center spacing prohibited the selection of any additional patches. We used random sampling because it eliminates (statistically) many of the biases that can occur with systematic sampling schemes.

The local luminance and the root-mean-squared (RMS) contrast of each patch (weighted by the raised cosine window) were measured. The local luminance of a patch is defined by

$$L = \frac{1}{\sum_{i=1}^N w_i} \sum_{i=1}^N w_i L_i, \quad (2)$$

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 of the raised cosine windowing function at the i th pixel. The RMS contrast of the patch is defined by

$$C_{\text{rms}} = \sqrt{\frac{1}{\sum_{i=1}^N w_i} \sum_{i=1}^N w_i \frac{(L_i - L)^2}{L^2}}. \quad (3)$$

We chose to measure RMS contrast (as opposed to some other definition of contrast) for several reasons: (1) it is a standard measure, (2) it has been used in contrast normalization models of cortical cell responses, and (3) it predicts human contrast detection thresholds for both natural scene patches and laboratory stimuli quite well and better than other common measures of contrast (see, for example, Bex & Makous, 2002; Watson, 2000).

In rare cases (e.g., when a relatively dark region had a small, but very bright region), very high RMS contrasts were obtained.² Although these outlier cases are rare, they point to a potential weakness of the RMS con-

trast measure as a plausible measure of the potential effectiveness of a stimulus in driving contrast adaptation. To evaluate the effect of this weakness in the standard RMS contrast measure, we also measured local contrast with a slightly modified version

$$C_{\text{rms}} = \sqrt{\frac{1}{\sum_{i=1}^N w_i} \sum_{i=1}^N w_i \frac{(L_i - L)^2}{(L + L_0)^2}}, \quad (4)$$

where L_0 is a “dark light” parameter, chosen to be 7 td (1 cd/m² assuming a 3 mm pupil), based on human (photopic) intensity discrimination data (e.g., Hood & Finkelstein, 1986). This dark light parameter takes into account the reduction in visual sensitivity at low luminance, which is due (presumably) to spontaneous neural activity and other sources of internal noise. As it turned out, using this modified measure had very little effect on the results of the data analyses and had no impact on the global trends.

For comparison with the contrast response functions of neurons in striate visual cortex, we also measured local contrast using a band-limited measure. Although striate cortex neurons often reach their maximum response (saturate) at low contrasts, their response tuning functions are approximately invariant as a function of contrast, for the dimensions of spatial frequency, orientation and phase (e.g., Albrecht & Hamilton, 1982; Geisler & Albrecht, 1997; Sclar & Freeman, 1982; see Section 4). Further, there is much evidence that both the invariant tuning and the half-saturation contrast of striate neurons are due to a fast acting contrast gain control (“normalization”) mechanism (e.g., Albrecht & Geisler, 1991; Heeger, 1991, 1992; see Section 4). In order to have invariant response tuning, the spatial frequency, orientation and phase tuning of the contrast normalization mechanism must be quite broad. If it were completely broad (flat), then RMS contrast would be an appropriate contrast measure of natural scenes to compare with the half-saturation contrast of cortical neurons. On the other hand, if the normalization mechanism were less broadly tuned, then a band-limited RMS contrast measure would presumably be more appropriate, because a smaller band-limited contrast is what the normalization mechanism would be encoding. The tuning functions of contrast normalization are uncertain, but they must be broad enough to allow invariant tuning, and thus they must (from computational considerations) be at least twice the bandwidth of a neuron’s response tuning functions. The average spatial frequency bandwidth of cortical neurons is approximately 1.5 octaves (De Valois, Albrecht, & Thorell, 1982) and the average orientation bandwidth is approximately 40 deg (De Valois, Yund, & Hepler, 1982). Therefore, in computing band-limited RMS contrast, each rural and urban image was filtered in the Fourier domain with log Gabor transfer functions (both even and odd phase) that had a 3 octave spatial frequency bandwidth and an 80 deg orientation bandwidth. After inverse Fourier transformation, the mean luminance of the image was restored, the local RMS contrast was measured as described above, and then combined from the even and odd phase filters. The peak spatial frequency of the log Gabor transfer function was set to two cycles per analysis patch width, and for each analysis patch width the measurements were made for peak orientations of 0, 45, 90, and 135 deg. The measurements were averaged across the four peak orientations.

In some of the analyses, we measured the joint statistics of the local contrast (or luminance) as a function of the distance between the centers of pairs of patches. For example, we measured how the correlation between the contrasts of two patches depends on the distance between the patches. To do this the patch pairs were binned as a function of distance. The distance bins were spaced by the radius of the image patch.

All of the analyses were carried out on whole images (both rural and urban). The analyses were also carried out for different kinds of physical constituents of the natural images. To do this each rural image was hand segmented into rectangular regions that contained only one kind of constituent: sky, ground, foliage, or backlit foliage (i.e., foliage where the background is primarily sky rather than foliage or ground). Fig. 2 shows a typical image that has been hand segmented in this fashion. Some images contained all kinds of physical constituents, but many contained only a subset. The segmentation judgments were subjective, but we tried to be as conservative as possible; that is, we minimized the contamination of one variety of physical constituent (e.g., foliage) with others (e.g., sky).

² For example, the RMS contrast of a delta function is infinite.




Fig. 2. Example of hand segmentation of an image into regions containing “sky”, “foliage”, “ground”, and “backlit foliage”.

Regions of an image that were ambiguous in terms of our categories were not included in the analysis. For each kind of constituent, all of the rectangular regions obtained from all the rural images were analyzed in the same way as the whole images.

3. Results

Fig. 3 shows the measurements of local luminance. Fig. 3A plots the average log luminance as a function of patch size. (Here, all log units are base 10.) As can be seen, average local luminance is independent of patch size, and it varies across the kind of physical constituent in the order one might expect intuitively (e.g., sky has the highest luminance, foliage the lowest). Fig. 3B shows the average range of local luminance within images, where the range represents a 95% range—the difference between the 97.5 percentile log luminance and 2.5 percentile log luminance. These ranges were computed separately for each image (or rectangular region) and the ranges averaged. Only the range is plotted because the distribution of log luminance was

found to be fairly symmetrical about the mean log luminance. Fig. 3B shows that the typical 95% range of local luminance within both rural and urban images is approximately an order of magnitude. Within foliage regions that are backlit with sky, the range is also nearly an order of magnitude, but decreases sharply with image patch size. Within foliage and ground regions the range is a factor of approximately 3, and within sky regions the range is a factor of approximately 2. These are substantial ranges and hence a visual system could potentially benefit from having rapid local luminance gain control mechanisms that could operate within the time frame of a single fixation.

The full luminance ranges are larger, usually more than 2 log units (see Fig. 5). In addition, there are some full luminance ranges that exceed the dynamic range of the camera (e.g., scenes with deep shadows and specular highlights). However, the fraction of pixels where this occurs is very small. Fig. 3B shows that in the typical image the vast majority of local luminance values are within a log unit of each other.

Fig. 4 shows the measurements of local contrast. Fig. 4A plots the average RMS contrast as a function of patch size and Fig. 4B plots the average band-limited RMS contrast. For both the rural and urban images the average RMS contrast is approximately 0.2 for small patch sizes and increases monotonically. Essentially the same pattern is observed for image regions containing only foliage or ground. Not surprisingly, the RMS contrast is considerably higher for image regions containing only backlit foliage, and considerably lower for image regions containing only sky. A similar pattern of results was obtained for band-limited contrast. Unlike RMS contrast, band-limited contrast is relatively constant with patch size, especially for rural images.

Figs. 4C and D show the average 95% range of local RMS contrast within images. Because the distribution of local contrasts is not symmetric about the mean, Fig. 4C plots the lower end of the range (2.5 percentiles) and Fig. 4D plots the upper end of the range (97.5 percentiles). The 95% range of contrasts in the average rural or urban

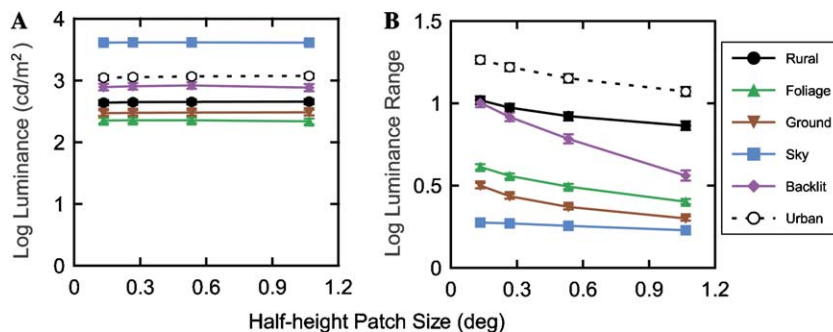


Fig. 3. Summary plots of local luminance as a function of image patch size and type of image region. The definition of local luminance is given by Eq. (2) in the text. (A) Average local luminance in \log_{10} units as function of patch size. Data points represent averages across all patches. Error bars represent ± 1 standard error computed across images. (B) Average range of local luminance within an image. The range is defined as the difference between the \log_{10} luminance at the 97.5 percentile and the \log_{10} luminance at the 2.5 percentile. Data points represent averages across images. Error bars represent ± 1 standard error computed across images.

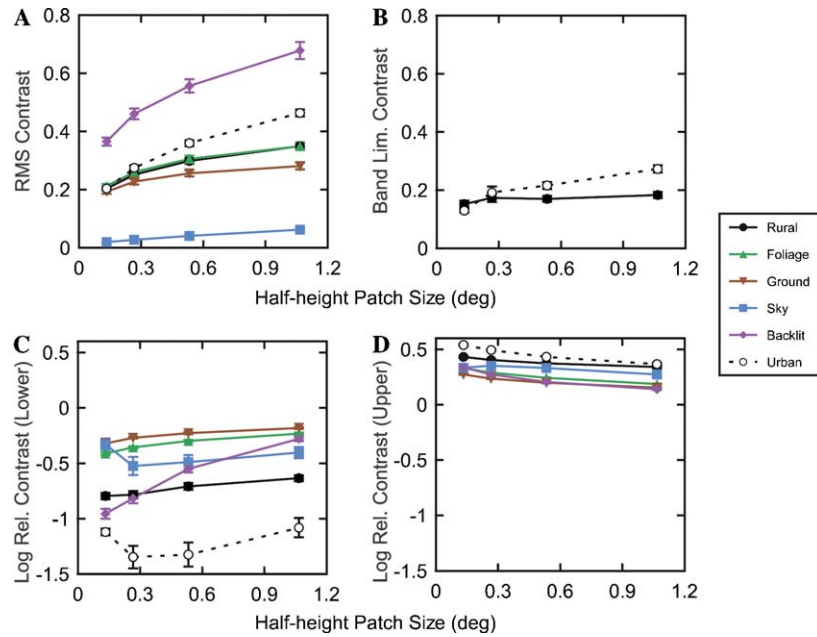


Fig. 4. Summary plots of local contrast as a function of image patch size and type of image region. (A) Average local RMS contrast. (B) Average band-limited RMS contrast (see Section 2 for definition of band-limited contrast). (C) Lower bound of 95% confidence interval of relative RMS contrast in \log_{10} units. (D) Upper bound of 95% confidence interval of relative RMS contrast in \log_{10} units. Error bars represent ± 1 standard error of the mean computed across images.

image is greater than an order of magnitude. As with local luminance, these are substantial ranges and hence a visual system could potentially benefit from having rapid local contrast gain control mechanisms that could operate within the time frame of a single fixation.

Figs. 3 and 4 show that there are substantial variations in both local luminance and local contrast within natural images. Thus, an important issue in evaluating potential luminance and contrast gain control mechanisms is the degree to which these variations in local luminance and contrast are correlated. If they are uncorrelated, then local contrast and luminance gain control mechanisms could potentially operate (and evolve) independently. For example, the contrast gain control mechanism would not need to take into account the local luminance. On the other hand, if they are highly correlated, then efficient luminance gain control and efficient contrast gain control mechanisms would need to share the same information, in order to exploit the redundancy implicit in the correlation. In this case, an efficient contrast gain control mechanism would presumably need to take into account local luminance.

Fig. 5 plots the average joint distributions of luminance and RMS contrast for a patch diameter of 0.54 deg. To obtain these distributions, each patch luminance and contrast was normalized by the mean patch luminance and mean patch contrast for that image. Then, the values were pooled across all the images, and the result scaled to match the mean patch luminance and the mean patch contrast across all images. Thus, the results in Fig. 5 are representative of a given single image. The contours in the plots show the regions corresponding to 40%, 65%, and 90% of the

volume under the distribution. There is relatively little systematic relationship between luminance and contrast for either the full images or their constituents. The only obvious asymmetry is the clusters of high probability at high luminance and low contrasts in the rural images, which appear to be due to sky.

Fig. 6 plots measurements of the correlation between luminance and contrast as a function of patch size. These are average correlations that were obtained by computing the correlation separately for each image and then averaging across images. The correlations are relatively small, but significant. For both rural and urban images, and for backlit foliage, there is a slight negative correlation of approximately -0.2 ; for ground there is an even smaller negative correlation of approximately -0.1 ; for sky the correlation is approximately 0; for foliage there is a slight positive correlation of approximately 0.15. These results suggest that local contrast gain control mechanisms could be efficient without taking into account the local luminance.

Interestingly, the low correlation between luminance and contrast is a result of the phase structure of real images. To examine the effect of the phase structure we randomized the phase spectrum of each natural image and then repeated the correlation measurements. The method of phase randomization was as follows: (1) generate a Gaussian white noise image and take its Fourier transform, (2) take the Fourier transform of the natural image, and determine its amplitude spectrum, (3) replace the amplitude spectrum of the white noise image with the amplitude spectrum of the natural image, and then take the inverse Fou-

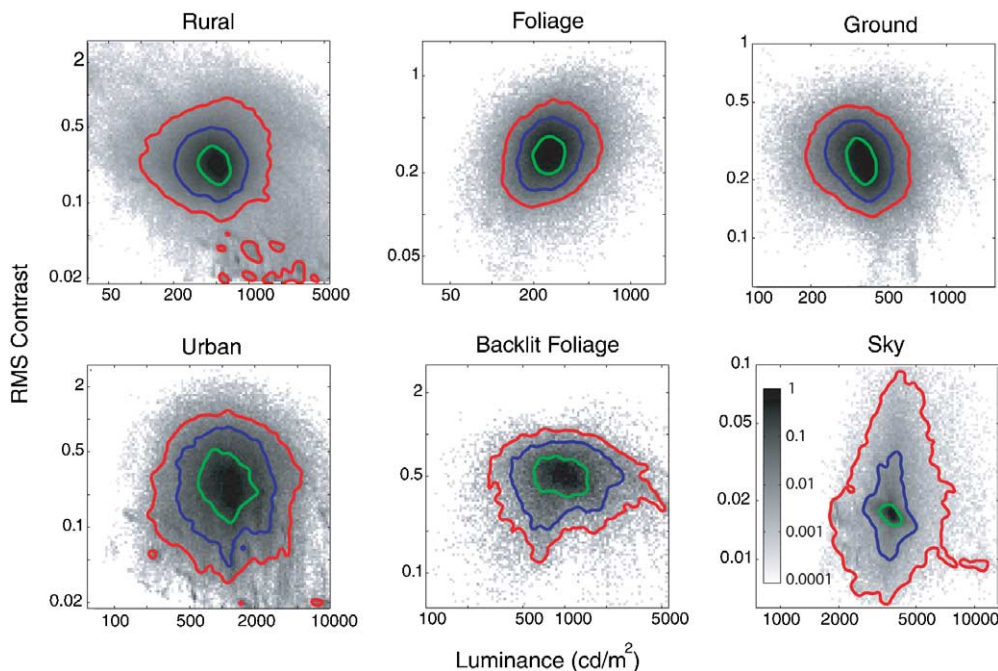


Fig. 5. Joint probability distributions of contrast and luminance for a patch diameter of 1 deg, for each type of image region. These distributions represent the variation of luminance and contrast within a typical image region; specifically, we first computed the overall average luminance and contrast across image regions, and then rescaled each image so that its average luminance and contrast would match the overall average. The contours delineate the areas containing 90% (red), 65% (blue), and 40% (green) of the observations.

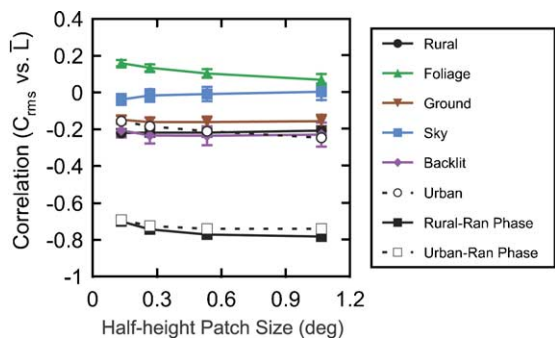


Fig. 6. Correlation between local luminance and local RMS contrast as a function of analysis patch size. The black and open squares show the correlations for images where the spatial phases have been randomized; a correlation of approximately -0.8 is also obtained for $1/f$ noise.

rier transform, and (4) scale the resulting image about its mean to eliminate any negative values. As can be seen in Fig. 6, there is a strong negative correlation (-0.7 to -0.8) between local luminance and RMS contrast in the phase-randomized images. Given that the amplitude spectra of natural images fall roughly as $1/f$ (Burton & Moorehead, 1987; Field, 1987) it is not surprising that we also obtained a correlation of approximately -0.8 for $1/f$ noise images. These results suggest that $1/f$ noise is not a good model of natural image statistics, at least for the purpose of understanding the computational requirements of luminance and contrast gain control mechanisms. Later, we describe several factors contributing to the low correlation between local luminance and contrast.

The dynamic requirements of luminance and contrast gain control mechanisms should depend upon the frequency and magnitude of changes in luminance and contrast that fall within a neuron's receptive field. If the changes are frequent and large then the gain control needs to be rapid and powerful. The frequency of saccadic eye movements (3–5 per second) implies that there are frequent changes, and we have seen that there are large variations in local luminance and contrast within rural and urban images. However, whether or not there are frequent large changes in the luminance and contrast falling within a receptive field depends on how rapidly local luminance and contrast vary across space. If they vary gradually across space, relative to the average distance between fixations, then the changes will be small and hence more sluggish gain control mechanisms might be adequate.

To evaluate how rapidly local luminance and contrast vary across space we measured pair-wise correlations as a function of distance. Fig. 7 plots the distance between image patches where the correlation falls to a value of 0.25, which we call the *decorrelation distance*. (Note that a correlation of 0.25 implies that the percentage of variation in one patch predicted by the other patch is about 6%.) For rural and urban images, the decorrelation distance for contrast is about 2 deg for small patch sizes and increases slightly with patch size. Interestingly, for all of the constituents of the rural images, the decorrelation distance for contrast is almost exactly the same—increasing from about 1 deg for the smallest patch size to about 2 deg for largest patch size. The decorrelation distance

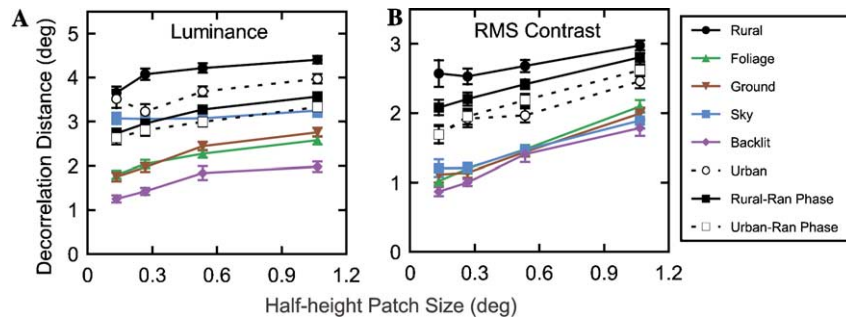


Fig. 7. Distance between image patches where the correlation drops on average to a value of 0.25, as a function of image patch size. (A) The correlations between the luminances of the image patches. (B) The correlations between the RMS contrasts of the image patches.

for luminance is 1.5–2 times larger than it is for contrast, but it varies less with patch size. Also, the decorrelation distance for luminance varies more across the constituents of the images than it does for contrast. Finally, note that the decorrelation distance is a little smaller for the phase-randomized images than for the original images. Given that the mean saccade length in complex search tasks is greater than 3.0 deg (see Section 4), it seems safe to conclude that there will often be little correlation between the contrasts (or, the luminances) within a given receptive field before and after a saccade.

4. Discussion

To obtain a better understanding of the statistical properties of receptive-field stimulation during typical saccadic inspection, and to obtain a better understanding of the functional requirements for effective luminance and contrast gain control, we measured the variation and the covariation of local contrast and local luminance in natural images as a function of analysis patch size, distance between patches, and type of image region (“foliage”, “ground”, “sky”, and “backlit foliage”). We found that (1) variations in local luminance and contrast within a given image are substantial (Figs. 3 and 4), (2) local luminance and contrast at the same spatial location are relatively uncorrelated within a image, which is not true for noise with the same amplitude spectrum as the natural image (Figs. 5 and 6), (3) the correlation between local contrasts falls rapidly with spatial distance (Fig. 7A), (4) the correlation between local luminances falls rapidly with spatial distance, but less rapidly than for contrast (Fig. 7B), and (5) the above hold for the different types of image region and for different patch sizes.

4.1. Eye movements and rapid contrast gain control

The average distance between eye fixations depends on the particular task that is being performed. If an observer is performing a task such as inspecting a small detail of some object, then the eye movements will be quite small. On the other hand, if an observer is performing a task such as scanning a prairie for trees, then the eye movements will

be quite large. As a representative task between these two extremes, consider a search task where the display size is similar to the image size analyzed here, which had a width of approximately 17 deg (see Section 2). Najemnik and Geisler (2005) measured eye movements while observers searched for Gabor targets that were randomly located in backgrounds of $1/f$ noise with a diameter of 15 deg. They varied the target and noise contrast parametrically and found that the average distance between fixations was more than 4 deg. We note that this estimate of the average distance between fixations in natural tasks is likely to be conservative; for example, under natural viewing conditions Becker (1975) (as described in Rodieck, 1998) finds that the average distance between successive fixations is greater than 7 deg (see also Land & Hayhoe, 2001). These results, in combination with the present study, strongly suggest that during many natural tasks there will be little correlation between the contrasts that fall within a given receptive field on successive fixations; furthermore, the jumps in contrast that occur from one fixation to the next will often be quite large (see, for example, Fig. 1).³

These facts have strong implications for the dynamics of contrast gain control mechanisms. Consider, for example, the contrast gain control mechanisms in the primary visual cortex (V1). It is well known that the sensitivity of V1 neurons decreases following the presentation of high contrast stimuli (Albrecht, Farrar, & Hamilton, 1984; Bonds, 1991; Ohzawa, Sclar, & Freeman, 1985; for a review see Albrecht, Geisler, Frazor, & Crane, 2002; Albrecht, Geisler, & Crane, 2003). The time course for the build up and decay of these sensitivity changes is on the order of seconds, and thus the underlying adaptation mechanisms are too sluggish to adjust for many of the rapid large changes in contrast that occur due to eye movements.

³ One possible objection to this conclusion is that, during natural search tasks, observers might select fixation locations on the basis of local contrast. However, what evidence is available suggests that fixated locations are only slightly higher in contrast (on average) than randomly selected locations (Reinagel & Zador, 1999). More importantly, most receptive fields are not centered at the fixation location and hence will receive a random sample of local contrast.

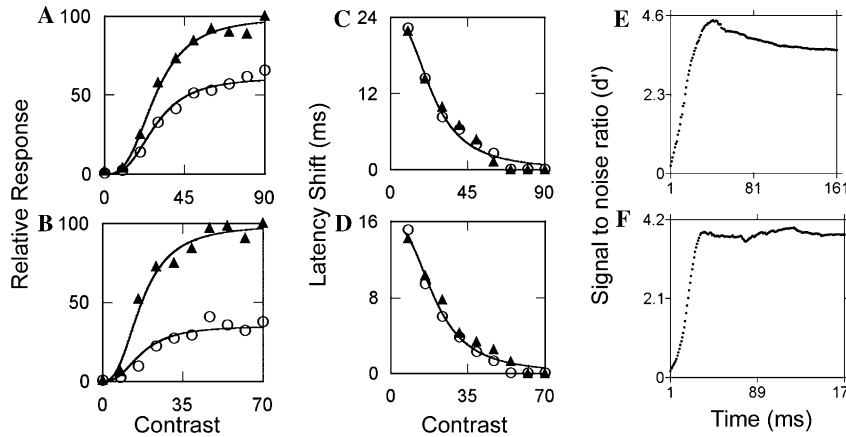


Fig. 8. Evidence for rapid contrast gain control in primary visual cortex. (A and B) Typical contrast response functions for an optimal and a non-optimal spatial phase during the first 20 ms of response to a 200 ms sine wave grating in monkey (A) and cat (B). (C and D) Typical contrast latency functions for an optimal and a non-optimal spatial phase for the first 20 ms of response to a 200 ms sine wave grating in monkey (C) and cat (D). (E and F) Signal-to-noise ratio (d') for pattern detection as a function of integration time starting at the onset of the response to a high contrast 200 ms sine wave grating of optimal spatial frequency, orientation and phase, in monkey (E) and cat (F). (Taken from Albrecht et al., 2002 and Frazor et al., 2004.)

There is, however, evidence for a much faster form of contrast gain control in the primary visual cortex (Albrecht & Geisler, 1991; Albrecht et al., 2002; Albrecht & Hamilton, 1982; Carandini & Heeger, 1994; Carandini, Heeger, & Movshon, 1997; Frazor, Albrecht, Geisler, & Crane, 2004; Geisler & Albrecht, 1992, 1997; Heeger, 1991, 1992; Sclar & Freeman, 1982). This rapid form of gain control, which is often referred to as “contrast normalization”, appears to play a fundamental role in cortical processing: (1) it creates invariant tuning characteristics of cortical neurons along various stimulus dimensions including orientation, spatial frequency, spatial phase, temporal frequency, and direction of motion, even at contrasts producing response saturation (see references above), (2) as a consequence, it also creates invariant population responses as a function of contrast, (3) it causes the relationship between the stimulus and response to become more constrained (unique) at high response rates (Albrecht & Geisler, 1991; Geisler & Albrecht, 1995), and (4) it increases the statistical independence of neural responses in a given local region thereby increasing the efficiency of the neural representation (Wainwright, Schwartz, & Simoncelli, 2002).

Two recent studies (Albrecht et al., 2002; Frazor et al., 2004) provide strong evidence that contrast normalization has very rapid temporal dynamics. For example, Fig. 8A shows the contrast response functions (response as a function of contrast) measured during the first 20 ms of the response of a typical neuron in monkey V1, for sine wave grating stimuli that have an optimal spatial phase (solid symbols) and a non-optimal spatial phase (open symbols). Fig. 8B shows the same measurements for a typical neuron in cat area 17. As can be seen, response saturation is reached at the same contrast for optimal and non-optimal stimuli, thus preserving selectivity to phase even in the saturated response range. The curves through the data have exactly the same shape (they differ by a scale factor), implying that selectivity (the phase tuning) is approximately con-

stant (invariant) independent of contrast. Figs. 8C and D show the contrast latency functions (the change in the latency to the peak response as a function of contrast) for the same cells shown in Figs. 8A and B. As can be seen, response latency declines with contrast in the same way for optimal and non-optimal stimuli. These two non-linear effects, which hold for all cells measured in the study, must be due to contrast-dependent mechanisms, because the response saturation and the latency changes are determined solely by the contrast of the stimuli, and not by the response rate of the cell (for more details see Albrecht et al., 2003, and the other references listed above). The fact that this full-blown pattern of contrast gain control effects occurs within tens of milliseconds of response onset implies that contrast normalization is very rapid.

It appears then, that the temporal dynamics of at least one component of contrast gain control are well matched to the temporal dynamics of contrast on the retina implied by the statistics of natural images and normal eye movement patterns. Rapid contrast gain control may be the result of feed-forward and/or feedback neural mechanisms in the retina, LGN, and cortex (Albrecht & Geisler, 1991). Regardless of the locus, our findings suggest that the rapid dynamics of these mechanisms may be the consequence of an evolutionary pressure created by the statistics of contrast in the natural environment, in conjunction with the eye movement requirements of foveated visual systems.

Further evidence for a match between the temporal dynamics of primary visual cortex neurons and the statistics of eye movements is shown in Figs. 8E and F, which illustrate how detection performance (d') grows as spikes are integrated during a 200 ms presentation of an optimal sine wave grating.⁴ As the integration interval increases,

⁴ A d' of 1.0 corresponds to 75% correct in the two-interval two-alternative forced choice detection task.

d' increases rapidly and then reaches a plateau well before the end of the 200 ms presentation. Over the population of neurons measured in Frazor et al. (2004), the average time to reach 90% of the maximum d' is approximately 50 ms in monkey and approximately 100 ms in cat. Thus, for stationary stimuli it appears that most of the spike rate information is transmitted by primary visual cortex neurons within a time interval that is well within the duration of a typical fixation during visual search. This time course would seem to be well matched to the eye movement system, under the assumption that recognition processes and eye movement planning/programming must occur before the end of the fixation. It is important to note that this rapid information saturation in the step response is not a reflection of the time constant of contrast normalization (which is considerably faster); rather, it is a reflection of the transient shape of the step response, which may be due to a combination of linear temporal filtering and rapid highly local light adaptation.

This rapid information saturation observed in primary visual cortex neurons is consistent with psychophysical studies showing that search (character detection) performance in sequences of random character displays is unaffected by decreasing the inter stimulus interval to less than half the duration of a single fixation (e.g., Sperling, Budiansky, Spivak, & Johnson, 1970), and with psychophysical studies showing rapid dynamics in contrast masking (e.g., Wilson & Kim, 1998).

4.2. Eye movements and rapid luminance gain control

The present results also have implications for the dynamics of luminance gain control/adaptation. Although the decorrelation distance for local luminance is larger than it is for local contrast (~ 4 deg vs. ~ 2.5 deg), there are still many fixation eye movements greater than 4 deg. For these fixations there will be little correlation in the local luminance before and after fixation. Therefore, given that there are substantial variations in local luminance within natural

images, it would presumably be useful for the visual system to have local luminance adaptation mechanisms that build up and decay rapidly enough to come to equilibrium in parallel with contrast normalization. There is psychophysical evidence for rapid multiplicative and subtractive gain control at photopic light levels (Geisler, 1981, 1983; Hayhoe, Benimoff, & Hood, 1987; Hayhoe, Levin, & Koshel, 1992; for reviews see Hood, 1998; Makous, 1997). However, there are few relevant neurophysiological studies in primates; Yeh, Lee, and Kremers (1996) report a rapid component of light adaptation in M and P cells, but the stimuli did not allow precise measurement of time constants.

4.3. Slow contrast gain control

If local contrast tends to be relatively uncorrelated across fixations within a scene, then how can we make sense of slow contrast adaptation? One possible explanation is that slow contrast adaptation adjusts for changes in contrast statistics that occur when the organism moves from one environment to another. However, the measurements shown in Fig. 9A suggest that this is probably not the case. This figure plots the mean local contrast and the standard deviation of local contrast for each of the 300 rural images. The fact that all but a handful of images cluster together implies that the distribution of contrast is similar from one image to the next. Thus, there would seem to be relatively little change in the distribution of contrast from one environment to the next, and hence relatively little need for slow contrast adaptation mechanisms.

Another possible explanation for slow contrast adaptation is that it provides useful sensitivity adjustments under fixation conditions that confine receptive fields to certain image regions for several seconds. For example, the average contrast of sky is nearly a log unit lower than the average contrast in other kinds of image region (see Fig. 4) and a receptive field may sometimes remain in a sky region for many seconds. The plausibility of this explanation depends

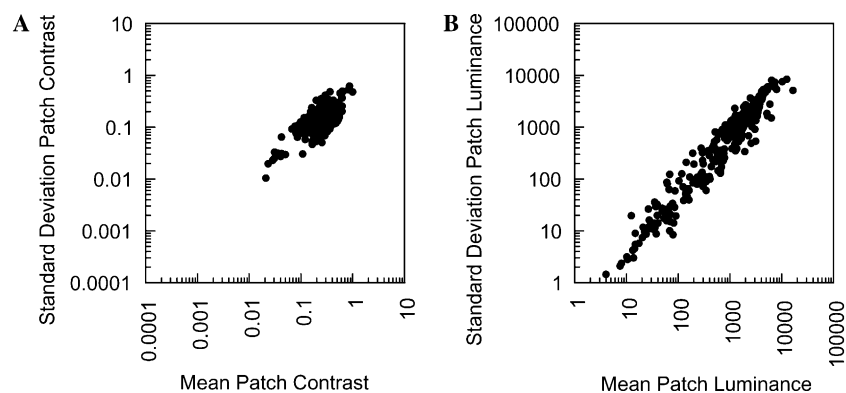


Fig. 9. Distributions of local contrast and local luminance for each of the 300 rural images. Each data point represents a distribution (a mean and a standard deviation) for a particular rural image. (A) Scatter plot of the mean local rms contrast and the standard deviation of local rms contrast. (B) Scatter plot of the mean local luminance and the standard deviation of local luminance. Note that all axes were set to be equal in numbers of log units. The analysis patch diameter was 0.54 deg.

upon the nature of the mechanisms underlying the slow contrast adaptation. In primary visual cortex, prolonged stimulation to high contrast patterns typically produces two changes in the contrast response functions of single neurons: an increase in the half-saturation contrast and a reduction in the maximum response rate (Albrecht et al., 1984). If these adaptation effects are primarily dependent upon the response rate (depolarization) of the neuron, then they are unlikely to be of much value in adjusting for different kinds of image regions. The high degree of selectivity of V1 neurons to spatial frequency, orientation and phase implies that the vast majority of fixations will produce little or no response from any given neuron, even if the eye movements are quite small and the receptive field is falling within a high contrast region of the image (Geisler & Albrecht, 1997). In other words, the average maintained activity of a cortical neuron will be quite small, even in a high contrast region.

On the other hand, if slow adaptation is primarily the result of network mechanisms that are relatively broad in their spatial frequency, orientation and phase tuning, then slow adaptation could be of benefit in adjusting to statistical properties of particular regions of the visual scene (Wainwright et al., 2002). There is some evidence that slow contrast adaptation involves network mechanisms (Albrecht et al., 1984; Movshon & Lennie, 1979), but the relative contribution of network and rate-dependent mechanisms is unknown. In any event, this is a potentially important role for slow adaptation.

Another possibility is that psychophysically measured contrast adaptation is a bi-product of high-level spatial pattern adaptation mechanisms that might serve various roles in object perception and recognition (Webster, 2004).

Finally, we note that it is possible that sluggish contrast adaptation is not an adjustment for statistical properties of natural scenes, but instead is a special case of a general mechanism whose purpose is to conserve metabolic energy by keeping the maintained activity of cortical neurons near zero. This is not implausible given that there appear to be severe limits on the average spike rate that can be supported by the metabolic systems in the brain (Attwell & Laughlin, 2001; Lennie, 2003).

4.4. Slow luminance gain control

There is less uncertainty about the role of slow luminance gain control. There is a long history of research measuring and characterizing the slow components of luminance adaptation that build up and decay on the order seconds or minutes (for reviews see Hood, 1998; Hood & Finkelstein, 1986; Shapley & Enroth-Cugell, 1984; Walraven, Enroth-Cugell, Hood, MacLeod, & Schnapf, 1990). These slow adaptation mechanisms undoubtedly reflect the environmental fact that changes in ambient illumination tend to occur slowly (dawn and dusk) or infrequently (e.g., moving out from under a forest canopy). The difference between luminance and contrast in this regard is dem-

onstrated in Figs. 9A and B. Fig. 9B plots the mean local luminance and the standard deviation of local luminance for each of the 300 rural images. Here we see that the natural images do not cluster together, but are spread across a wide range (note that the axes in Fig. 9A and B are equal in numbers of log units), and hence there would seem to be considerable value in having slow luminance adaptation mechanisms. Based on Fig. 9, it would seem that the statistics of natural images provide considerable evolutionary pressure for slow luminance adaptation, but perhaps less pressure for slow contrast adaptation.

4.5. Contrast response functions

Laughlin (1981) reported a close relationship between the distribution of contrast values in natural images and the shape of the contrast response function of the large monopolar cells (LP cells) in the blowfly. He found that the contrast response function effectively performs a form of “histogram equalization”—each possible response rate of an LP cell occurs equally often (on average) in the natural environment. This is an efficient way to use the full response range of a neuron, and is consistent with the notion that a goal in the early visual system is to efficiently encode visual information (e.g., Barlow, 1961). Tadmor and Tolhurst (2000) reported a similar close relationship between the distributions of equivalent Michelson contrast in natural images and the shapes of the contrast response functions of neurons in the lateral geniculate nucleus (LGN) of cats, and in magnocellular layers of the LGN in primates, although this simple relationship does not appear to hold as well for parvocellular neurons in the LGN or for neurons in primary visual cortex (Brady & Field, 2000; Clatworthy et al., 2003; Tadmor & Tolhurst, 2000).

Comparisons of the contrast in natural scenes and the contrast response function of neural populations depend to some extent on the measure of contrast. Most previous studies have used an equivalent contrast measure that is designed to represent the contrast in natural images that would activate (excite) a typical neuron; that is, they attempt to measure the fraction of local contrast that is matched to the typical receptive field. However, in the primary visual cortex (and perhaps in the retina as well) the shape of the contrast response function appears to be determined by the inhibitory (normalization) effects of contrast gain control mechanisms. For example, response saturation occurs at the same contrast for optimal and non-optimal stimuli, implying that a wide range of spatial frequencies, orientations and phases controls the half-saturation contrast of cortical neurons (see Fig. 8 and associated references). This suggests that RMS contrast, or some other broad band measure of local contrast, might also be an appropriate measure for comparison with neural contrast response functions. Fig. 2A shows that the average RMS contrast in rural images is in the range 0.2–0.34 (depending on analysis patch size). Fig. 2B shows that

the band-limited RMS contrast is in the range of 0.15–0.18 for rural images. (Recall that this band-limited RMS contrast is probably at or below the lower limit of plausible equivalent contrasts for the contrast normalization mechanisms evident in primary visual cortex.) If the contrast response functions (more specifically the contrast normalization mechanism) of cortical neurons were well matched to the contrasts in natural scenes we might expect the half-saturation contrast (c_{50}) to match the median contrast (Brady & Field, 2000; Clatworthy et al., 2003). The median half-saturation RMS contrast of neurons in the primary visual cortex of monkey is in the range of 0.18–0.24 (Albrecht & Hamilton, 1982; Geisler & Albrecht, 1997; Sclar, Maunsell, & Lennie, 1990), which would seem to be in reasonable agreement with the contrasts in natural scenes.

The median half-saturation RMS contrast for neurons in cat primary visual cortex is approximately half that in the monkey (Albrecht & Hamilton, 1982; Clatworthy et al., 2003; Geisler & Albrecht, 1997). However, optics and retinal center mechanisms create an effective cutoff frequency of 6–8 cpd (Blake, 1988), thereby reducing the effective RMS contrast of visual images. Blurring with a Gaussian kernel that cuts off at 8 cpd (assuming a peak contrast sensitivity of 100) reduces the effective RMS contrast by a factor of approximately 2. Thus, it is possible that the lower half-saturation contrasts of cat cortical neurons are matched to the effectively reduced image contrasts.

We note, however, that the rough match between the median half-saturation contrasts of cortical neurons and the contrast in natural images may have little to do with histogram equalization (in the contrast normalization mechanism). For example, the match could be the result of evolutionary pressure to maximize the signal-to-noise ratio of single neuron responses; that is, the match could reflect a compromise between the competing sub-goals of increasing gain to produce large responses to natural contrasts and decreasing the gain to avoid amplifying neural noise.

4.6. Independence of local luminance and contrast

Fig. 6 shows that there is relatively little correlation between local luminance and local contrast. Interestingly, there is a large negative correlation (approximately -0.8) in images that have the same amplitude spectra as natural images, but randomized phase spectra. It has been suggested that the amplitude spectra of natural images may be sufficient to understand retinal function: “The retina, being the first major stage in visual processing, is not expected to have knowledge beyond the simplest aspects of natural scenes and hence for understanding the retina the power spectrum (of the image) may be sufficient” (Atick & Redlich, 1992). The large difference in the luminance vs. contrast correlation between phase-scrambled and unscrambled images suggests that this is not the case.

If there were a large negative correlation between local luminance and contrast, then there would presumably have

been substantial evolutionary pressure to exploit the redundancy inherent in the negative correlation. For example, it would be possible to improve the contrast resolution of neurons (i.e., neurons in the early visual system could have steeper contrast response functions). Specifically, because of the strong correlation, the local luminance could be used to shift a very steep contrast response function to the proper location on the contrast axis. However, the independence of luminance and contrast eliminates this possibility; neurons are required to have less steep contrast response functions.

What is the reason for the large negative correlation between local luminance and contrast in phase-scrambled natural images? Consider the modulations in pixel luminance above and below the mean luminance of the whole image. By the principle of symmetry, randomizing the phase spectrum guarantees that the pixel luminance variations above and below the mean are on average statistically identical (i.e., inverting the pixel contrasts about the mean cannot change the statistics of a phase-scrambled image). Thus, a local image patch with luminance below the image mean will contain the same pixel luminance standard deviation as one above the image mean. But, the local contrast is by definition the pixel luminance standard deviation divided by the local mean, and therefore the contrast of the patch with luminance below the image mean will (on average) have the greater contrast.

The above argument shows that the low correlation between local luminance and contrast in unscrambled natural images is not a trivial result. What factors are responsible for the low correlation? The most obvious factor is based upon the classic view that the luminance distribution at the eye is the product of a surface reflectance function and a more-or-less statistically independent illumination function. To explore this possibility, we modeled the retinal luminance distribution as a product of a $1/f$ noise reflectance function and a $1/f^n$ illumination function. The functions were independent random samples and we varied the exponent of the illumination function from 1 to 3 (as the exponent increases the random texture becomes smoother). For all exponents there remains a large negative correlation between local luminance and contrast. However, the correlation is in the range of -0.5 to -0.7 , less negative than for $1/f$ noise. Thus, although independence of the reflectance and illumination function must be a contributing factor, it does not appear to be sufficient.

Another plausible factor is an effect due to shadows and shading. Surfaces (e.g., the surfaces of leaves) that are in direct sunlight will have greater luminance on average than those in shaded regions of a scene. Further, because of the directionality of sunlight, the shadows created by (and hence next to) the intense surfaces will tend to form higher contrasts than the shadows created by the less intense surfaces in the shaded regions. Similarly, because of the directionality of sunlight, the shading patterns on surfaces in sunlight will tend to have higher contrast than those in the shaded regions of the scene. This may even hold under

overcast conditions because the illumination is still (on average) more diffuse in shaded regions. These effects will tend to create a positive correlation between local luminance and contrast, counteracting the negative correlation expected for pixel luminance distributions that are symmetric about the mean. There is some circumstantial evidence for this hypothesis in our statistics. The foliage regions (which contain many shadows) should display this effect more than other types of region, and indeed foliage is the only type of region where we observed a positive correlation (0.15).

Another possible factor is that reflectance functions of natural environments are not well modeled as $1/f$ noise. More work will be required to determine the relative importance of these different factors.

In order for there to be a low correlation there must be greater variation in pixel luminance for local image patches with luminance above the image mean than for those with luminance below the image mean (this must be due to factors such as those discussed above). It is well known that the distribution of pixel luminance in natural images is skewed toward higher luminance (e.g., see Brady & Field, 2000). This must create greater variation in pixel luminance for local image patches with luminance above the image mean than for those with luminance below the mean. Could this first-order statistical property of natural images, in conjunction with the $1/f$ second-order statistics, account for the low correlation between local luminance and contrast or are higher-order statistics critical? To test this hypothesis we measured the correlation between local luminance and contrast for noise that had only the first- and second-order statistics of natural images. This “first-order $1/f$ noise” was created as follows:

1. Generate standard $1/f$ noise by filtering white noise in the Fourier domain with the average amplitude spectrum of the natural images. This step gives the noise the second-order statistics of natural images. The cumulative pixel luminance distribution of this noise is a cumulative normal distribution function with a mean of u and a standard deviation of σ , $N(y; u, \sigma)$. We set $u = 0.5$ and $\sigma = 0.1$.
2. Measure the normalized cumulative pixel-luminance distribution, $H(x)$, of the natural images. This is done by forming the cumulative histogram of pixel luminance values from all images, after normalizing each pixel value by the average luminance of the image to which it belonged.
3. Find the monotonic point non-linearity $x = g(y)$ that maps $N(y; u, \sigma)$ onto $H(x)$. The function $g(\cdot)$ is given by: $g^{-1}(x) = N^{-1}(H(x); u, \sigma)$. To obtain a smooth, monotonic and invertible function, we fit the raw values $N^{-1}(H(x); u, \sigma)$ with a Naka-Rushton equation, $g^{-1}(x) = ax^n / (x^n + b^n)$, where a , b , and n are free parameters. This function provided a good fit, with generally better than 98% of variance accounted for (see caption of Fig. 10).

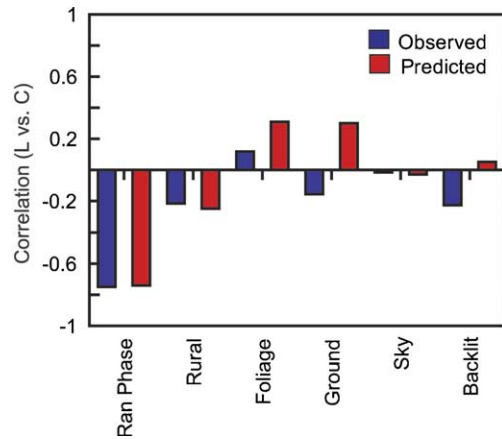


Fig. 10. Predicted effect of the first-order statistics (the pixel luminance distributions) on the correlation between local luminance and contrast. Natural images where the phase is randomized have a pixel luminance distribution that is symmetric about the mean (i.e., approximately Gaussian). In this case, the predicted correlation is highly negative (nearly -0.8). The skewed pixel luminance distributions of natural images shift the correlations near to or above zero. The parameters describing the pixel luminance distributions are as follows (the last number is the percentage of variance accounted for; see text for equations): ran phase ($a = 74$, $b = 1719$, $n = 0.66$, 99.1%), rural ($a = 1.28$, $b = 2.04$, $n = 0.41$, 99.4%), foliage ($a = 0.91$, $b = 0.65$, $n = 0.86$, 99.8%), ground ($a = 1.0$, $b = 0.86$, $n = 0.87$, 99.7%), sky ($a = 1.12$, $b = 1.16$, $n = 0.98$, 98.2%), and backlit ($a = 1.05$, $b = 0.82$, $n = 0.5$, 98.9%).

4. Apply the point non-linearity $g(y) = b\sqrt[n]{y/(a-y)}$ to the standard $1/f$ noise generated in step 1, where y is the pixel luminance. This point non-linearity gives the noise the first-order statistics of natural images, but does not alter the shape of the amplitude spectrum.⁵

Fig. 10 shows the predicted and observed correlations between luminance and contrast. The predictions for the random phase case are trivial, but provide a check on the procedure for estimating the point non-linearity. The key observation is that the predicted correlations for all the image types are either near zero or greater than zero. This implies that the first-order statistics of natural images reflect the lack of a negative correlation between local luminance and contrast. Fig. 11A shows a sample of $1/f$ noise with Gaussian first-order statistics (i.e., standard $1/f$ noise) and Fig. 11B shows the same sample of $1/f$ noise with first-order statistics that match the average rural image.

The low correlation between local contrast and luminance suggests that the mechanisms of luminance and contrast gain control could operate relatively independently, in the sense that local contrast gain control mechanisms could

⁵ In general, point non-linearities do affect the Fourier amplitude spectrum. For example, in the extreme case of a high threshold, all but a few of the most intense pixels would be set to zero, and hence the spectrum would become flat. However, for the smooth point non-linearities estimated here there was no measurable change in the shape of the amplitude spectrum.




Fig. 11. Samples of $1/f$ noise. (A) Standard $1/f$ noise, which has a Gaussian pixel luminance distribution. In this noise there is a large negative correlation between local luminance and contrast. (B) First-order $1/f$ noise, which has the average pixel luminance distribution of natural (rural) images. In this noise there is a weak negative correlation between local luminance and contrast.

be efficient without taking into account the local luminance. Thus, for example, initial processing could normalize for the local mean luminance by forming a local relative luminance (contrast) signal, $c(x, y) = (|L(x, y) - \bar{L}|) / (\bar{L} + L_0)$, and subsequent mechanisms could normalize for local contrast by forming a local relative contrast signal, $r(x, y) = (c(x, y)) / (\bar{c} + c_0)$, using only the local relative luminance signal as input. This is a traditional view of luminance and contrast processing in the visual system. The statistical properties of natural images suggest that this traditional view, which is simple and parsimonious, could also be efficient. Recent measurements in the lateral geniculate nucleus of the cat suggest that luminance and contrast gain control indeed act independently in the retina (Mante et al., 2005).

5. Conclusion

We find that there are substantial variations in local luminance and contrast in natural images and that the correlation of luminance and contrast as a function of distance falls fairly rapidly with respect to the average distance between fixations. This implies that the dynamics of at least some components of luminance and contrast gain control need to be very rapid, and thus the statistical properties of natural images lend support to recent physiological evidence for rapid contrast and luminance gain control. In addition, we find that there is little correlation between local luminance and contrast (even though one would expect a large negative correlation for $1/f$ noise). This suggests that the rapid contrast gain control mechanisms should not depend on local luminance. Thus, the statistical properties of natural images also lend support to recent physiological evidence for independent luminance and contrast gain control mechanisms in the retina. Finally, we found that the first-order statistics and the amplitude spectra of natural images account for most of the statistical regularities we observed in local contrast and luminance.

Acknowledgment

Supported by NIH Grant EY11747.

References

- Albrecht, D. G., & Geisler, W. S. (1991). Motion selectivity and the contrast-response function of simple cells in the visual cortex. *Visual Neuroscience*, 7, 531–546.
- Albrecht, D. G., Geisler, W. S., Frazor, R. A., & Crane, A. M. (2002). Visual cortex neurons of monkeys and cats: Temporal dynamics of the contrast response function. *Journal of Neurophysiology*, 88, 888–913.
- Albrecht, D. G., Geisler, W. S., & Crane, A. M. (2003). Nonlinear properties of visual cortex neurons: Temporal dynamics, stimulus selectivity, neural performance. In L. Chalupa & J. Werner (Eds.), *The visual neurosciences* (pp. 747–764). Boston: MIT Press.
- Albrecht, D. G., & Hamilton, D. H. (1982). Striate cortex of monkey and cat: Contrast response function. *Journal of Neurophysiology*, 48(1), 217–237.
- Albrecht, D. G., Farrar, S. B., & Hamilton, D. B. (1984). Spatial contrast adaptation characteristics of neurones recorded in the cat's visual cortex. *Journal of Physiology*, 347, 713–739.
- Atick, J. J., & Redlich, A. N. (1992). What does the retina know about natural scenes? *Neural Computation*, 4, 196–210.
- Attwell, D., & Laughlin, S. B. (2001). An energy budget for signaling in the grey matter of the brain. *Journal of Cerebral Blood Flow and Metabolism*, 21, 1133–1145.
- Balboa, R. M., & Grzywacz, N. M. (2003). Power spectra and distribution of contrasts of natural images from different habitats. *Vision Research*, 43, 2527–2537.
- Barlow, H. B. (1961). Possible principles underlying the transformations of sensory messages. In W. A. Rosenblith (Ed.), *Sensory communication* (pp. 217–234). Cambridge, MA: MIT Press.
- Becker, W. (1975). Die regelung der konjugierten horizontalen augenbewegung des menschen: Untersuchungen am sakkadischen system und am fixationssystem der menschlichen okulomortorik. *Fachbereich Elektrotechnik*. Munich: Technische Universitat Munchen.
- Bell, A. J., & Sejnowski, T. J. (1997). The 'independent components' of natural scenes are edge filters. *Vision Research*, 37, 3327–3338.
- Bex, P. J., & Makous, W. (2002). Spatial frequency, phase, and the contrast of natural images. *Journal of the Optical Society of America*, 19, 1096–1106.
- Blake, R. (1988). Cat spatial vision. *TINS*, 11(2), 78–83.
- Bonds, A. B. (1991). Temporal dynamics of contrast gain in single cells of the cat striate cortex. *Visual Neuroscience*, 6, 239–255.

- Brady, N., & Field, D. J. (2000). Local contrast in natural images: Normalization and coding efficiency. *Perception*, *29*, 1041–1055.
- Burton, G. J., & Moorehead, I. R. (1987). Color and spatial structure in natural scenes. *Applied Optics*, *26*, 157–170.
- Carandini, M., & Heeger, D. J. (1994). Summation and division by neurons in primate visual cortex. *Science*, *264*, 1333–1336.
- Carandini, M., Heeger, D. J., & Movshon, J. A. (1997). Linearity and normalization in simple cells of the macaque primary visual cortex. *The Journal of Neuroscience*, *17*(21), 8621–8644.
- Clatworthy, P. L., Chirimuuta, M., Lauritzen, J. S., & Tolhurst, D. J. (2003). Coding of the contrasts in natural images by populations of neurons in primary visual cortex (V1). *Vision Research*, *43*, 1983–2001.
- De Valois, R. L., Albrecht, D. G., & Thorell, L. G. (1982). Spatial frequency selectivity of cells in macaque visual cortex. *Vision Research*, *22*, 545–559.
- De Valois, R. L., Yund, E. W., & Hepler, N. (1982). The orientation and direction selectivity of cells in macaque visual cortex. *Vision Research*, *22*, 531–544.
- Field, D. J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America A*, *4*, 2379–2394.
- Frazor, R. A., Albrecht, D. G., Geisler, W. S., & Crane, A. M. (2004). Visual cortex neurons of monkeys and cats: Temporal dynamics of the spatial frequency response function. *Journal of Neurophysiology*, *91*, 2607–2627.
- Geisler, W. S. (1981). Effects of bleaching and backgrounds on the flash response of the visual system. *Journal of Physiology, London*, *312*, 413–434.
- Geisler, W. S. (1983). Mechanisms of visual sensitivity: Backgrounds and early dark adaptation. *Vision Research*, *23*, 1423–1432.
- Geisler, W. S., & Albrecht, D. G. (1992). Cortical neurons: Isolation of contrast gain control. *Vision Research*, *32*, 1409–1410.
- Geisler, W. S., & Albrecht, D. G. (1995). Bayesian analysis of identification performance in monkey visual cortex: Nonlinear mechanisms and stimulus certainty. *Vision Research*, *35*, 2723–2730.
- Geisler, W. S., & Albrecht, D. G. (1997). Visual cortex neurons in monkeys and cats: Detection, discrimination, and identification. *Visual Neuroscience*, *14*, 897–919.
- Geisler, W. S., & Diehl, R. L. (2002). Bayesian natural selection and the evolution of perceptual systems. *Philosophical Transactions of the Royal Society of London B*, *357*, 419–448.
- Geisler, W. S., Perry, J. S., Super, B. J., & Gallogly, D. P. (2001). Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research*, *41*, 711–724.
- Hayhoe, M. M., Benimoff, N. E., & Hood, D. C. (1987). The time-course of multiplicative and subtractive adaptation process. *Vision Research*, *27*, 1981–1996.
- Hayhoe, M. M., Levin, M. E., & Koshel, R. J. (1992). Subtractive processes in light adaptation. *Vision Research*, *32*, 323–333.
- Heeger, D. J. (1991). Nonlinear model of neural responses in cat visual cortex. In M. S. Landy & J. A. Movshon (Eds.), *Computational models of visual perception* (pp. 119–133). Cambridge, MA: The MIT Press.
- Heeger, D. J. (1992). Normalization of cell responses in cat striate cortex. *Visual Neuroscience*, *9*, 191–197.
- Hood, D. C. (1998). Lower-level visual processing and models of light adaptation. *Annual Review of Psychology*, *49*, 503–535.
- Hood, D. C., & Finkelstein, M. A. (1986). Sensitivity to light. In K. R. Boff, L. Kaufman, & J. P. Thomas (Eds.), *Handbook of perception and human performance* (Vol. 1). New York: Wiley.
- Land, M., & Hayhoe, M. M. (2001). In what ways do eye movements contribute to everyday activities? *Vision Research*, *41*, 3559–3566.
- Laughlin, S. B. (1981). A simple coding procedure enhances a neuron's information capacity. *Zeitschrift fur Naturforschung C*, *36*, 910–912.
- Lennie, P. (2003). The cost of cortical computation. *Current Biology*, *13*, 493–497.
- Makous, W. (1997). Fourier models and the loci of adaptation. *Journal of the Optical Society of America*, *14*, 2323–2345.
- Mante, V., Bonin, V., Frazor, R. A., Geisler, W. S., & Carandini, M. (2005). Independence of gain control mechanisms in early visual system matches the statistics of natural images. *Nature Neuroscience*, *8*, 1690–1697.
- Movshon, J. A., & Lennie, P. (1979). Pattern-selective adaptation in visual cortical neurons. *Nature*, *278*, 850–852.
- Najemnik, J., & Geisler, W. S. (2005). Optimal eye movement strategies in visual search. *Nature*, *434*, 387–391.
- Ohzawa, I., Sclar, G., & Freeman, R. D. (1985). Contrast gain control in the cat's visual system. *Journal of Neurophysiology*, *54*(3), 651–667.
- Olshausen, B. A., & Field, D. J. (1997). Sparse coding with an overcomplete basis set: A strategy by V1? *Vision Research*, *37*(23), 3311–3325.
- Purves, D., & Lotto, R. B. (2003). *Why we see what we do: An empirical theory of vision*. Sunderland, MA: Sinauer Associates.
- Reinagel, P., & Zador, A. M. (1999). Natural scene statistics at the centre of gaze. *Computation in Neural Systems*, *10*, 1–10.
- Rodieck, R. W. (1998). *The first steps in seeing*. Sunderland, MA: Sinauer Associates.
- Ruderman, D. L. (1994). The statistics of natural images. *Computation in Neural Systems*, *5*, 517–548.
- Sclar, G., & Freeman, R. D. (1982). Orientation selectivity in the cat's striate cortex is invariant with contrast. *Experimental Brain Research*, *46*, 457–461.
- Sclar, G., Maunsell, J. H. R., & Lennie, P. (1990). Coding of image contrast in central visual pathways of macaque monkey. *Vision Research*, *30*, 1–10.
- Shapley, R., & Enroth-Cugell, C. (1984). Visual adaptation and retinal gain controls. *Progress in Retinal Research*, *3*, 263–346.
- Simoncelli, E. P., & Olshausen, B. A. (2001). Natural image statistics and neural representation. *Annual Review of Neuroscience*, *24*, 1193–1215.
- Sperling, G., Budiansky, J., Spivak, J. G., & Johnson, M. C. (1970). Extremely rapid visual search: The maximum rate of scanning letters for the presence of a numeral. *Science*, *174*, 307–311.
- Tadmor, Y., & Tolhurst, D. J. (2000). Calculating the contrasts that retinal ganglion cells and LGN neurones encounter in natural scenes. *Vision Research*, *40*, 3145–3157.
- Tolhurst, D. J., Tadmor, Y., & Chao, T. (1992). The amplitude spectra of natural images. *Ophthalmic and Physiological Optics*, *12*, 229–232.
- van Hateren, J. H. (1992). Real and optimal neural images in early vision. *Nature*, *360*, 68–70.
- van Hateren, J. H., & van der Schaaf, A. (1998). Independent component filters of natural images compared with simple cells in primary visual cortex. *Proceedings of the Royal Society of London B*, *265*, 359–366.
- Wainwright, M. J., Schwartz, O., & Simoncelli, E. P. (2002). Natural image statistics and divisive normalization: Modeling nonlinearities and adaptation in cortical neurons. In R. Rao, B. A. Olshausen, & M. S. Lewicki (Eds.), *Statistical theories of the brain* (pp. 203–222). Cambridge, MA: MIT Press.
- Walraven, J., Enroth-Cugell, C., Hood, D. C., MacLeod, D. I. A., & Schnapf, J. L. (1990). The control of visual sensitivity: Receptor and postreceptor processes. In L. Spillman & J. S. Werner (Eds.), *Visual perception: The neurophysiological foundations*. San Diego: Academic Press.
- Watson, A. B. (2000). Visual detection of spatial contrast patterns: Evaluation of five simple models. *Optics Express*, *6*, 12–33.
- Webster, M. A. (2004). Pattern-selective adaptation in color and form perception. In L. Chalupa & J. Werner (Eds.), *The visual neurosciences* (pp. 936–947). Cambridge, MA: MIT Press.
- Wilson, H. R., & Kim, J. (1998). Dynamics of a divisive gain control in human vision. *Vision Research*, *38*, 2735–2741.
- Yeh, T., Lee, B. B., & Kremers, J. (1996). The time course of adaptation in macaque retinal ganglion cells. *Vision Research*, *36*, 913–931.