

Color and luminance information in natural scenes

C. A. Párraga, G. Brelstaff⁺ T. Troscianko and I. Moorhead⁺⁺.

Perceptual Systems Research Centre, Dept. of Psychology, 8 Woodland Rd,
University of Bristol, BS8 1TN, UK

Abstract

The spatial filtering applied by the human visual system appears to be low-pass for chromatic stimuli and band-pass for luminance stimuli. Here we explore whether this observed difference in contrast sensitivity reflects a real difference in the components of chrominance and luminance in natural scenes. For this purpose a digital set of 29 hyper-spectral images of natural scenes has been acquired and its spatial frequency content analyzed in terms of chrominance and luminance defined according to existing models of the human cone responses and visual signal processing. The statistical $1/f$ amplitude spatial frequency distribution is confirmed for a variety of chromatic conditions across the visible spectrum. Our analysis suggests that natural scenes are relatively rich in high spatial-frequency chrominance information which does not appear to be transmitted by the human visual system. This result is unlikely to have arisen from errors in the original measurements. Several reasons may combine to explain a failure to transmit high spatial-frequency chrominance: (a) its minor importance for primate visual tasks, (b) its removal by filtering applied to compensate for chromatic aberration of the eye's optics, or (c) a biological bottleneck blocking its transmission. In addition, we graphically compare the ratios of luminance to chrominance measured by our hyperspectral camera and those measured psychophysically over an equivalent spatial frequency range.

1. Introduction

1.1 Context

A series of studies of how the human visual system (HVS) encodes and transmits retinal imagery has increasingly advocated a role for natural scenes as a formative force. In this approach, evolutionary pressure may provide the impetus for the HVS to adopt efficient visual representations in an information theoretic sense^{1,2,3}. By coding first the things that matter most (food, predators, mates, etc.) the fittest visual system ought to survive. In this scenario, the statistical spatio-chromatic properties of natural scenes help mold the spatio-chromatic characteristics of the early visual pathways so that they reduce redundancy and separate essential signal from noise. This approach has been successfully used by Atick³ to predict the physiologically measured responses of certain cells in lateral geniculate nucleus (LGN). The assumptions underlying this prediction include a constraint on the spatial frequency (SF) make-up of natural scenes over a range of chromatic conditions. This constraint, commonly termed the 1/f law, requires that the average Fourier amplitude of natural images falls off inversely with SF. Measurement has shown this to be roughly true for achromatic images^{4,5}. However, as yet, the evidence that 1/f applies in a variety of chromatic conditions has been, at best, inconclusive^{6,7,8}. We present here new measurements that further confirm the 1/f assumption over a wider variety of chromatic conditions.

Note, that a 1/f law does not imply that luminance contrast sensitivity (CSF) should simply match 1/f, because as is implied by Atick such a CSF would transmit a lot of noise. By explicitly modeling an early noise suppression stage he shows how a recognizable luminance CSF can be predicted.

Atick investigated only linear transformations as candidate coding schemes. Basically, this means that his LGN cells simply compute weighted sums or differences of their inputs - in our terms these may represent luminance or chrominance (e.g. L+M or L-M, respectively where L derives from the red cones and M from the green). As recent developments in artificial neural networks⁹ and single cell recording¹⁰ techniques are revealing - single cells might be computing sophisticated non-linear functions of their inputs. With this in mind, we present results for a simple non-linear (divisive) transformation in addition to some previously established linear transformations. In the absence of other guidance, our non-linear transformation is motivated by the desire to represent colored objects

in a way that is independent of the shadowing that pervades natural scenes. In particular, we have investigated the definitions of luminance = $L+M$ and chrominance = $(L-M)/(L+M)$.

This formulation has the advantage of retaining the L-M subtractive element to chrominance whilst ensuring that a cast shadow boundary that typically attenuates L and M by similar amounts effects a luminance step but produces little or no change in chrominance. We precede details of our investigation by a review some evidence and models under-pinning it.

1.2 Evidence

Both physiological and psychological evidence suggests that the HVS analyzes images employing neural subsystems tuned to different attributes of the stimulus (color, texture, and movement) that vary in their degrees of spatial and temporal resolution. In particular, the psychophysically measured contrast sensitivity functions for visual stimuli modulated in luminance and chrominance show different profiles.

Chrominance modulation (of red/green isoluminant color design) is not seen at high SF above about 12 cycles/deg, whereas luminance modulation has been reported to be visible up to¹¹ and beyond¹² 30 cycles/deg. Furthermore, luminance sensitivity appears to peak around 2 cycles/deg and drops off at lower SFs (thus achieving band-pass filtering), whereas sensitivity to isoluminant color is highest at those low SFs (and is low-pass filtered). We restrict our consideration to red-green opponent definitions and do not address the blue-yellow system with its complication of sparser retinal coverage, nor peripheral vision.

A qualitative division into band-pass luminance and low-pass chrominance transmission prevails also in physiological measurements^{13,14} of the single cells in parvocellular layers of LGN in macaque (with much similarity to the HVS). The responses measured to modulations in luminance or chrominance show that the majority of cells (coined "Type-I R-G") exhibit center-surround receptive fields that simultaneously encode band-pass luminance and low-pass chrominance signals. It is not clear precisely how the signals from these cells might propagate to generate psychophysically measured sensation, nor whether other visual pathways might be involved to significant extents¹⁵. However, it is interesting to note that these "Type-I R-G" cells, considered leading candidates for the early transmission of chromatic and luminance information, do seem to share similar SF characteristics as the final output (single cells however have narrower bandwidths).

1.3 Multiplex Model

A multiplex model^{16,17} (analogous to that of AM radio transmission) may be used to picture the SF banding in the signal transmitted from a single "Type-I R-G" cells. Demultiplexing this signal after transmission to striate cortex might provide components of luminance and chrominance from which the psychophysical CSFs could be derived. Summing the signals transmitted by opposite pairs of co-located opponent cells (i.e. a red-center/green-surround and a green-center/red-surround) could fully recover both components¹⁸. However, this would require two neurons transmitting down the optic nerve instead of one and so some alternatives have been considered. Demultiplexing the signal from a single neuron into two discrete SF bands can only partially recover the original components: as their SF domains do overlap the reconstructed chrominance would be contaminated by any low SF luminance modulations, and reconstructed luminance would contain any high SF chrominance modulation. In either case the degree of cross-talk (contamination) depends on the SF content of the signals actually carried.

A large amount of cross-talk might normally be expected in low SF chrominance because low SF luminance stimuli are abundant in natural scenes. A single cell operating noise suppression according to Atick's model should theoretically transmit less low SF luminance and thus cause less cross-talk.

Conversely, the amount of cross-talk in the luminance signal will depend on the quantity of high SF chromatic signal arriving at the "Type-I R-G" cell. This in turn will depend on the relative abundance of high SF chrominance in natural scenes - as well as the optical and retinal factors.

As a provisional quantification of the cross-talk problem that might be faced by the HVS we describe below an analysis of our hyper-spectral measurements in terms of the ratio of the Fourier amplitudes of luminance to chrominance - over SF, for various definitions.

Bear in mind that - as we have not explicitly modeled noise suppression - the ratios may be biased and difficult to interpret. Further, note that our measuring device precludes analysis of ratios nominally beyond 9 cycles/deg.

1.4 Definitions

Previous studies have motivated various definitions of luminance and chrominance in terms of the Smith and Pokorny cone responses^{19,20}. We adopt three linear definitions from the literature and add

our own non-linear definition. In our case only the L and M cone responses are involved and our computations assume their spectral curves normalized so to peak at unity. The definitions are:

a) *Simple* definition:

$$lum_s = L + M$$

$$chrom_s = L - M.$$

b) *Inglis and Tsou²¹* definition - motivated by flicker and acuity criteria:

$$lum_{IT} = 1.02.L + M$$

$$chrom_{IT} = 0.41.L - M.$$

c) *Buchsbaum and Gottschalk²²* definition - derived from a principal component analysis of the cone responses:

$$lum_{BG} = 0.887.L + 0.461.M$$

$$chrom_{BG} = 0.46.L - 0.88.M$$

d) *Shadow-removing* definition:

$$lum_{sr} = L + M$$

$$chrom_{sr} = \frac{L - M}{L + M} = \frac{lum_s}{chrom_s}$$

In practice, our hyperspectral measurements for each scene allow values of L and M to be computed from 31 narrow-band radiance images. The data-set of all 29 scenes is being made available to other researchers at an ftp site posted at <http://www.crs4.it/~gjb/ftpJOSA.html>.

2. Methods

2.1 The hyper-spectral camera.

The hyper-spectral camera, further described elsewhere²³, measures radiance as a function of wavelength for each image pixel. It consists of an electro-optic mechanism built around a Pasecon integrating camera tube, a camera control unit (CCU), a carousel slide changing filter mechanism, a portable PC, a video monitor and a battery power supply. The Pasecon has good linearity throughout the visible spectrum and is able to integrate signal over a computer selected interval - thus permitting a reasonable signal-to-noise ratio even at the low light levels incurred by using

narrow band filters. A light-tight, slide changing, filter mechanism allows the camera to sequentially acquire images through the set of 31 optical interference filters that span the range of 400 to 700 nm and with nominal 10 nm spacing. The entire system is controlled by the PC and mounted on a trolley along with power supply kit. A manual, fixed focal length lens (Fuji CF25B, f/1.4, 25 mm) with a field angle of 28.71h deg x 21.73v deg is employed. An MS-DOS C program sends a value for the integration time to the CCU which digitizes each image and transfers it to the PC frame card. The whole system acquires a sequence of 31 chromatically narrow-band filtered 8-bit, 256 x 256 pixel images in less than 5 minutes.

Concerning the MTF of the system, there are two possible issues: (a) the possibility of aliasing of spatial frequencies above the Nyquist limit (9.0 cycles/degree) and (b) effects of wavelength (e.g. due to chromatic aberration of the camera lens). For (a) tests were performed on synthetic scenes that showed no intrusion of aliasing of higher SF's in the Fourier spectra of the images in the range used in this study; for (b) there was no effect of wavelength in the relevant range. There was a small reduction in modulation for high SF's below 450 nm but these wavelength do not provide a significant input to the L and M cone primaries.

Individual interference filters differ in their degree of flatness and thus may systematically displace their images slightly on the camera target. These shifts (no more than 2 pixels) once calibrated are routinely corrected on the frame card. Each image corresponds only to the central part of the visual field supplied by the lens. The field angle is equal to 14.35 x 14.35 deg and each pixel subtends an angle of approximately 0.056 x 0.056 deg (3.36 arcmin) - which is not quite the same order of the size as a foveal cone (around 0.5 arcmin). Once recorded, the image-set is archived on a UNIX workstation where measured radiance is computed for each interference filter using a different conversion function. Each function depends the pixel gray-level, pixel position, filter transmission, the lens f-stop and the tube integration time. It had been pre-calibrated in the dark room. The calibration process involved a Kodak standard gray card, a TopCon SR1 spectrometer, and a target collage of colored papers under illumination by a steady current light source. The relative error of the camera compared to the SR1 was less than 5% in the range 400-570 nm, and 10% in the range 580-700 nm. The worst case (20% relative error) was found in the extreme red²³ (580-700 nm).

Given the long acquisition time, scenes were recorded under stable weather and lighting conditions. Outdoor scenes were recorded around noon to minimize changes of illumination, and sometimes it was necessary to wait until the sky was completely clear or completely overcast before starting. Unfortunately this greatly restricts the quantity of cast shadows in our data. Some scenes were

recorded inside the glass houses of the Bristol Botanical Gardens to avoid small movements of branches and leaves occasioned by the wind, and out of doors some objects at short-distance were avoided.

2.2 Characteristics of the data-set

Although there is no formal agreement about what is considered a representative sampling of the visual environment, we hoped to ensure that some of the most common natural objects are represented in our data-set. These include plants with different shapes, textures and colors, flowers (often bright red) , trunks, branches, grass, green and yellow leaves, trees, bushes, rocks and sky. Figure 1 shows four achromatic representations of images used for this work. Many scenes are of the English countryside and gardens - in which buildings or other human artifacts are avoided. Others are of natural objects arranged in the laboratory to look fairly naturalistic. Our general aim was to acquire a set of scenes which could conceivably be representative of the environment in which primate vision evolved. Of course, we cannot be sure that we have achieved this aim. However, the inclusion of vegetation seen from different distances goes some way towards satisfying this requirement.



Figure 1: Example of typical images from our data-set.

Several practical constraints also limited our choices. Bright sky and reflecting water result in large saturated areas of the image and are deliberately avoided. Only four of our scenes contain regions of sky. Drifting objects such as clouds and strong shadows are also avoided. All scenes were recorded between October 1993 and January 1994.

Table 1:

Range and central spatial frequency (in cycles/deg) of SF rings used to divide the Fourier space.

	Range	Central SF
band	cycles/deg	cycles/deg
1	0.03 - 0.07	0.05
2	0.07 - 0.14	0.10
3	0.14 - 0.28	0.21
4	0.28 - 0.55	0.41
5	0.55 - 1.10	0.83
6	1.10 - 2.20	1.65
7	2.20 - 4.40	3.30
8	4.40 - 8.80	6.60

2.4 Image processing.

A standard 2D FFT algorithm was available to derive for a given scene the Fourier amplitude spectrum of either each individual narrow-band radiance image, or of combinations of them - in the all forms of *lum* and *chrom* image as defined previously. From all such images it was then possible to measure the total amplitude contained within a set of chosen SF bands. These bands labeled 1 to 8 in Table 1, are centered at equal intervals on a logarithmic scale and do not overlap. Band 1 omits the zeroth SF so to exclude the direct influence of a scene's illumination level. To favor equality in the subsequent statistical representation of scenes all *lum* and *chrom* images are individually normalized to contain unit total Fourier amplitude. We do not mean, however, to imply that chrominance is apriori of equal importance as luminance: the normalization is merely a graphical convenience and does not unduly affect our conclusions.

3. Results

3.1 Measurements of 1/f.

The results summarized in Figure 2 demonstrate that the 1/f law approximately holds even for 10 nm chromatic bands. Expressing the Fourier amplitude distribution of spectral radiance as f^α allows a measurement of the slope α which should be -1 if 1/f holds. The value of α , plotted for each 10 nm interval, results from fitting a straight line to the SF distribution of the mean Fourier amplitude. In particular, the SF distribution is sampled at the center of each of the 8 bands in Table 1, and the mean amplitude is computed from all 29 scenes. In this case, each scene is given equal statistical weight by first normalizing its total SF amplitude across all wavelengths to be unity (again ignoring zeroth SF). As the plot shows slope α varies between about -1.08 and -1.16. This result is very similar to that previously reported by the authors for derived reflectance²³.

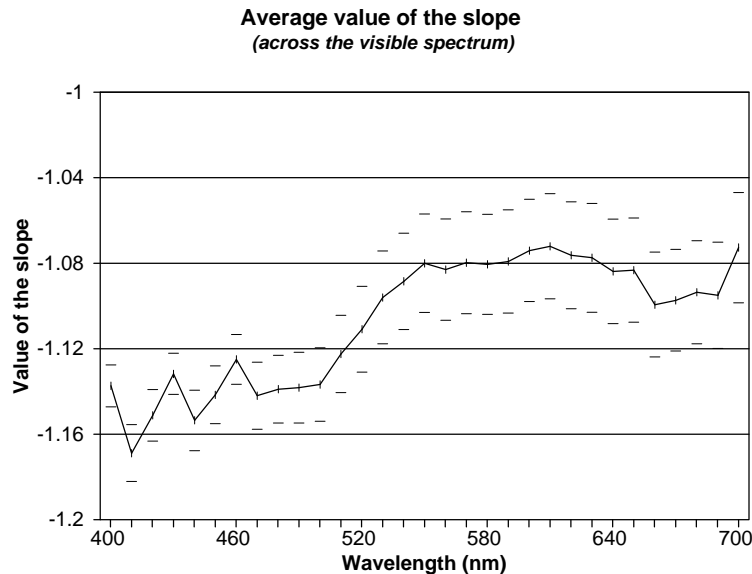


Figure 2: Mean value of the slope (α) across wavelength. Scenes were converted into radiance and their Fourier amplitude spectra normalized to 1. Standard error is shown in the plot.

As might be expected, computed luminance and chrominance also roughly obey 1/f. Figure 3 gives an idea of the variability encountered between scenes. It plots amplitude spectra in *lum* for four example images of the data-set. Figure 4 plots the mean amplitude over all 29 scenes in *lum* and *chrom* - each exhibiting 1/f like behavior. In this case, the simple definition of *chrom* is used.

Figure 5 plots the same for the shadow-removing definition - with similar results. Table 2 quantifies the mean values of the slope α for both definitions.

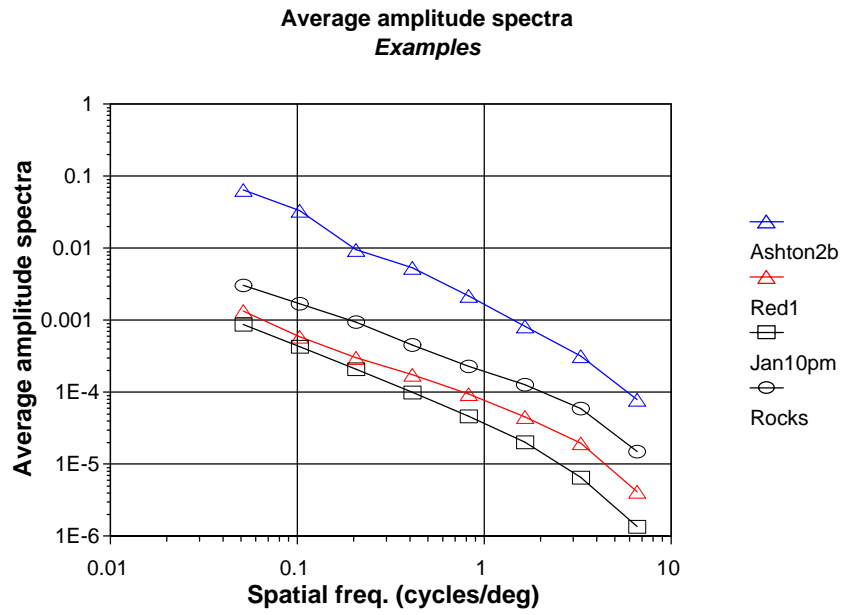


Figure 3: Amplitude Spectrum (*Lum*) for some scenes of our data-set (simple definition).




Figure 4: Amplitude Spectrum for *Lum* and *Chrom* for all the data-set (29 scenes) (simple definition).

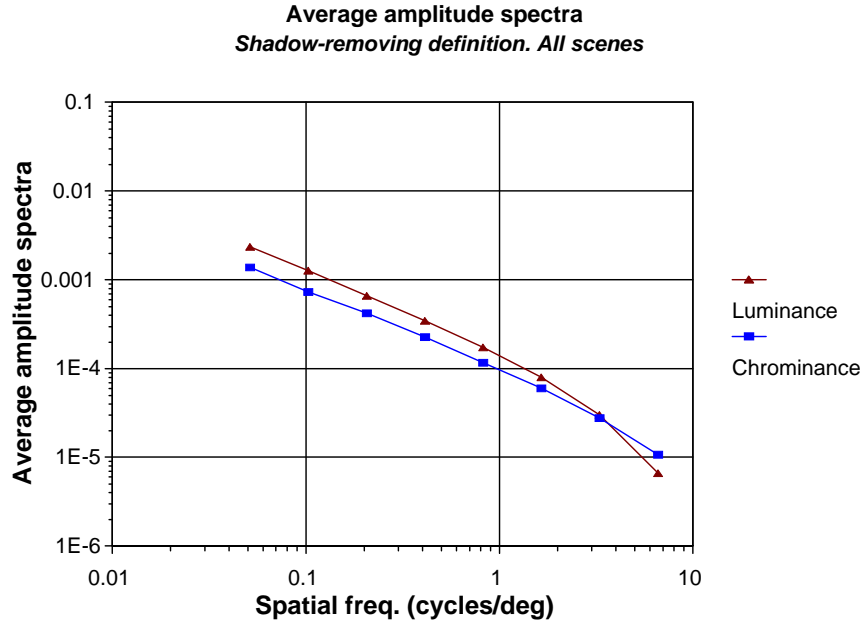


Figure 5: Amplitude Spectrum for *Lum* and *Chrom* for all the data-set (29 scenes) (shadow-removing).

Any discrepancy from $1/f$ occurs at high SF and could possibly be artifacts of our imaging system such as fine scale variations during acquisition perhaps due to wind effects. These effects would have been expect to be most profound in short distant outdoor scenes and thus motivated our analysis of bias described below.

Table 2:

Mean values of the slope (α) for the simple and shadow-removing definitions of *lum* and *chrom*.

Slope (α)	Simple definition	Shadow-removing definition
<i>Lum</i>	-1.11 ± 0.13	-1.11 ± 0.13
<i>Chrom</i>	-1.06 ± 0.11	-0.94 ± 0.12

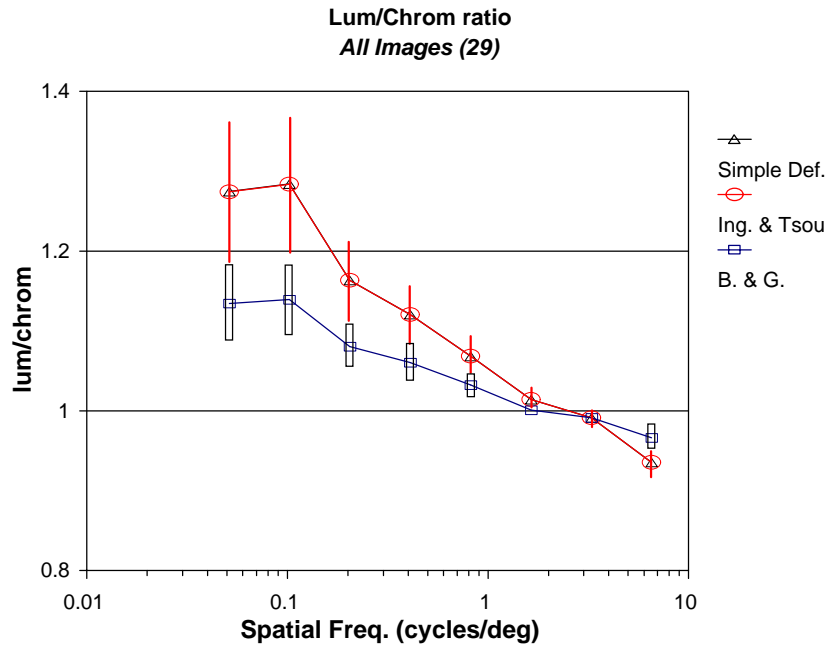


Figure 6: Mean ratio of the *lum* image to that of the *chrom* image amplitude for all the data-set. The three definitions are used. Standard error is shown on the plot.

3.2 Measurements of lum to chrom ratio

The ratio of the *lum* image amplitude to that of the *chrom* image amplitude is plotted for each spatial frequency band. Figure 6 shows the mean across all scenes of this ratio for all three linear definitions.

In the case of the *simple* and the *Ingling and Tsou* definition the results are coincident. For low SF there is relatively more luminance than chrominance. This is reversed only for the highest SF considered in the plot (the value of *lum/chrom* is lower than 1).

The *Buchsbaum and Gottschalk* definition, while not coincident, produces a similar trend. Figure 7 shows the results for the divisive *shadow-removing* definition. The basic features are the same. The only notable new feature is the change in the shape of the plot. One explanation for the similarity could be the avoidance of sharp shadows in our data-set.

3.3 Analysis of the bias of the data-set.

Figure 8 was produced analyzing two sets of scenes separately. The first set (called *short distance* viewing scenes set) consists of scenes in where the objects were up to 10 meters away from the camera. The second set (called *long-distance* viewing set) consists of scenes in where most of the

objects (e.g. trees) were farther than 10 meters. The plot suggests that the statistics of this *long distance* viewing set of scenes may be different from the others. This biasing might be due to the considerable proportion (50% approx.) of grass and sky in the scene.

Other classifications of the scenes of our data-set were attempted in order to check whether the quality of the pictures or the lighting (artificial or natural) might influence or results. None of these proved to be determinant over the appearance our plots. A “best quality” set of ten scenes was chosen in order to evaluate the presence of artifacts in our data-set. Figure 9 shows the results averaging over short distance viewing scenes. The data-set is divided into artificially lit (indoor) scenes, naturally lit (outdoor) scenes and “best quality” scenes. However, there seems to be little difference between these potentially different types of scene.

4. Discussion

Regarding the whole of the HVS simply as a black box, whose CSFs may be revealed by psychophysical experiment, leads to the inevitable view that the CSFs mismatch the SF content of natural scenes. Our results using four different measures of the luminance chrominance ratio leave little doubt of this - compare them with the corresponding ratios from Mullen in Figure 10. The “black box” must be acting to discard high SF chrominance and low SF luminance. As Atick illustrated the latter is probably achieved, in part, by quantum noise suppression acting at, or before, cells in the LGN (possibly by retinal adaptation to overall light level involving both cones and bipolar cells). Atick adopted a specific linear model of a single cell when making his predictions. We chose not to adopt that, or any other, explicit model of a single cell because: (a) it is becoming increasingly apparent that some single cells may be computing sophisticated non-linear functions of their inputs, (b) the logistics of cascading the outputs of such cells to generate CSF-like profiles is not yet well established. Further analyses (e.g. measuring correlation statistics) of data-sets such as ours might provide some insight into these issues. In this paper, the divisive definition of chrominance was unable to allow a fuller understanding of these issues. However, the execution of its designed purpose was probably frustrated by the lack of shadows in our data-set, for practical reasons.

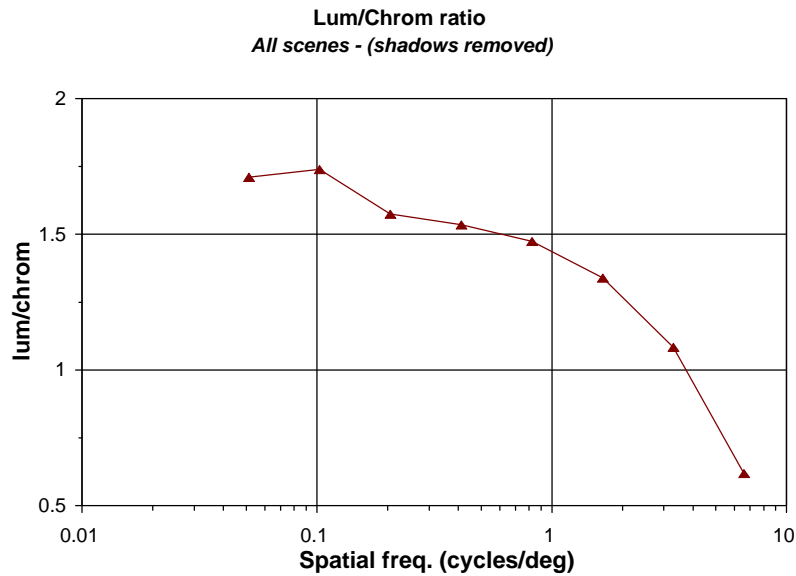


Figure 7: Mean ratio of the *lum* image amplitude to that of the *chrom* image amplitude using the *shadow-removing* definition of *lum* and *chrom*. This average is done across all the data-set.

The fact that the 1/f law seems to statistically hold across the many chromatic conditions that we present here (narrow and wide spectral bands) helps confirm a key assumption upon which Atick's single cell predictions depend.

Below we suggest some reasons to why - however, it may be achieved, the HVS appears to discard high SF chrominance.

(a) *Its irrelevance to primate vision tasks:* An early task of the visual system is the segregation of different parts of a visual scene in order to identify discrete objects. The spatial distribution of wavelength differences (color) would contribute to the differentiation of some objects from their surrounds and make them more identifiable despite the irrelevant contours produced by the shadowing²⁴. Consider the problem of finding a cherry in a cherry tree: the relevant information about the whole object is provided by global properties of the cherry. If we omit sensitivity to local changes in color (high SF) we are more likely to encode this global color which is all we need for the segregation task. Information about the local properties of the cherry comes from the luminance channel and it is not relevant for this early segregation task. Another consideration that our analysis of static scenes has not addressed is how the temporal statistics of natural scenes might also eclipse the relevance of color.

(b) *Its removal by the neural machinery that may compensate for the chromatic aberrations of the eye:* The HVS may be predisposed to discard high SF color information as it is differentially

blurred on to the retina by the optics of the eye. This blurring, termed chromatic aberrations²⁵, might not affect luminance to the same degree - although this remains unclear. For the record, however, our data-set is focused most sharply around the peaks of the L and M cones; and Mullen took considerable care to avoid chromatic aberration influencing her CSFs. So, although we have no concrete support: it is possible that the HVS discards high SF color because it forms part of the message that carries aberrated data about the real world²⁶.

(c)*Physiological constraints*: The optic nerve is a bottleneck in the information path from the retina to subsequent brain centers: there are insufficient neurons to transmit the retinal information. There is therefore strong pressure to maximize the amount of information transmitted per optic nerve neurons, which implies strong post-receptoral re-coding of the stimulus information designed to maximize the efficiency of this information transmission. As discussed in the Introduction this kind of data compression might have to be lossy and the HVS may have evolved to lose high SF color, and this might be achieved by some sort of multiplexing technique.

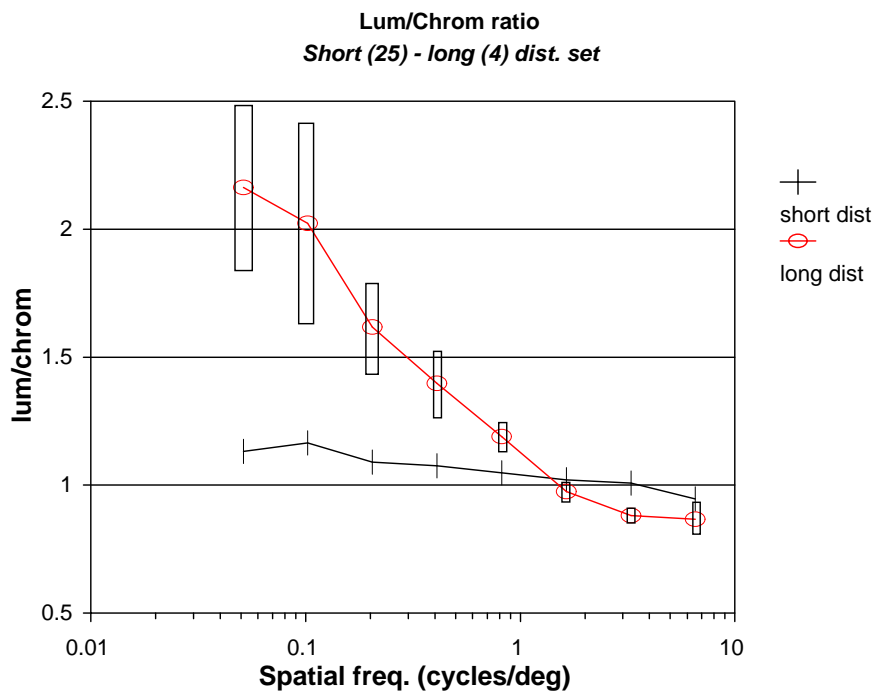


Figure 8: Mean ratio of the *lum* image amplitude to that of the *chrom* image amplitude for the *long-distance* viewing set (objects in the range 10m to 4 km and including grass and sky) and the *short-distance* viewing set (objects up to 10 m far away from the camera)

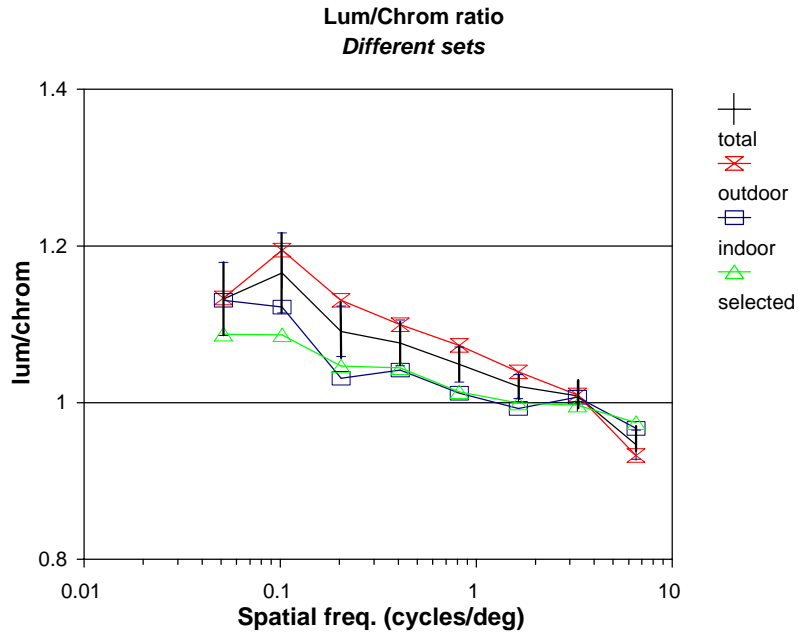


Figure 9: Mean ratio of the *lum* image amplitude to that of the *chrom* image amplitude for short distance scenes. The scenes are divided again into different groups (*outdoor*, *indoor* and *selected*) and plotted along with the average for all groups (*total*).

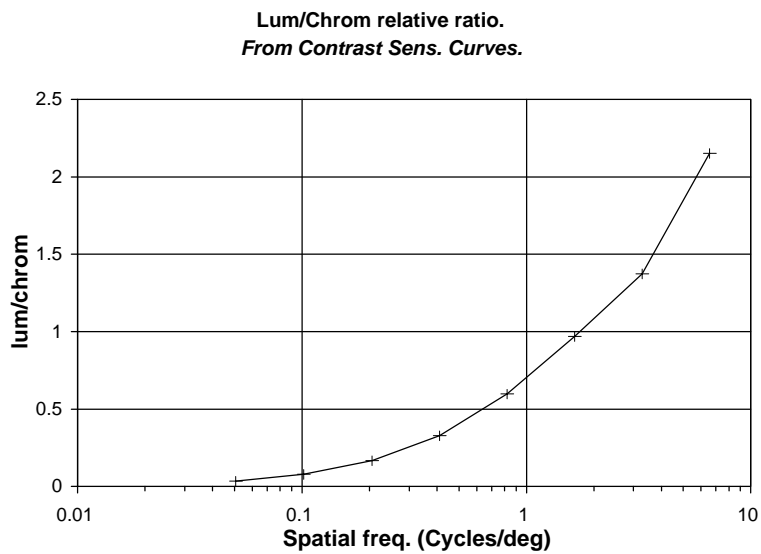


Figure 10: Plot of the *Luminance to Chrominance* ratio obtained from Mullen's¹¹ measured contrast sensitivity curves. The total area under each of the contrast sensitivity curves has been made equal to 1. To match these measurements, the amount of spatial energy for luminance in the high SF range must exceed that for chrominance and the opposite must happen for the low SF range.

5. Conclusions

A new hyperspectral data-set of 29 natural scenes has been acquired and made available for spatio-chromatic analysis. Our preliminary analysis of the ratio of luminance to chrominance Fourier amplitudes in these scenes suggests a mismatch between spatial frequency content of the natural world and the human visual system's contrast sensitivity functions. The mismatch might be due to quantum noise suppression as discussed by Atick. The $1/f$ law that Atick's analysis assumes has been confirmed to approximately hold for both narrow and wide-band chromatic spectra. We have discussed the following possible reasons why the HVS appears to discard high SF chrominance: (a) its irrelevance to primate vision; (b) its removal the neural machinery that may compensate for the chromatic aberrations; (c) lossy data compression in response to physiological constraints.

6. Acknowledgments

The construction of the hyper-spectral camera was supported by a grant from the UK Defence Research Agency. Grant No D/ER1/9/4/2034/102/RARDE. We thank Derek Carr for technical support and Kay Nelson and Julian Partridge for help with image acquisition. Comments from two anonymous reviewers have been particularly helpful. The "*shadow-removing*" definition of chrominance benefited from discussion with Fred Kingdom.

6. References

- 1- H.B. Barlow, "Possible principles underlying the transformation of sensory messages", Sensory Communications, W.A.Rosenblith (Ed.), MIT Press, Cambridge MA (1961).
- 2- M.V.Srinivisan, S.B.Laughlin and A. Dubs, "Predictive coding: a fresh view of inhibition in the retina.", Proc. R. Soc. Lond. B 216, 427-459 (1982).
- 3- J.J. Atick, "Could information theory provide an ecological theory for sensory processing?", Network 3, 231-251 (1992).
- 4- D.J. Field, "Relation between the statistics of natural scenes and the response properties of cortical cells", J.Opt. Soc. Am. A, Vol 4, 12, 2379-2394 (1987).

- 5- D.J.Tolhurst, Y.Tadmor and Tang Chao, "Amplitude spectra of natural images", *Ophthalm. Physiol. Opt.*, Vol.12, 229-232 (1992).
- 6- G.J. Burton and I.R. Moorhead, "Color and Spatial structure in natural scenes" *Applied Optics*, 26, 157-170 (1987).
- 7- G. Brelstaff and T. Troscianko, "Information content of natural scenes: implications for neural coding of color and luminance", *S.P.I.E. Vol. 1666, Human Visual Processing and Digital Display*, B.E.Rogowitz (Ed.), 302-309 (1992).
- 8- J.B. Derrico and G. Buchsbaum, "A computational model of spatiochromatic image coding in early vision," *J. of Visual Com. and Im. Rep.* 2, 31-37 (1991).
- 9- S.M. Courtney, L.H. Finkel and G.Buchsbaum "A multistage neural network for colour constancy and color induction", *IEEE Trans. on Neural Networks*, 6, No 4, 972-985 (1995).
- 10- C. Koch, "Computation and the single neuron", *Nature* 385, No 16, 207-210 (1997).
- 11- K.T. Mullen, "Contrast sensitivity of human color vision to red-green and blue-yellow chromatic gratings," *J. Physiol.* 359, 381-400 (1985).
- 12- R.L. De Valois, H. Morhan, D.M. Snodderly "Psychophysical studies of monkey vision III. Spatial luminance contrast sensitivity tests of macaque and human observers", *Vision Res* 14, 75-81 (1974).
- 13- A.M. Derrington and P. Lennie. "Spatial and temporal contrast sensitivities of neurons in lateral geniculate nucleus of macaque," *J. Physiol.* 357, 219-240 (1984).
- 14- A.M. Derrington, J. Krauskopf and P. Lennie. "Chromatic Mechanisms in lateral geniculate nucleus of macaque," *J. Physiol.* 357, 241-265 (1984).
- 15- R.L De Valois and K.K De Valois, "A Multi-stage Color Model", *Vision Research* 33, No 8, 1053-1065 (1993).
- 16- C.R. Ingling Jr. and E. Martinez, "The spatiochromatic signal of the r-g channel," *Color Vision: Physiology and Psychophysics*, Ed. by J.D. Mollon & L.T. Sharpe, Academic Press Inc. (London) Ltd., 433-444 (1983).
- 17- F.A.A. Kingdom and K.T. Mullen, "Separating color and luminance information in the visual system", *Spatial Vision*, Vol 9, No 2, 191-219 (1995).

- 18- M. D'Zmura and P. Lennie,. "Mechanisms of color constancy", J.Opt. Soc. Am. A, Vol. 3, No 10 (1986).
- 19- C. Smith and J. Pokorny. "Spectral sensitivity of the foveal cone photopigments between 400 and 500 nm," Vis. Res. 15, 161-171 (1975).
- 20- D. Travis, "Effective Color Displays", Academic Press, London, Appendix 4, 271-272 (1991).
- 21- C.R. Ingling Jr. and B. H. Tsou, "Spectral sensitivity for flicker and acuity criteria," J. Opt. Soc. Am. A 5, 1374-1378 (1988).
- 22- G. Buchsbaum and A. Gottschalk, "Trichomacy, opponent color coding and optimum color information transmission in the retina," Proc. R. Soc. Lond. B 220, 80-113 (1983).
- 23- G. Brelstaff, A. Párraga, T. Troscianko and D. Carr, "Hyperspectral camera system:- acquisition and analysis," S.P.I.E. Vol. 2587, Geographic Information Systems, Photogrammetry, and Geological/Geophysical Remote Sensing, J.B.Lurie, J.J.Pearson, E.Zillioli (Ed.), 150-159 (1995).
- 24- J.D. Mollon, "Tho' she kneel'd in that Place where they grew... The uses and origins of primate colour vision", J. Experimental Biology 146, 21-38 (1989).
- 25- Y. Le Grand. "Form and space vision," Translated by Michel Millodot and Gordon G. Heat, Indiana Univ. Press, Bloomington and London, 5-35 (1967).
- 26- M.V. Berry and A.N. Wilson, "Black and white fringes and the colours of caustics", Applied Optics 33, No 21, 4714-4964 (1994).