Supplementary Material

Spatiochromatic Properties of Natural Images and Human Vision

C.A. Párraga, T. Troscianko, and D.J. Tolhurst

Supplementary Experimental Procedures

Obtaining the Images and Image Analysis

Camera Calibration

We carefully calibrated a digital camera (Nikon Coolpix 950) so that each pixel in the (uncompressed) output image represented the capture of each of the three human cone types. The calibration process utilized a set of narrowband optical interference filters spanning the range from 400 to 700 nm, a constant-current light source, a white target made of cyanoacrylate powder, and a measuring device (a Topcon SR-1 telespectroradiometer calibrated against a National Physical Laboratory standard light source). This calibration process removed the various nonlinearities introduced by the camera electronics and compensated for the spectral responses of the three camera CCD sensors. The output of this process was a trichromatic representation of each scene giving, for each pixel, the relative capture ratios of the three human cone types, based on the Smith and Pokorny cone fundamentals. These scenes (labeled L, M, and S) were then combined to obtain luminance and chromatic representations.

Since the spectral sensitivities of the RGB elements in the CCD camera do not exactly match the spectral sensitivities of the LMS cone fundamentals, the three LMS cone catches cannot be computed unambiguously. This becomes evident, as camera metamers will differ from human metamers. Given the way our camera was calibrated, we cannot compensate for this, but we can provide an estimation of how large the errors are likely to be. To estimate the camera's errors in determining the L, M, and S cone catches, we photographed eight different objects under controlled illumination conditions in our lab. These objects were five different leaves and three red fruits (commonly found in our picture data set). Figure S1 shows the thumbnails of all of them.

The pictures were then converted to LMS cone representations using our standard procedure. The central 50×50 pixels (avoiding strong specularities, especially in the red fruit) were averaged, and the standard deviation was calculated. The value of this average was normalized by dividing by the total cone response ($L + M + S$). A similar value was obtained by measuring approximately the same spot with the Topcon SR-1 telespectroradiometer and converting its spectral radiance output to the equivalent L, M, and S cone catches. Both normalized values (obtained with the camera and the SR-1, respectively) were compared, and their mean differences was found to be 9.2% for the L signals and 9.5% for the green signal. The largest differences in both measurements coincided with the largest standard deviation of the average of pixels within the 50×50 central box. This is due to the fact that the red surfaces were producing some specular reflections and their color was not uniform. The same experiments performed on a nonspecular, evenly colored surface (Macbeth color chart) give smaller differences (less than 5%) for the L, M, and S signals obtained with the two methods.

Image Processing

A standard two-dimensional fast-Fourier transform algorithm was used to derive the amplitude spectrum for a given *luminance* or *chromatic* image. *Luminance* (lum) was defined as ($L + M$), and *chromatic* images were either red-green [$(L - M)/\mathrm{lum}$] or blue-yellow [$(S - \mathrm{lum})/(S + \mathrm{lum})$]. The DC component was removed. The usual way of calculating the amplitude slope requires that all sf components of the picture are averaged regardless of orientation, and this average is plotted on log-log axes as shown in Figure S2.

The usual next step would be to calculate the slope of the line and to report the number as α. In this approach, the high-sf range of the spectrum is overrepresented when we measure the slope, and these points will determine the value of α. Our approach is somewhat different.

Spectral slopes were measured by dividing the Fourier space into nine circularly symmetric, logarithmically spaced, one-octave spatial frequency bands (see Figure S3A) and averaging the Fourier content within each of the bands. The averages were then plotted against the midspatial frequency of the band in log-log coordinates, and the slope (α) of this line was calculated by linear regression. This prevented any bias due to oversampling of the high spatial frequencies. The ordinate axis in Figure S2 shows the distribution of these bands in the spatial frequency domain.

Figure S1. Eight Different Objects Used for the Estimation of the Camera's Error

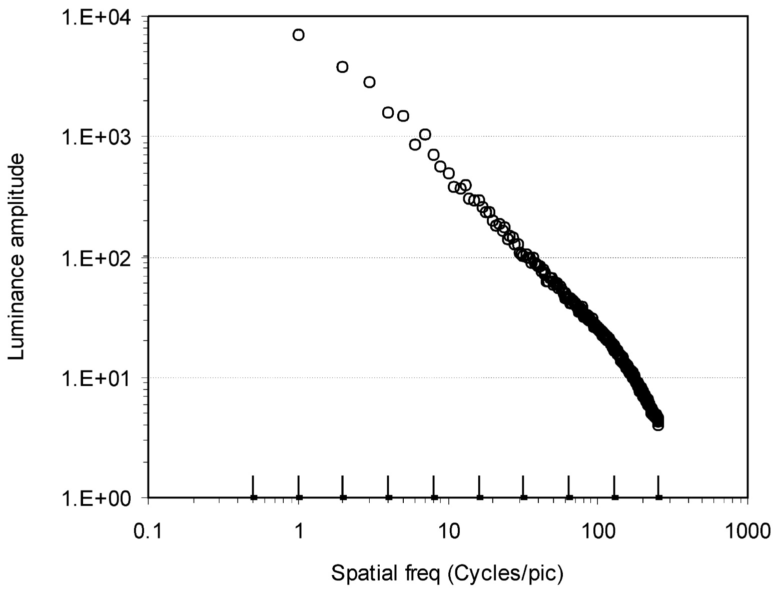

Figure S2. Example of a Log-Log Plot Representation of Fourier Amplitude Averaged across Orientations

Counting the Number of "Red" Pixels

Figure S4 shows the distribution of pixels according to their L and M cone catches for all of our data sets. To produce this figure, all images were converted to LMS, and then the relative values $l = L/(L + M + S)$ and $m = M/(L + M + S)$ were computed. Pixels that comply with the condition $l > 1.5^* m$ are shown in red, and the others are shown in green. The vertical axis shows the log values of the number of pixels. The figure shows that there is a bimodal distribution of the two types of pixels, and this distribution is optimally bisected by the criterion $l > 1.5^* m$, as shown in the plots.

Calculation of the Optimal Ratio of Spectral Slopes

To calculate this, we fitted functions to Mullen's (1985) contrast sensitivity functions (csfs) data. Figure S5 shows these. The sf axis (*x* axis in Figure S5) was divided into logarithmic sf bands similar to those in Figure S3A. We used the angular subtense of the camera optics to convert the bands from cycles/picture to cycles/°. After this, we calculated the areas under each of the human contrast sensitivity functions for every band. The total areas were normalized to unity in both cases.

After this, we proceeded in similar fashion with a pair of exponential functions (which are similar to the linear Fourier amplitude functions when plotted in log-log space, [see Figure S6]). One of these (plotted in blue in Figure S6) had a slope of −1.1, corresponding to our average value of α for "normal scenery". The slope of the second function (plotted in red in Figure S6) was our free parameter. We found that a ratio of slopes of 0.76 for these two curves makes the ratio of the areas under them fit optimally with the ratios of the areas under Mullen's csf data.

Errors Introduced by the Camera's Optics

Estimation of the Variability of the Amplitude Slope with Optical Characteristics of the Camera

Optical characteristics of the camera (aperture, zoom settings, etc.) were varied as little as possible to avoid introducing artifacts into the estimation of α. The smallest possible aperture (around f/11) was used in most cases in order to maximize depth of focus, and the preferred focal length was "telephoto" to avoid interference between the camera operator and the scene. The use of the smallest aperture ensured that most of the objects in the scene were in focus (largest depth of field). Of the scenes in our data set, 52 had a focal length of 19 mm (aperture f/11) and 67 had a focal length of 20.4

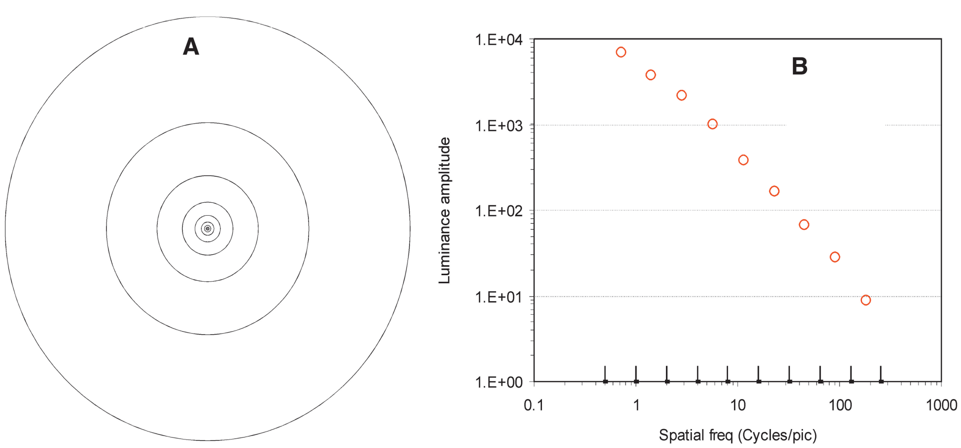

Figure S3. Calculation of the Slope of the Amplitude Spectrum

(A) A scheme of the logarithmically spaced one-octave sf bands used to calculate the slope of the amplitude spectrum (α).

(B) Data from the previous figure after the Fourier amplitude has been averaged within the sf bands.

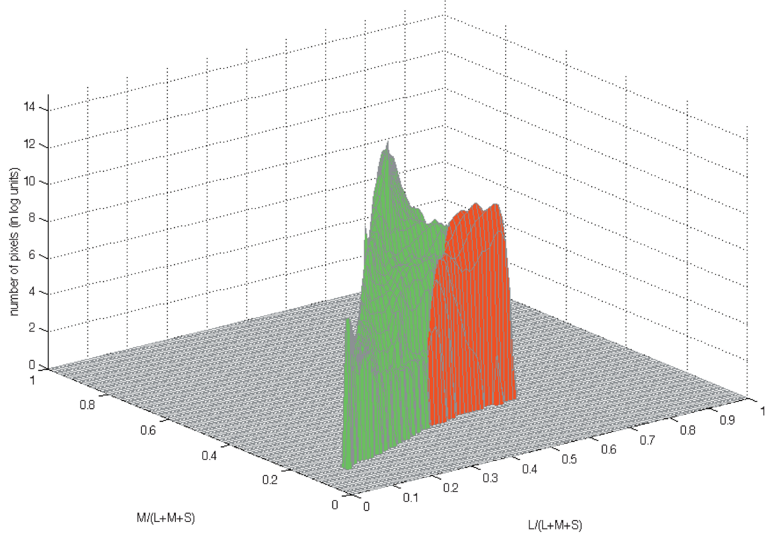

Figure S4. Distribution of the Pixels for All the Image Data Set in the Form of a Cone-Activation Chromaticity Diagram

The axes correspond to L and M cone ratios. Here, it is possible to see the bimodal distribution of "red" and "green" pixels in our data set.

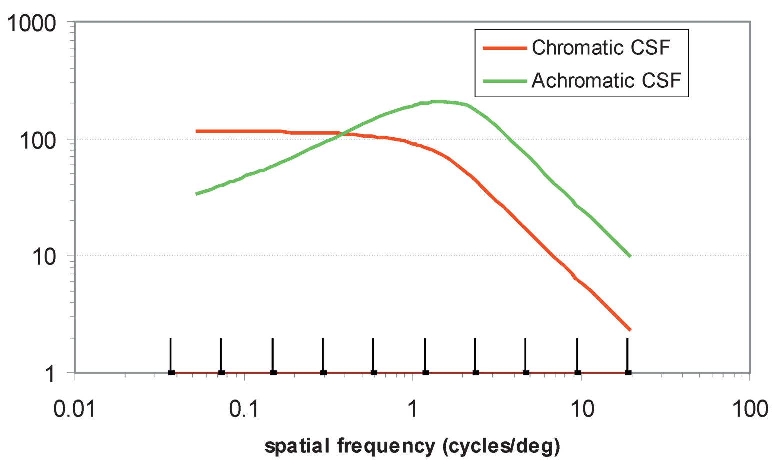

Figure S5. Plot of the Functions Fitted to the Mullen Achromatic and Chromatic Contrast Sensitivity Functions Data

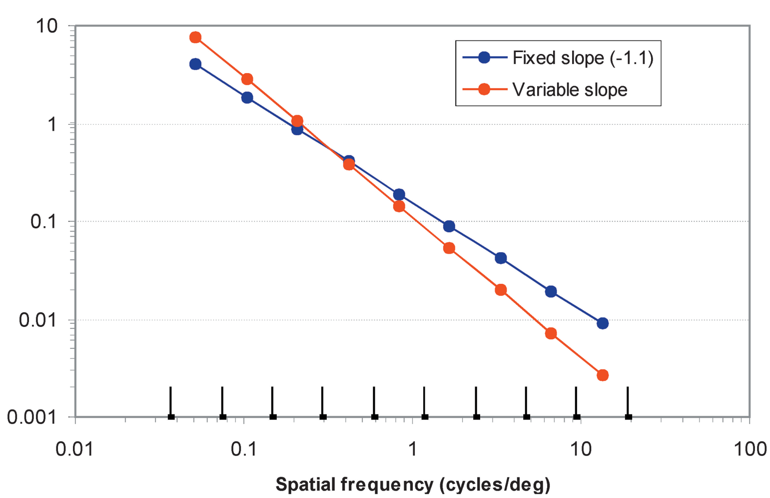

Figure S6. Plot of Two Exponential Functions that Become Straight Lines when Plotted in Log-Log Space

The slope of the line in blue was fixed to −1.1. The slope of the other line was varied so that the ratio of areas under each of the lines optimally fit the ratio of the areas under the csf data. The sf bands were chosen to be the same as in Figure S5. The areas were normalized to unity and calculated for each of the sf bands in all cases.

mm (aperture f/11.4). Only five had a 7.2-mm focal length (aperture f/7), and here we explore how this may have influenced our results.

To estimate the variability in the measurement of the luminance, red-green, and yellow-blue α slopes, we recorded a typical image of a red fruit and green foliage at different viewing distances but using different zoom settings (focal length = 7.2, 9.3, 10.8, 12.4,

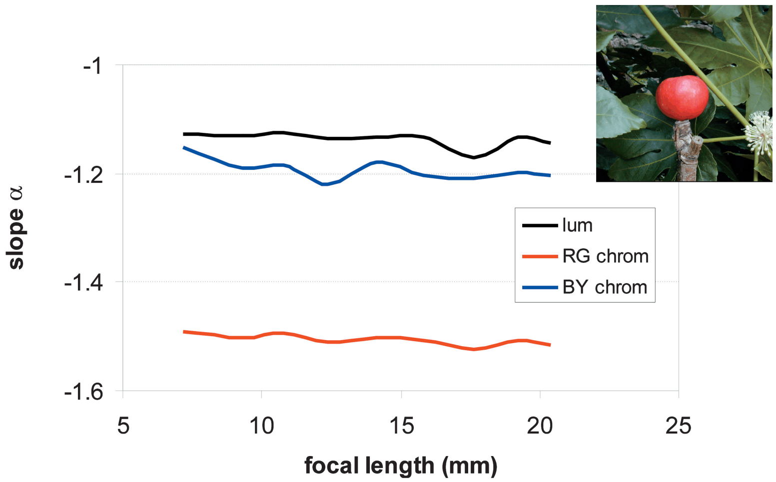

Figure S7. Variation of the Slope α with the Focal Length of the Camera Lens

14.1, 15.8, 17.6, 19.2, and 20.4 mm) in such a way that the resultant subtense of the fruit was the same in all of the images. What varies is the degree of foreshortening of linear perspective in the images. We wished to know whether these effects on linear perspective would have any effect on our dependent variable, the spectral slope α. The values of α of the luminance and chromatic representations were measured in all cases. Figure S7 shows these results. It is clear that there are small fluctuations of α around the average value (possibly due to small changes in the distribution of the objects in the different pictures or changes in the illumination), but it is also clear that there is no correlation with increasing values of focal length, and therefore no effect of linear perspective on the spectral slope of this kind of image.

Errors Introduced by the Camera Lens

To estimate the effects of chromatic aberration introduced by the camera lens, we took a picture (under controlled incandescent light conditions and under cloudy skylight) of a target that consisted of a white piece of paper with some black letters in bold characters. The lens was set to small aperture (f/11.4). This produced a picture with a broad spatial frequency content. The three R, G, and B planes that constitute the raw picture were separated, and their contents were Fourier analyzed. The slope (α) of the amplitude spectrum was measured for the three planes separately. The plots of Fourier amplitude versus spatial frequency in log-log coordinates were very similar for the three planes in each of the pictures. Table S1 (below) shows the values of α for two lighting conditions.

Results in Table S1 show that the differences in amplitude slope between the different planes are very small. In the case of the cloudy skylight illumination, this difference is about 3% in the worst case. Moreover, a comparison of the slope sets for the two illuminants suggests that the small differences in α are a result of signal-noise ratio effects arising from the illuminant rather than because of the lens. A very similar set of numbers was also obtained for a much larger aperture (f/4), suggesting that the lens is generally well designed to eliminate chromatic aberration.

Chromatic aberration occurs when light of one wavelength is in focus when light of another wavelength is not. Since the target was black and white and was illuminated by incandescent light, it is possible that this measure of total chromatic aberration averages all values of intensity at different parts of the spectrum into one single value. To test whether this might be the case, a second experiment was performed.

Table S1. Slopes of the Amplitude Spectrum for Each of the RGB Planes that Constitute the Test Picture

Slope (α)	R	G	B
Incandescent light	−1.054	−1.0368	−0.995
Cloudy skylight	−1.077	−1.087	−1.111

Using the same test chart, two pictures were taken through chromatic filters. The first picture of the chart was taken through a 610-nm narrowband chromatic filter, and the second picture was taken through a 520-nm narrowband chromatic filter. The filters were 10-nm narrowband and were chosen to be close to the points of maximum sensitivity for the L-M (red-green) system. The focal distance was constantly maintained throughout the whole procedure so that the first image was in focus. The luminance pictures were produced in both cases (see above for details), and their spectral slopes were measured and compared. The measured values of α were −1.151 and −1.126 for the 610- and 520-nm filters, respectively. The difference is about 2%, and it can be safely ignored in our calculations for the purposes of this paper.