

Hyper-spectral camera system:- acquisition and analysis.

Gavin Brelstaff<sup>†</sup>, Alejandro Párraga, Tom Troscianko & Derek Carr

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

### ABSTRACT

A low-cost, portable, video-camera system built by University of Bristol for the UK-DRA, RARDE Fort Halstead, permits in-field acquisition of terrestrial hyper-spectral image sets. Each set is captured as a sequence of thirty-one images through a set of different interference filters which span the visible spectrum, at 10 nm intervals: effectively providing a spectrogram of 256x256 pixels. The system is customised from off-the-shelf components. A database of twenty-nine hyper-spectral images sets was acquired and analysed as a sample of natural environment. We report the manifest information capacity with respect to spatial and optical frequency drawing implications for management of hyper-spectral data and visual processing.

### 1. INTRODUCTION

Much recent effort has been expended in the development of hyper-spectral imaging devices for both terrestrial and remote-sensing applications. For example, the SpectraCube 1000 - now available from Metax in the UK permits terrestrial scenes to be acquired at 5 nm intervals within the spectral range 400 to 1000 nm. Hyper-spectral airborne devices such as AVIRIS from NASA provide 31 x 10 nm bands - and more recently from TRW the TRWIS system is delivering 92 x 3nm bands. Hyper-spectral space-borne imaging sensors are also planned in the near future - e.g. the MODIS sensor with 36 bands. These developments are motivated by the desire to see structures in the world at a chromatic resolution higher than that of the human eye perhaps to reveal oil spills, forest canopy features, oil and gas reserves and other environmentally interesting sites. Naturally there is also a military interest in breaking camouflage on the battle field. The progression from fairly wide-band, tri-chromatic and multi-spectral imaging to hyper-spectral imaging with its many narrow-bands is accompanied by a massive increase in the amount of data that can potentially be stored. Data management decisions must be addressed to determine which data are stored and archived. For particular applications a data analysis may determine which combinations of bands provides the best results - e.g. to prospect for a particular mineral. However, when the plan is to provide high chromatic imagery for multiple purposes, then the decision is more difficult to make. One factor to take into account is the average information content carried in various parts of the spectrum when viewing sites of interest. If contrast modulation is greater in one part of the spectrum than another then it would make sense to favour it in the data representation. Here we measure the average spatio-chromatic information content in a sample of natural images of terrestrial scenes in order to explore suitable hyper-spectral data representations. It is interesting to speculate how the human visual system has evolved into a three-band system in response to the spatio-chromatic characteristics of the natural environment.

Several workers<sup>1-7</sup> in the field of human vision have explored efficient representations of the information in the visible environment and their potential evolutionary advantages. From this they have made predictions of the type of processing occurring in visual neurones. Computational strategies<sup>2-5</sup> for such processing have been proposed from physical properties of the stimulus environment. Again, these strategies ought to be based on quantitative knowledge of statistical properties of natural scenes.

A first attempt to relate the statistical structure of the environment to the coding properties of the visual system was made by D. Field<sup>4, 5</sup>. He processed a set of achromatic natural images and showed that the information content varies with spatial frequency. Roughly equal information is contained within each octave of spatial frequency - i.e. on average the Fourier amplitude spectrum falls off inversely with the spatial frequency. This is commonly known as the 1/f law of

---

<sup>†</sup> Now at National Remote Sensing Centre Ltd, Farnborough, GU14 ONL, UK

fractal-like natural textures. An analysis of hyper-spectral terrestrial scenes, made here, appears to confirm that this law holds for narrow-wavelength bands uniformly throughout the visible spectrum. This analysis is of 29 scenes acquired using a specially designed camera system. The cost of the system was kept low by combining and adapting several off-the-shelf components:


- 1- Standard 2 inch square interference filters.
- 2- A modified Kodak circular-carrousel slide changer
- 3- Customised stepper motor drive and electronics
- 4- Standard Pasecon integrating camera tube
- 5- Standard lens optics
- 6 -Data Translation image grabbing card
- 7 -Compaq portable PC, battery power supply and inverter.

The system was completed, assembled, calibrated and operated by members of the Perceptual Systems Research Centre of the University of Bristol between 1990 and 1993 - under funding from the UK Defence Research Agency.

## 2. ACQUISITION

### 2.1 Characteristics of the hyper-spectral camera

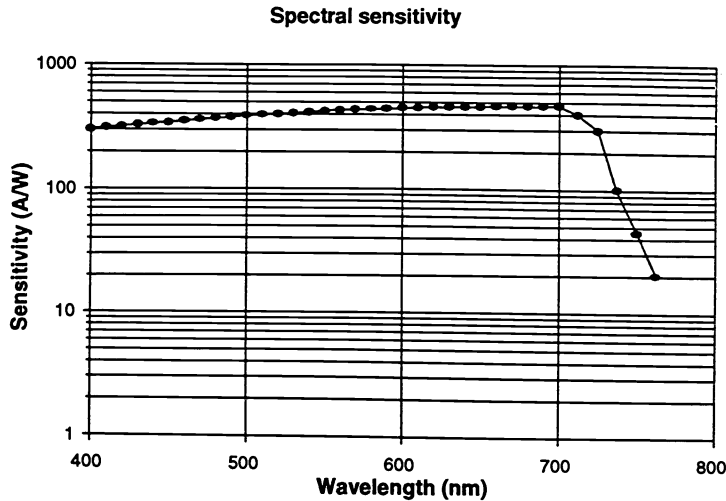
The hyper-spectral camera consists of an electro-optical mechanism built around a "Pasecon" tube, a camera control unit (CCU), a set of narrow-band interference filters, a carousel slide changing filter mechanism, a portable 386-PC, a real-time video monitor and a battery power supply. The scheme of the hyper-spectral camera is shown in Figure 1.



**Figure 1:** Scheme of the hyper-spectral Camera.


The Pasecon tube was chosen because it has good linearity throughout the visible spectral spectrum ( 400-700 nm) and an ability to integrate signal over a user-selectable integration time. The latter is important because it allows a reasonable signal to noise ratio to be achieved even at the low light levels incurred by placing a narrow band filter in front of the target. Figure 2 illustrates the camera tube sensitivity between 400 and 800 nm (from manufacturer's data). The camera converts the video signal to digital form in the CCU. A special card was design to enable clean and noise-free digital images to be transferred directly into the frame-store on the Data Translation frame grabber card on the PC.

The slide changing filter mechanism was added to the e-o mechanism to allow the hyper-spectral camera to sequentially grab images through a set of 31 different optical interference filters. These are chromatically narrow-band in the range of 400 to 700 nm and with nominal 10 nm spacing between their peaks.



**Figure 2:** Camera tube sensitivity between 400 and 800 nm (from manufacturer's data).

Figure 3 shows a plot of transmittance versus wavelength for the whole set of filters employed. The entire system is controlled by the portable 386-PC and mounted on a trolley along with a 12 volts inverter to supply 240 volts mains in the field. Manual fixed focal length lenses (Fuji CF25B, f/1.4, 25 mm) with a field angle of 28.71 x 21.73 deg were used. The marked aperture settings on the lens are: f1.4, f2, f2.8, f4, f5.6, f8, f11 and f16.



**Figure 3:** Transmittance of the spectral interference filters employed.

A PC program (written in 'C' under MS-DOS 3.0) sends a value for the integration time to the hyper-spectral camera control unit which grabs the image and transfers it from its frame card to the PC frame card. The whole system allows a sequence of 31 chromatically narrow-band filtered 8-bit images (256 x 256 pixels x 256 grey-levels) to be grabbed. These images correspond only to the central part of the visual field supplied by the lens. The field angle of this picture is thus equal to 14.35 x 14.35 deg. Given this, each pixel subtends an angle of 0.056 x 0.056 deg (3.36 arc minutes) approximately. This value is of the order of the size of a foveal cone diameter (1 arc minute approximately). Once

recorded, the image set is transferred to an IBM (RS/6000) workstation for processing.

A graphical user interface (written in 'C' under AIX on the RS/6000 accessing the OSF/Motif X-Windows System) allows convenient access to the data as grabbed image sets. Using this interface it is possible to display each image, in turn, on screen and to display a spectrograph of the light at the point in the image where the mouse cursor has been pressed in the display window. Either grey level, radiance or derived reflectance can be plotted. Figure 4 illustrates this facility. Radiance estimation and reflectance derivation depend on the calibration process described below.




Figure 4: Display window (OSF/X-Motif) of the IBM-RS/6000 employed to process the dataset.

## 2.2 The hyper-spectral camera calibration

The hyper-spectral camera system outputs images as disk files containing the relevant header information including the integration time and lens aperture with which multi-spectral light-measurements are to be reconstructed. All measurements described below were obtained after a warm-up period of two hours - after which stable signals were produced.

The calibration strategy consists of three main steps:

- Finding for any given filter and aperture setting the number of integration frames required to achieve a reasonably large dynamic range.
- Correcting the non-linear signal characteristics of the hyper-spectral camera and the off-axis variation across the lens and camera target.
- Converting the grey level output of the hyper-spectral camera into measurements of the spectral radiance (and then spectral reflectance) across the target.

### 2.2.1 Selecting the dynamic range for each image

Very bright regions like sky and specularities are not of direct interest, and would severely compress the dynamic range (0, 255) if properly represented. No sophisticated strategy was used to identify these bright regions. Instead, the system was created to allow the adjustment of the integration time until bright regions, other than sky and specularities, produce responses near to, but not above, the maximum. The real-time video monitor provided an effective way to identify these bright regions on the field. For example, it is possible to choose a ceiling value of 254 (in grey-levels) so that 90% of the recorded scene will not exceed it. Alternatively, it is possible to choose among nine small square regions regularly distributed on the scene and manipulate its statistics (e.g. median, mean, maximum). For example, one might select the median value within the small square region in the centre not to exceed the ceiling of 240.

Once both the statistics (S) and ceiling (C) are defined an algorithm automatically find a value for the integration time (IT)

so that  $S(IT)$  is near to, but not above,  $C$ . The accepted difference between  $C$  and  $S$  can be specified by a tolerance parameter. Given that this calibration is repeated for each filter on the carousel, the algorithm has been designed to accelerate the process. This algorithm can predict a suitable value for the  $IT$  using the results of previous attempts, and check if the light levels change between individual trials. If that happens then the calibration process is halted and the problem reported to the user.

### 2.2.2 Corrections of non-linear characteristics of the hyper-spectral camera

To correct for the non-linear characteristics of the camera signal a look up table (LUT) was produced as follows. The hyper-spectral camera was pointed at a white paper test card, the radiance of which was  $0.002 \text{ Wsr}^{-1} \text{ nm}^{-2} \text{ m}^{-1}$  as measured with a TopCon Spectra-radiometer (model SR1). The test card was lit using a steady illuminant, a tungsten bulb with constant current through it. To avoid possible contributions from light with wavelengths outside the visible range the 570 nm interference filter was always used and all measurements were taken in a dark room. As Figure 2 shows, the camera tube sensitivity decays rapidly in the infra-red (IR) part of the spectra ensuring that even if the filter has a secondary peak in the IR - which was difficult to determine - then it would have little or no effect on the measurements made. Figure 5 shows the apparatus set up.




Figure 5: Apparatus set up for calibration.

The system output within a central square region of the image was measured whilst varying the integration time from 1 to 200 frames for all 8 lens aperture settings. This provides a series of eight isophote curves. To obtain a further eight curves a transmission filter with neutral density 0.5 was slotted into the camera behind the 570 nm filter. The procedure described above was repeated 10 times to obtain variability data and hence 16 exemplary isophotes (one for each of the above conditions) were selected and used to create the look-up table by a process of inter-isophote interpolation.

The same procedure used for calibrating the centre of the scene was repeated for nine other regions evenly distributed across the image. An additional set of 8 further regions lying on the perimeter of the image was used to provide extra information about the off-axis output. A bi-quadratic spatial interpolation was used to approximate values between the measured regions - thus resulting in an off-axis attenuation mask for the image. As before, this data was stored in an off-axis correction file that was simultaneously used with the LUT to convert the hyper-spectral camera output into light measurement.

### 2.2.3 Conversion of camera signal into spectral radiance

To interpret camera grey-level signal as spectral radiance a conversion factor is required for each filter. This was obtained by matching the output of the hyper-spectral camera through each of its 31 filters to that of the TopCon SR1 when both systems were pointing at a standard Kodak grey card. Further measurements of the spectral radiance by the hyper-spectral camera were compared to that obtained with the TopCon SR1 using a collage of coloured papers illuminated with a constant current tungsten lamp. The relative error of the matching between both systems was less than 5% in the range

400-570 nm. In the range 580-700 nm the relative error was larger but always less than 10%. For one red sheet the match for some unknown reason was out by as much as 20% in the spectral range 650-700 nm. Figures 6(a) and 6(b) provide an example of this matching.




Figure 6(a): Spectral radiance obtained with the hyper-spectral camera.

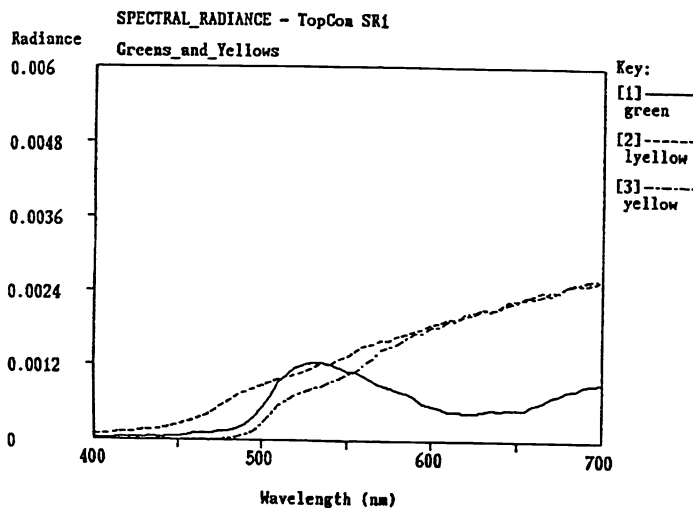


Figure 6(b): Spectral radiance obtained with the TopCon SR1 point spectral radiometer.

To convert the hyper-spectral camera output into measurements of spectral reflectance, it is necessary to know the spectral characteristics of the light falling over the scene. These characteristics can be estimated by placing an object of known reflectance (e.g. the Kodak standard grey card) into the scene and measuring its radiance. By keying-in the location of the grey card in the images our system is then able to perform the conversion automatically. It does this by inverting the following approximation

$$\text{scene radiance} = \text{illumination} \times \text{reflectance}.$$

This relationship assumes that light falling over the grey card is an accurate representation of that falling over the rest of the scene and that the reflections are Lambertian - i.e. that they are independent of the viewing angle as expected for diffuse, not specular, surfaces.

### 2.3 Scene acquisition

A dataset of 29 scenes was obtained between September 1993 and January 1994. Each of these scenes contains 31 images taken through a different spectral filter. The normal recording procedure was as follows:

- a) Positioning of the Kodak grey card in a visible place in the scene - to allow dynamic calibration and light checking.
- b) Measuring the spectral radiance of the card using the TopCon SR1. This was done to detect changes in the spectral characteristics of the illumination over the time each scene was acquired. The measurements were taken from approximately the same place as that of the hyper-spectral camera and just before starting the recording of the scene.
- c) Recording of the scene with the camera. Considering the relative long time (about 5 minutes) needed to acquire each scene it was often necessary to wait until weather conditions appeared acceptable before starting.
- d) Making the second spectral measurement of the grey card radiance - to be compared with the first.
- e) Back in the laboratory the PC was connected to the network and the grabbed images were transferred into the UNIX workstation for analysis.

### 2.4 Common practical problems and their solutions

Most of the practical problems we faced were related to changes in the scene during the recording. We consider each in turn below.

- a) Linear light changes. These were generally caused by changes in the sun's position during the recording of the scene. These alterations had little effect on the spectral properties of the light and were detected by comparing the SR1 measurements before and after the recording. To avoid them, most of the recording was done around noon and a special algorithm "reilluminate" was developed to minimise their effects. The algorithm works with reference to the camera signal of the grey card - as described below.
- b) Light fluctuation, mainly due to short-term variations in cloud cover. This generally affected one or two of the images in any scene. To detect the fluctuation it was necessary to compare the SR1 spectral measurements on the grey card with those of the hyper-spectral camera in the same place. The most effective way to avoid fluctuation was to wait until the sky was completely clear or completely overcast before grabbing the scene. Small errors were corrected with the "reilluminate" program.
- c) Small movements of objects, such as tree branches in the wind, in the scene during acquisition. At short range it was necessary to either wait for the movement to stop or work in a sheltered location - e.g. in the glass houses of the University of Bristol Botanical Gardens. Otherwise, the effect was reduced by taking only long and middle distance shots.

### 2.5 Corrections made by software

The "reilluminate" program was created to correct for small light fluctuations and linear light changes of particular scenes of the dataset. It uses the spectral radiance image of the scene along with one of the SR1 measurements taken of the grey card. It "reilluminates" the scene with the light spectrum falling over the grey card so that the radiance measured with the hyper-spectral camera on the grey card matches that of the SR1 measurement. This correction was only deemed necessary for five particular scenes of the dataset. To decide whether it was necessary to use this program or not, the following points were considered:

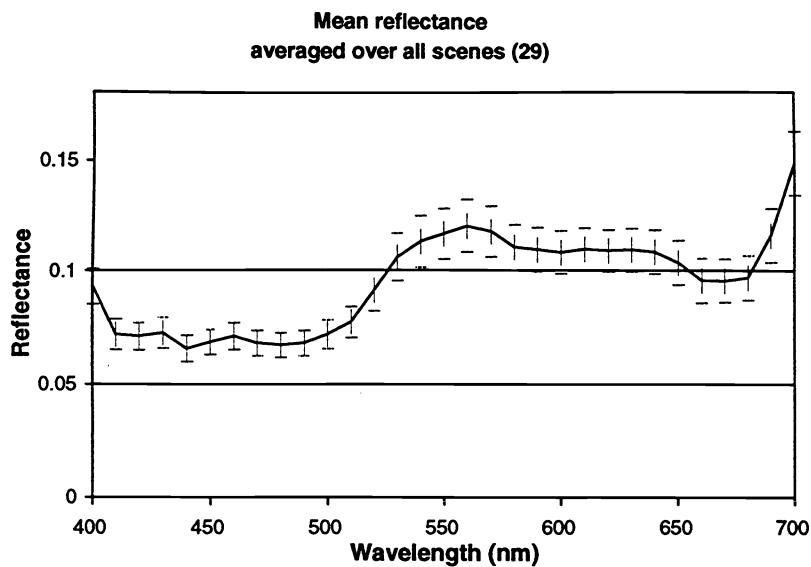
- a) The goodness of fit between the two SR1 measurements taken on the grey card before and after the recording of each scene. This fit models a linear change in the overall brightness of the light before and after.
- b) Comparison between the spectral radiance of the grey card as measured by the SR1 and the spectrum obtained by the hyper-spectral camera in the same place.
- c) Light and wind conditions on the day when the scene was recorded. This information was registered in a table for each scene of the dataset.

### 3. ANALYSIS

For the 29 scenes acquired by the camera system we report the average information content with respect to wavelength. For this purpose we report the average variation over wavelength of:

- a) the magnitude of the signal.
- b) the modulation of the signal.

We chose the derived reflectance as the signal, rather than the radiance, so that we aim to report properties of the objects in the scene rather than those of the particular illumination that happened to be prevailing. Thus the average magnitude reported is that of mean derived reflectance - this varies across the 31 wavebands as plotted in Figure 7. It is interesting to note that the minima in the mean reflectance are roughly coincident with the absorption peaks of the two most common types of chlorophyll.



**Figure 7:** Distribution of mean spectral reflectance from all scenes.

For the average modulation we report average Fourier amplitude occurring within each narrow-band over different ranges of spatial frequency. For this purpose, a Fourier transform is made of the 31 narrow-band reflectance images of each scene. It is desirable for the average modulation to represent all scenes equally. For this reason we normalise each scene - by scaling its constituent images - so that it contains unit Fourier amplitude. Here we exclude the zeroth frequency component as this is actually the reflectance reported above - and would otherwise bias our measure of modulation. In order, to report the full dynamic range of amplitude with spatial frequency a logarithmic scale is necessary. For this purpose Fourier space is divided into 8 concentric rings each of them representing a given range of the spatial frequency spectrum - as tabulated in Table 1 and illustrated in Figure 8. These rings are labelled as band 1 to 8 in order of increasing spatial frequency. Figure 9 shows the average amplitude over all scenes occurring within each of these rings per wavelength band - i.e. it shows 8 x 31 average values. As wavelength varies the values within each ring appear remarkably constant. Figure 10 shows a three-dimensional plot of the same data from which it is easier to visualise the slope of amplitude against spatial frequency. When this slope is 1.0 then Field's 1/f law is obeyed<sup>4</sup>. The value of the slope varies very little with wavelength - exhibiting a range between 1.08 and 1.16. Thus we conclude that the 1/f law seems to be a good approximation throughout the visual spectrum for the data that we acquired. The uniformity of modulation with wavelength found within our sample of natural scenes is in contrast to prediction that might be made from psychophysical experiments<sup>8</sup> and theories of human visual processing<sup>9</sup>.






Figure 8: Division of the Fourier space into 'logarithmic' ranges.

Spatial Frequency	Range	Centre
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

Table 1: Table of the range and centre of each spatial frequency ring used to divide Fourier space..

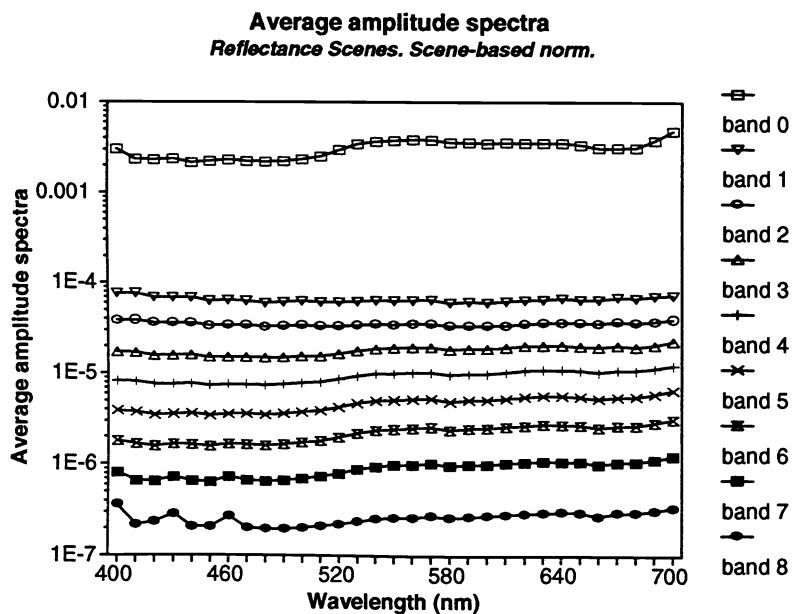
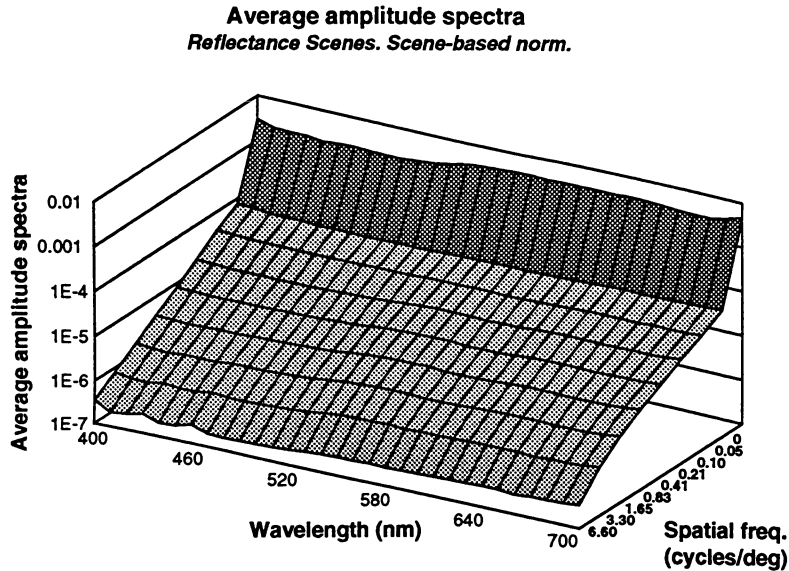


Figure 9: Variation of the average Fourier amplitude across the visible spectrum for all spatial frequency ranges: band 1 to 8. Note band 0 is the mean reflectance - reported in figure 7.



**Figure 10:** *Three-dimensional plot of the previous figure.. One of the axes displays the central value of spatial frequency (in cycles/deg) corresponding to the SF bands.*

#### ACKNOWLEDGEMENTS

This work was funded by the UK Defence Research Agency under grant D/ER/1/9/4/2034/102. We thank the Bristol University Botanical Gardens and Julian Partridge for assistance in acquiring scenes including plants and animals.

#### REFERENCES

1. Brelstaff G. and Troscianko T. 1992. *Information content of natural scenes: implications for neural coding of colour and luminance*. S.P.I.E. Vol.1666, Human Visual Processing and Digital Display, pp 302 - 309.
2. Derrico J. B. and Buchsbaum G. 1991. *A computational model of spatiochromatic image coding in early vision*. Journal of Visual Communication and Image Representation, 2, 31-37.
3. Buchsbaum G. and Gottschalk A. 1983. Trichomacy, opponent colour coding and optimum colour information transmission in the retina. Proc. R. Soc. Lond. B220, pp. 80-113.
4. Field D. J. 1987. *Relation between the statistics of natural images and the response properties of cortical cells*. J. Opt. Soc. Am., Vol.4, No. 12, pp 2379-2394.
5. Field D. J. 1989. What the statistics of natural images tell us about visual coding. S.P.I.E. 1077, pp. 269-276.
6. Burton G. J. and Moorhead I. R. 1987. *Color and spatial structure in natural scenes*, Applied Optics vol. 26, pp. 157-170.
7. Atick J. J. 1992. Could information theory provide an ecological theory of sensory processing. Network 3, pp. 213-251.
8. Mullen K. T. 1985. Contrast sensitivity of human colour vision to red-green and blue-yellow chromatic gratings. J. Physiol. 359, 381-400.
9. Ingling Jr. C. R. and Martinez E. 1983. *The spatiochromatic signal of the r-g channel*. Colour Vision: Physiology and Psychophysics. ed. by J.D. Mollon & L.T. Sharpe, pp. 433-444.