Chapter 2

Methods

Technical background

For this exploration of the statistical properties of the natural environment we require a representative sample of natural scenes. This implies also a system for image acquisition (camera), image processing (computers), storage (disks, tapes, etc.) and display (monitors). We require for the system to provide a measure of radiance as a function of wavelength for each pixel of our images. Using this information we calculate the statistics of the environment and contrast them with some of the known features of the h.v.s. An ordinary colour (CCD) TV camera seems to be the most obvious choice except for the fact that it makes use of its own colour model. Colour models are 3D co-ordinate systems in where each colour is represented by a single point which facilitates the specification of colours in some generally accepted (standard) way. Most colour models in use today are either hardware-oriented or applications-oriented and their assumptions do not represent the world directly. Furthermore, any trichromatic sampling of this kind does not allow subsequent recoding into, for example, more than three channels or into channels of varying bandwidths. We wished to have the freedom to be able to model any possible *n-channel* sampling where *n* is less than 31. Another possible option is the use of photographic film and a digitizer. Here the process of calibration is more complicated because the output has to be compensated for the non-linearities of the photographic process. The main reason, however, for the choice of video rather than photographic technology was the greater ease of digitising the images of a video representation.

In order to provide a more direct (and free of erroneous assumptions) representation of the natural environment we employ a multispectral digital camera as described in the following sections.

2.1 The Defence Research Agency-funded multispectral camera

The gathering of all statistical properties of a large sample of natural visual images used in this work was made by a video camera capable of recording images through narrowband interference filters and storing the resultant images on a computer disk. The development of the camera was funded by the Defence Research Agency (DRA).

The DRA-camera system and controller device were partially constructed by Custom Camera Designs Ltd. (of Wells, Somerset) and completed, assembled, calibrated and operated by members of the Perceptual Systems Research Centre (PSRC) of the University of Bristol between 1990 and 1993. Originally, it was constructed to achieve the following aims:

a) To construct an electro-optical (e-o) system capable of representing the spectral distribution of light in terrain scenes.

b) To validate the e-o system with calibrated images.

2.2 Characteristics of the DRA-camera

The DRA-camera consists of an e-o mechanism built around a "Pasecon" tube, a camera control unit (CCU), a carousel slide changing filter mechanism, a portable 386-PC, a real-time video monitor and a battery power supply. The scheme of the DRA-camera is shown in Figure 2.1. The Pasecon tube was chosen because it has good linearity throughout the full spectral range and a low-noise output at low light inputs. Figure 2.2 illustrates the camera tube sensitivity between 400 and 800 nm (from manufacturer's data). The original Camera used to convert the video signal to digital form (analogue/digital conversion) for storage in the internal frame-store card and then reconvert to analogue form (digital/analogue conversion) for both display and as input to the Data Translation frame card. Given that the final digital/analogue device proved to be very prone to drift with temperature and since the required images were to be stored and processed in digital form, digital images were transferred directly from the camera after the first analogue/digital conversion. A card was designed and built to enable clean and noise-free images to be transferred digitally directly from the camera's first a/d converter into the frame-store on the Data Translation card.

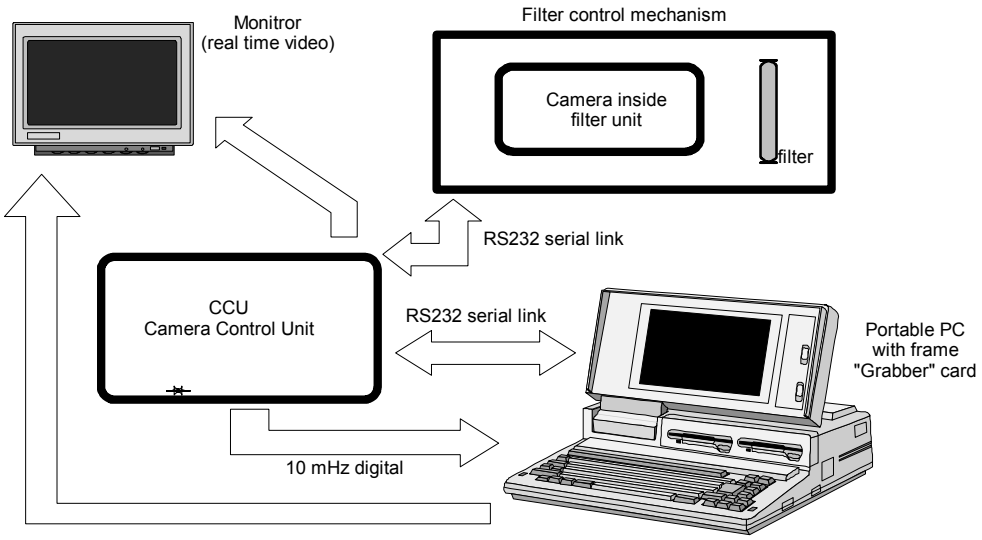

Figure 2.1: Scheme of the DRA-Camera.

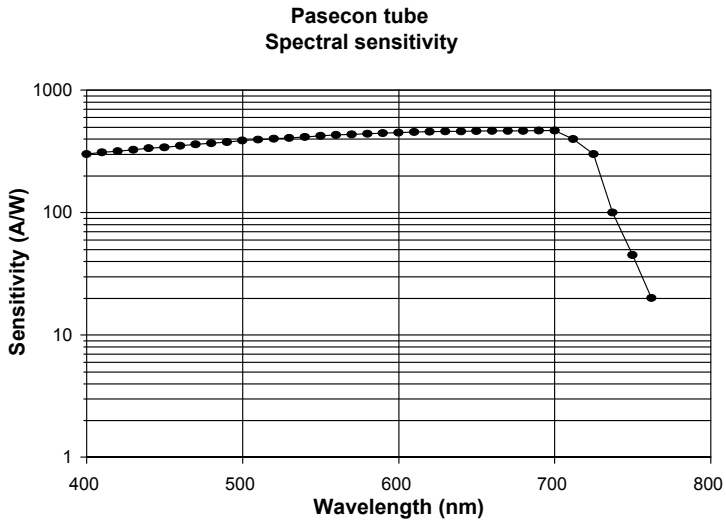

Figure 2.2: Camera tube sensitivity between 400 and 800 nm (from manufacturer's data).

The slide changing filter mechanism was added to the e-o mechanism to allow the DRA-camera to sequentially grab images through a set of 31 optical filters. These are chromatically narrowband in the range of 400 to 700 nm and with 10 nm spacing between their peaks. Figure 2.3 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.

A PC program (written in 'C' under MS-DOS 3.0) sends a value for the integration time (see below) to the DRA-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 narrowband 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 lenses. 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.

Figure 2.3: Transmittance of the filters employed.

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 it is possible to display each image 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 reflectance can be shown on the graph. Figure 2.4 illustrates this facility. This requires the images to be calibrated for reflectance. This was achieved in the manner described below.

Figure 2.4: Display window of the IBM-RS/6000 employed to process the dataset.

2.3 The DRA-camera calibration

The DRA-camera system was calibrated to write out image files along with relevant header information such as the integration time and lens aperture from which multi-spectral light-measurements can be reconstructed. Figure 2.5 shows a measurement of the system output versus time since system boot. Notice that the minimum time necessary to ensure a stable output was around 6000 seconds (1hour 36 min). All measurements described below were obtained several hours after the system boot.

The calibration strategy consisted basically of:

a) Finding for any given filter and aperture setting the number of integration frames required to achieve a reasonably large dynamic range.

b) Correcting the non-linear characteristics of the DRA-camera and the off-axis variation across the camera target.

c) Converting the grey level output of the DRA-camera into measurements of the spectral radiance (and spectral reflectance) across the target.

Figure 2.5: System output versus time since system boot.

(a) Choosing the right 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 exceeding the ceiling.

Once both the statistics (S) and ceiling (C) are defined the algorithm tries to automatically find the 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, an algorithm has been developed to accelerate the process. This algorithm can predict a suitable value for IT using the results of previous attempts, and check if the light levels change between individual trials. If that happens, the calibration process is halted and the problem reported.

(b) Corrections of non-linear characteristics of the DRA-camera

To correct the non-linear characteristics of the DRA-camera, a look up table (LUT) was produced as follows:

The DRA-camera was pointed to a blank piece of white paper (test card) the radiance of which was 0.002 WSr^{-1} nm^{-2} m^{-1} as measured with a TopCon Spectro-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 2.2 shows, the camera tube sensitivity decays drastically in the infra red (IR) part of the spectra ensuring that no IR radiation will affect the measurements in spite of any secondary peak the filter could have in this region. Figure 2.6 shows the apparatus set up.

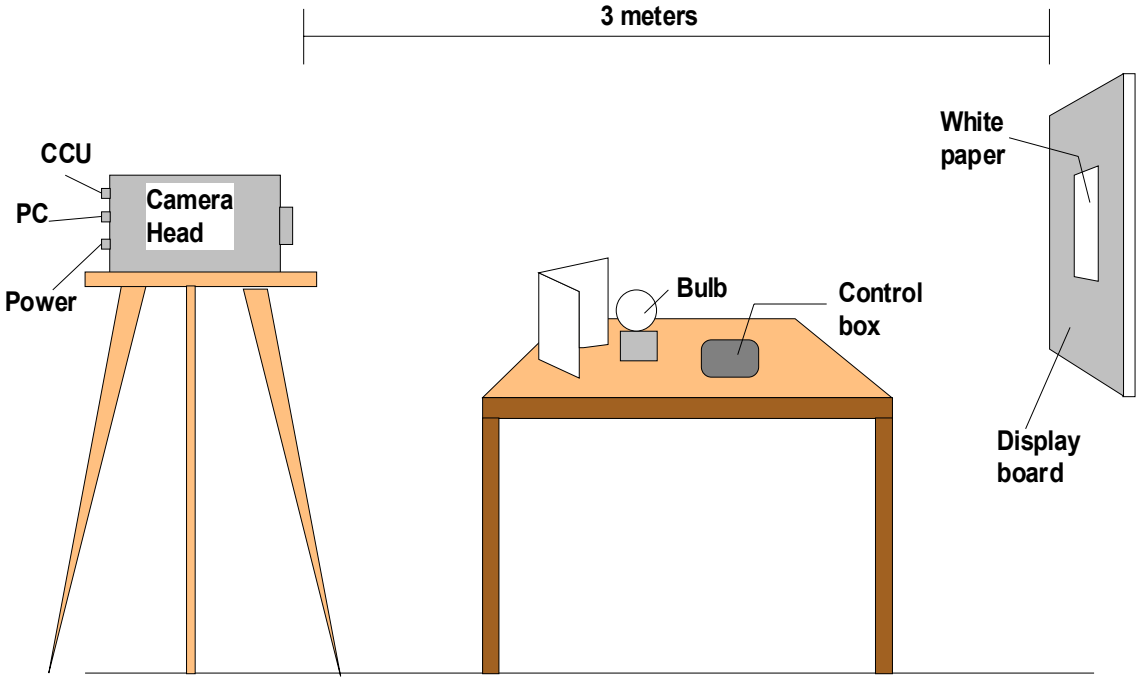

Figure 2.6: apparatus set up for calibration.

The system output within the central square region was estimated by varying the integration time from 1 to 200 frames for all 8 lens aperture settings. This provides a series of eight 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 16 exemplary curves (one for each of the above conditions) were selected and used to create the look-

up table.

The same procedure used for calibrating the centre of the scene was repeated for all nine regions of the scene. An additional set of 8 further regions lying on the perimeter of the scene was used to provide extra information about the off-axis output. As before, this data was stored in an off-axis correction file that was simultaneously used with the LUT to convert the DRA-camera output into light measurement.

(c) Conversions into spectral radiance

This was obtained by matching the output of the DRA-camera through each of its 31 filters to that of the TopCon SR1 when both systems were pointing to a standard Kodak grey card. Further measurements of the spectral radiance by the DRA-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 bigger 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 2.7(a) and 2.7(b) provide an example of this matching.

Figure 2.7(a): Spectral radiance obtained with the DRA-camera

Figure 2.7(b): Spectral radiance obtained with the TopCon SR1.

(d) Conversions into spectral reflectance

To convert the DRA-camera output into measurements of spectral reflectance, it is necessary to know the spectral characteristics of the light falling over the scene.

$$\textit{Scene Radiance} = \textit{Illumination} \times \textit{Reflectance}$$

This information can be obtained by placing an object of known reflectance (like the Kodak standard grey card) into the scene and measuring its radiance. The algorithm allows the introduction of the XY position of the grey card on the screen and performs the conversion automatically. This assumes that light falling over the grey card is an accurate sample of the light falling over the whole scene and that the reflections are approximately independent of the subtended angle (Rayleghian or diffuse reflections).

2.4 Scene grabbing

A dataset of 31 scenes was obtained for this analysis 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) Warming up of the DRA-camera systems and battery checking. This was usually done the day before to ensure that the system would be fully warmed during the recording. As shown on Figure 2.5 these temperature changes can alter the system

output.

b) Selection of the place and time of the day for the recording. This was done the day before (after checking the forecast information). To avoid changes in light because of changes in sun position, all pictures (except the ones taken inside the lab) were generally taken around noon. Several places were chosen depending on their facilities for moving the trolley with the equipment and the drift of people around. Several scenes were recorded near the Royal Fort House of the University of Bristol, others in the University of Bristol Botanical Gardens. Some practical problems that influenced our selection of the place and daytime will be discussed in the next section.

c) Positioning of the grey card in a visible place in the scene. Most of the pictures (29) contained no artificial objects with the exception of the Kodak grey card, and (in two cases) a tripod. The card was placed in all scenes for calibration and light checking purposes.

d) Recording of the spectral radiance of the grey card using the TopCon SR1. This was made to detect changes in the spectral characteristics of the light before and after every picture recording. The measurements were taken from approximately the same place as that of the DRA-camera and just before starting the proper recording of the scene.

e) Recording of the scene. Considering the relative long time (about 5 minutes) needed to grab each scene it was often necessary to wait until weather conditions were acceptable (see next section about practical problems) before starting.

f) Second spectral measurement of the grey card radiance. The procedure was similar to (d) and was made immediately after finishing the scene recording.

g) A photograph was taken of approximately the same scene to be used as a colour image reference.

h) Back in the laboratory the grabbed images were transferred into the IBM/6000 workstation for further analysis.

2.5 *Common practical problems and their solutions*

Most of the practical problems we faced were related to changes in the scene during the

recording. A different solution was formulated for each one.

a) Linear light changes. These were generally caused by changes in the sun position during the recording of the scene. These alterations never affected 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 its effects (see below).

b) Light fluctuation, mainly due to short-term variations in cloud cover. This generally affected one or two of the images in the scene. To detect them, it was necessary to compare the SR1 spectral measurements on the grey card with those of the DRA-camera in the same place. The most effective way to avoid these fluctuations was to wait until the sky was completely clear or completely overcast before grabbing the scene. Small errors were corrected with the "*reilluminate*" program (see below).

c) Small movements of objects (such as tree branches) in the scene during imaging. These were due to wind and their negative effects were more intense for objects at a short distance. Negative effects were minimised by grabbing medium and long distance pictures but often it was necessary to wait for these movements to cease. All the naturally lit short distance scenes were taken in the glass houses of the University of Bristol Botanical Gardens in order to avoid this problem.

2.6 *Corrections made by software*

The "*reilluminate*" program was created to correct small light fluctuation and linear light changes of particular scenes of the dataset. It uses the spectral reflectance image of the scene along with one of the SR1 measurements (the program allows one to choose between any of them) taken on the grey card. It "reilluminates" the scene with the light spectrum falling over the grey card so that the radiance measured with the DRA-camera on the grey card matches the SR1 measurement. Appendix A gives a detailed explanation of its fundamentals along with the DRA-camera mathematical theory.

Modifications of the data were avoided as much as possible and the use of this kind of correction was limited to five particular cases of the dataset. To decide whether it was necessary to use this program or not, the following points were considered:

a) The fit between SR1 measurements taken on the grey card before and after the recording of each scene. A special algorithm (Press *et al.* 1991) was applied to compare both spectral measurements.

b) Comparison between the spectral radiance of the grey card as measured by the SR1 and the spectrum obtained by the DRA-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.

After correcting a scene, the comparison described in (b) was performed again in order to evaluate the effects of the "*reilluminate*" program. However, all further analyses in terms of Fourier amplitude were performed separately on the corrected dataset and on the original one, and there does not appear to be a marked difference between both sets of results.

2.7 *Some statistics of the dataset*

Although there is no formal agreement about what is considered a representative sampling of the visual environment, we ensure that (within our limitations) some of the most common natural objects are represented in our dataset. These include plants with different shapes, textures and colours, flowers (mostly in bright red colours), trunks, branches, grass, yellow leaves, trees, bushes, rocks and sky. Some of these objects were artificially arranged in the laboratory trying to make this arrangement to look as casual as possible. Others are just images of the British countryside (in which buildings or other human artifacts were avoided) or of gardens (either taken in the Botanical Gardens or in the University of Bristol Royal Fort Gardens). Our general aim was to generate a set of images 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, we felt that the inclusion of vegetation seen from different distances might satisfy this requirement in part. Later on we shall discuss the extent to which our images (obtained in typical diffuse Northern Temperate illumination) may differ from images obtained under more directional illumination. Practical constraints limited our choice of possible images. Bright parts such as sky result in large saturated portions of the image and were deliberately avoided. Only four of our scenes contain regions of

sky. Moving objects such as clouds are also avoided. Very strong shadows produced by solar illumination move during the period of image acquisition and are consequently avoided. Water (i.e. lakes, ponds, etc.) is also avoided because both of the bright reflections and the movement. All the scenes were recorded in Autumn and Winter (between October 1993 and January 1994) and this can be considered as a limitation in the range of possible natural colours and textures.

Of the 31 scenes recorded, 29 contained no other artificial object than the grey card, and in two cases a tripod to attach it. These were the ones used in our analysis. The remaining two included artificial objects and were used to test the quality of the lenses (as is described in Chapter 5).

Different classifications of the dataset can be performed depending on the characteristic considered:

a) In terms of lighting, nineteen scenes were illuminated by the sun (these are called *sun-illuminated*), directly or through cloud cover, nine were illuminated by incandescent lighting and one by fluorescent lighting (these are called *artificially-illuminated*).

b) In terms of distance to the main objects in the scene, four *long distance* scenes were taken. This means that the main objects were in the range 0.1-4.0 km. The rest of the dataset includes objects in the range 0.5-50.0 m (called *short-distance* scenes).

c) In terms of the corrections performed, the dataset can be classified in three different categories:

Corrected scenes (5 in total). "*Reilluminate*" was used to correct light fluctuation on several filtered images.
Slightly corrected (14 in total). "*Reilluminate*" was used to produce the adjustment of the DRA-camera measurement to that of the SR1 on the grey card.
Without corrections (10 in total) These were also called *selected scenes* in later chapters.

Several software tools were developed in 'C' to run on scenes on the RS/6000. The most frequently used in this analysis were:

"DO_LMS": This uses as input the 31 narrowband images that constitutes a given scene. The program first converts the value of each pixel into radiance or reflectance (as described in section 2.3) according to a selection by the user. A special option allows the user to implement the corrections described in section 2.6 (i.e. the "reilluminate" algorithm). Following this, the program weights each of the 31 monochromatic images and adds them together to generate each of the L, M, S cone sensitivity images. The weighting is produced according to the Smith and Pokorny (1975) cone sensitivity functions (see Figure 1.3) which are normalised to have a maximum sensitivity value equal to one. The output of this program is a set of three floating-point images, each of them corresponding to the L, M and S cone sensitivity (see Figure 2.8).

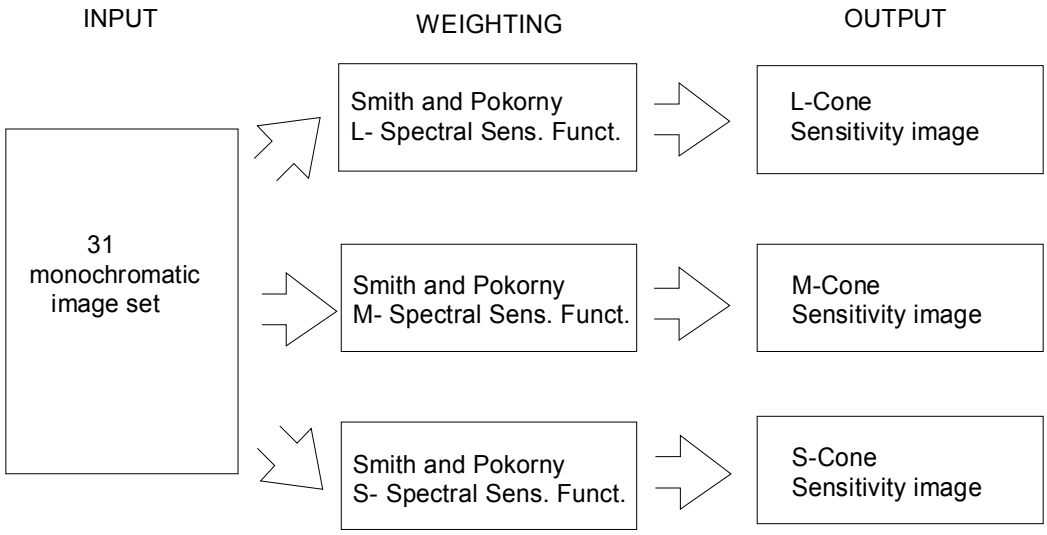

Figure 2.8: Scheme of the "DO-LMS" algorithm.

"FBANDS": This produces the Fourier amplitude spectra of an image and displays it as a 2-dimensional intensity function. For example, if "FBANDS" is applied to a floating-point image A, the output will be another floating-point image where brightness is proportional to the amplitude spectrum of A. In this representation, the distance from the centre of the image is proportional to the spatial frequency value and the angle from the horizontal represents the orientation (see Appendix B). Given that the dynamic range of Fourier spectra is much higher than the display device is able to faithfully

reproduce (255 grey levels), the centre of the plot has been removed. This corresponds to the highest value and is analysed separately in our analysis (it also represents the average luminance of the image A). Figure 2.9 shows a 3-dimensional diagram of the Fourier amplitude spectra of an image.

Figure 2.9: 3-dimensional diagram of the Fourier amplitude spectra of an image. Taken from Field (1987).

The program also divides the Fourier space into 8 concentric annuli (bands), each of them representing a given range of the spatial frequency spectrum. The total Fourier amplitude and its average value within the band (averaged across orientation) are measured into these bands and the results presented in a table. It is possible to choose between two different sets of concentric bands, the *"linear"* and the "*logarithmic"*. Tables 2.1 and 2.2 show the range (in pixels measured from the centre of the screen and their equivalent in cycles/deg) covered by each band in both cases. Figure 2.10 shows a scheme of this annuli-shaped division. These *linear* and *logarithmic* band spacing produce different results when applied to our dataset (see later in Chapter 5 for a more extensive discussion of this).

We choose the number of logarithmic bands to be equal to 8 because this is a convenient number for dividing the 256 x 256 pixels screen. In this case the minimum bandwidth corresponds to one pixel (see Table 2.2). To keep the symmetry between both versions, the number of bands is equal to 8 in the linear case too.

All the software tools were developed by the PSRC between 1991-1993.

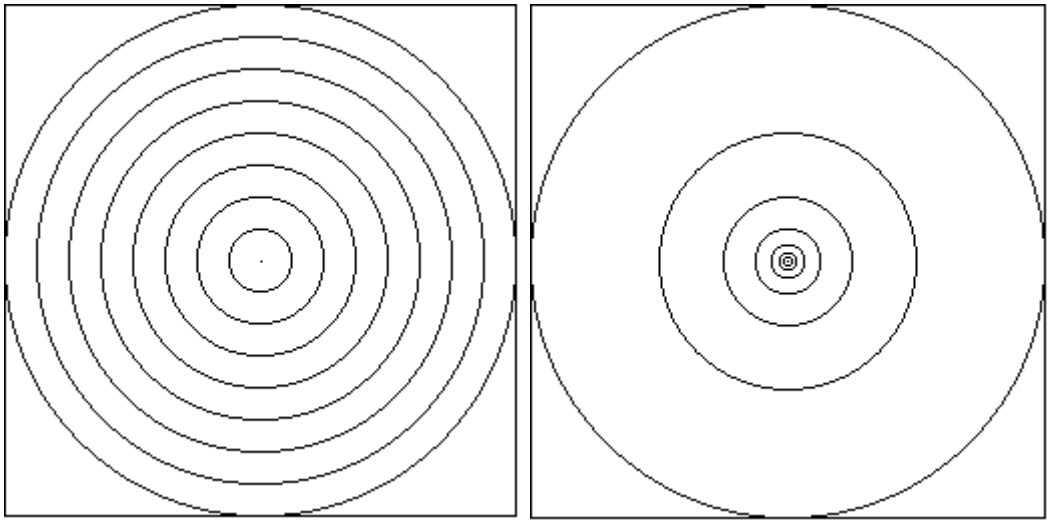

Figure 2.10: Annuli-shaped division of the Fourier space. *Linear* (left) and *logarithmic* (right).

2.9 *Validity tests*

Several validity and repeatability tests were performed to check whether the system (DRA-camera plus software tools) was working properly.

The DRA-camera output was compared to that of the TopCon Spectro-radiometer for several coloured objects to ensure that radiance/reflectance measurements were reliable.

Synthetic images of sine gratings with spatial frequencies of 0.008, 0.1 and 0.5 cycles/pixel were constructed using a special algorithm designed for this purpose. Following this, they were processed using FBANDS. The amplitude spectrum image showed the peaks corresponding to the spatial frequencies mentioned above. Logarithmic bands 2, 6 and 8 respectively showed amplitude values 1000 times greater than the rest.

A natural image was filtered using a spatial frequency low-pass filter and then processed using FBANDS. The results showed the corresponding shift in the Fourier amplitude spectra towards the low spatial frequencies. A similar analysis was performed using high-pass filters. The data obtained using the "*reilluminate*" algorithm was compared with the results obtained without it and the difference proved to be non significant for this analysis.

Further lab testing of the camera lenses were performed to see whether there is any defocus on the blue side of the spectrum that could affect our analysis. These results are more intensively discussed in Chapter 5.

Table 2.1: Linear Bands				
Band	**Range**		**Centre**	**Nb**
	pixels	***cycles/deg***	***cycles/deg***	***Pixels***
1	1 - 16	0.03 - 1.10	0.57	888
2	17 - 32	1.10 - 2.20	1.65	2516
3	33 - 48	2.20 - 3.30	2.75	4116
4	49 - 64	3.30 - 4.40	3.85	5716
5	65 - 80	4.40 - 5.50	4.95	7352
6	81 - 96	5.50 - 6.60	6.05	8924
7	97 - 112	6.60 - 7.70	7.15	10564
8	113 - 128	7.70 - 8.80	8.25	12110

Table 2.2: Logarithmic Bands				
Band	**Range**		**Centre**	***Nb***
	pixels	***cycles/deg***	***cycles/deg***	***Pixels***
1	1 - 1	0.03 - 0.07	0.05	8
2	2 - 2	0.07 - 0.14	0.10	16
3	3 - 4	0.14 - 0.28	0.21	44
4	5 - 8	0.28 - 0.55	0.41	180
5	9 - 16	0.55 - 1.10	0.83	640
6	17 - 32	1.10 - 2.20	1.65	2516
7	33 - 64	2.20 - 4.40	3.30	9832
8	65 - 128	4.40 - 8.80	6.60	38950

The tables above show the range covered by each band of the FBANDS algorithm in the 8 bit representation of the Fourier space and their equivalent in cycles/degree. The band centre is the value used to represent a given band in all the plots of this work. Nb represents the total number of pixels included by each band. All conversions were done using the field of view of the DRA-camera.