Using digital photography to study animal coloration.

Authors: Martin Stevens[1], C. Alejandro Párraga[2], Innes C. Cuthill[1], Julian C. Partridge[1] & Tom Troscianko[2].

Institutions: [1]School of Biological Sciences, University of Bristol, Woodland Road, BS8 1UG. UK. [2]Department of Experimental Psychology, University of Bristol, Woodland Road, Bristol BS8 1TN, UK.

Running Title: Using cameras to study animal coloration.

ABSTRACT

In understanding how visual signals function, quantifying the components of those patterns is vital. With the ever-increasing power and availability of digital photography, many studies are utilising this technique to study the content of animal colour signals. Digital photography has many advantages over other techniques, such as spectrometry, for measuring chromatic information, particularly in terms of the speed of data acquisition and its relatively cheap cost. Not only do digital photographs provide a method of quantifying the chromatic and achromatic content of spatially complex markings, but they can also be incorporated into powerful models of animal vision. Unfortunately, many studies utilising digital photography appear to be unaware of several crucial issues involved in the acquisition of images, notably the non-linearity of many cameras' responses to light intensity, and biases in a camera's processing of the images towards particular wavebands. In this paper, we set out step-by-step guidelines for the use of digital photography to obtain accurate data, either independent of any particular visual system (such as reflection values), or for particular models of non-human visual processing (such as that of a passerine bird). These guidelines include how to: linearise the camera's response to changes in light intensity; equalise the different colour channels to obtain reflectance information or how to produce a mapping from camera colour space to that of another colour space (such as photon catches for the cone types of a specific animal species).

Key Words: animal coloration - camera calibration - colour - colour measurement – digital cameras - imaging - photography - radiance - reflection - signals.

Introduction

Investigations into the adaptive functions of animal coloration are widespread in behavioural and evolutionary biology. Probably because humans are 'visual animals' themselves, studies of colour dominate functional and evolutionary investigations of camouflage, aposematism, mimicry, and both sexual and social signalling. However, with advances in our knowledge of how colour vision functions and varies across species, it becomes increasingly important to find means of quantifying the spatial and chromatic properties of visual signals as they are perceived by other animals or, at the very least, in a manner independent of human perception. This is non-trivial because colour is not a physical property, but rather a function of the nervous system of the animal perceiving the object (Newton, 1718: 'For the rays, to speak properly, are not coloured'; Endler, 1990; Bennett, Cuthill & Norris, 1994). One way to produce an objective measure of the properties of a colour signal is to measure surface reflectance using spectrophotometry, which provides precise information on the intensity distribution of wavelengths reflected (e.g. Endler, 1990; Zuk & Decruyenaere, 1994; Cuthill *et al.*, 1999; Gerald *et al*., 2001; Endler & Mielke, 2005). Reflectance data can also be combined with information on the illuminant and the photoreceptor sensitivities of the receiver (and, if available, neural processing) to model the colours perceived by non-human animals (e.g. Kelber, Vorobyev & Osorio, 2003; Endler & Meilke, 2005). However, conventional spectrometers provide only point samples, so to characterise adequately the colour of a heterogeneous object requires multiple samples across an appropriately designed sampling array, such as multiple transects or pre-specified regions (e.g. Cuthill *et al.*, 1999; Endler & Mielke, 2005). This not only has a cost in terms of sampling time, but the information about spatial relationships between colours then needs to be reconstructed from the geometry of the sampling

array (e.g. Endler, 1984) and the spatial resolution is generally crude. Spectrometry also usually requires a static subject, either because of the need to sample an array or because the measuring probe often needs to be close to or touching the colour patch, a particular problem in the field or with delicate museum specimens. Focussing optics can obviate the need for contact with the animal or plant and offer a degree of 'remote sensing' (e.g. Marshall *et al*., 2003; Sumner, Arrese & Partridge, 2005), but this approach is rare.

An alternative to spectrometry is photography, which has a long history of use in studies of animal coloration (A. Thayer, 1896; G. Thayer, 1909; Cott, 1940; Tinbergen, 1974; Pietrewicz & Kamil, 1979) but is becoming increasingly used because of the flexibility and apparent precision that digital imaging provides. Colour change in the common surgeonfish (Goda & Fujii, 1998), markings in a population of Mediterranean monk seals (Samaranch & Gonzalez, 2000), egg crypsis in blackbirds (Westmoreland & Kiltie, 1996), the role of ultraviolet (UV) reflective markings and sexual selection in guppies (Kodric-Brown & Johnson, 2002), and the functions of primate colour patterns (Gerald *et al.*, 2001), are a few recent examples. Digital photography bears many advantages over spectrometry, particularly in the ability to utilise powerful and complex image processing algorithms to analyse entire spatial patterns, without the need to reconstruct topography from point samples. More obviously, photographing specimens is relatively quick, allowing rapid collection of large quantities of data, from unrestrained targets and with minimal equipment. Imaging programs can be used to obtain various forms of data, including colour patch size and distribution measures, diverse 'brightness' and colour metrics, or broadband reflection values (such as in the long-, medium-, and short wavebands). Video imaging can provide temporal information too. Digital technology also has the potential for manipulating stimuli for use in experiments, the most impressive

examples being in animations within video playback experiments (e.g. Rosenthal & Evans, 1998; Künzler & Bakker, 1998), although there are problems with these methods that need to be understood (D'Eath, 1998; Fleishman *et al*., 1998; Cuthill *et al*., 2000; Fleishman & Endler, 2000).

Digital photography is increasingly incorporated into many studies of animal coloration due to its perceived suitability for objectively quantifying colour and colour patterns. However, many studies appear to be unaware of the complex image processing algorithms incorporated into many digital cameras, and make a series of assumptions about the data acquired that are rarely met. The images recorded by a camera are not only dependent upon the characteristics of the object photographed, the ambient light, and its geometry, but also upon the characteristics of the camera (Barnard & Funt, 2002; Westland & Ripamonti, 2004). Therefore, the properties of colour images are device-dependent, and images of the same natural scene will vary when taken with different cameras, since the spectral sensitivity of the sensors and firmware/software in different cameras varies (Hong, Lou & Rhodes, 2001; Yin & Cooperstock, 2004). Finally, the images are frequently modified in inappropriate ways (e.g. through 'lossy' image compression; for a glossary of some technical terms see Appendix 1) and 'off-the-shelf' colour metrics applied without consideration of the assumptions behind them. At best, most current applications of digital photography to studies of animal coloration fail to utilise the full potential of the technology; more commonly, they yield data that are qualitative at best and uninterpretable at worst. This paper aims to provide an accessible guide to addressing these problems. We assume the reader has two possible goals: to reconstruct the reflectance spectrum of the object (maybe just in broad terms such as the relative amounts of long-, medium- and short-wave light; although we will also consider something more ambitious), or to model the object's colour as perceived by a non-human animal. As we are considering

applications of the accessible and affordable technology of conventional digital colour cameras, we are primarily focussed on the human-visible spectrum of ca. 400 – 700 nm, but we also consider ultraviolet imaging and combining this information with that from a standard camera. Our examples come from an investigation of colour patterns on lepidopteran wings, and how these might be viewed by avian predators. This is a challenging problem (birds are potentially tetrachomatic and have an ultraviolet-sensitive cone type; reviewed by Cuthill *et al.*, 2000), yet it is both tractable and informative, because much of the avian colour world overlaps with ours and birds are the focal organisms in many studies of animal coloration (whether their sexual signals, or the defensive coloration of their prey).

Conceptual background

The light coming from an object, its radiance spectrum, is a continuous distribution of different intensities at different wavelengths. No animal eye, or camera, quantifies the entire radiance spectrum at a given point, but instead estimates the intensity of light in a (very) few broad wavebands. Humans and many other primates use just three samples, corresponding to the longwave (LW or 'red'), mediumwave (MW or 'green') and shortwave (SW or 'blue') cone types in the retina (Figure 1A); bees and most other insects also use three samples, but in the ultraviolet (UV), SW and MW wavebands; birds and some reptiles, fish and butterflies use four samples (UV, SW, MW and LW; Figure 1B). A corollary of colour vision based on such few, broadband, spectral samples is that the colour appearance of an object can be matched, perfectly, by an appropriate mixture of narrow waveband lights ('primary colours') that differentially stimulate the photoreceptors. Three primary colours (e.g. red, green and blue in video display monitors) are required for colour matching by normally-sighted

humans. All that is required is that the mix of primary colours stimulates the photoreceptors in the same way as the radiance spectrum of the real object (without actually having to mimic the radiance spectrum *per se*). The additive mixing of three primaries is the basis of all video and cinematographic colour reproduction, and colour specification in terms of the amounts of these primaries, the so-called tristimulus values, lies at the base of most human colour science (Wyscecki & Stiles, 1982; Mollon, 1999; Westland & Ripamonti, 2004). That said, RGB values from a camera are not standardised tristimulus values and so, although they are easily obtained with packages such as Paintshop Pro (Jasc Software, Minneapolis, USA) or Photoshop (Adobe Systems Inc., San Jose, USA), simply knowing the RGB values for a point in a photograph is not sufficient to specify the colour of the corresponding point in the real object.

An over-riding principle to consider when using digital cameras for scientific purposes is that most digital cameras are designed to produce images that look good, not to record reality. So, just as Kodachrome and Fujichrome produce differing colour tones in 'analogue' film-based cameras, each film type having its own advocates for preferred colour rendition, the same is true of digital cameras. The values of R, G and B that are output from a camera need not be linearly related to the light intensity in these three wavebands. In technical and high-specification cameras they are, and the sensors themselves (the CCDs or Charge Coupled Devices) generally have linear outputs. In contrast most cameras designed for non-analytical use have non-linear responses (Cardei, Funt & Barnard, 1999; Lauziére, Gingras & Ferrie, 1999; Cardei & Funt, 2000; Barnard & Funt, 2002; Martinez-Verdú, Pujol & Capilla, 2002; Westland & Ripamonti, 2004). This is a function of post-CCD processing to enhance image quality, given the likely cross-section of printers, monitors and televisions that will be used to view the photographs (these devices themselves having diverse, designed-in,

non-linearities; Westland & Ripamonti, 2004). Most digital images will display well on most monitors as the two non-linearities approximately cancel each other out. The first step in analysing digital images is therefore to linearise the RGB values.

Even with RGB values that relate linearly to R, G and B light intensity, there is no single standard for what constitutes 'red', 'green' and 'blue' wavebands; nor need there be, as different triplets of primary colours can (and, historically, have been) used in experiments to determine which ratios of primaries match a given human-perceptible colour (Mollon, 1999; Westland & Ripamonti, 2004). The spectral sensitivities of the sensors in a digital camera need not, and usually do not, match those of human visual pigments, as was the case with the Nikon 5700 Coolpix camera primarily used in this study (Figure 1C). The RGB values in images from a given camera are specific to that camera. Indeed, the values are not necessarily even specific to a particular make and model, but rather specific to an individual camera, because of inherent variability in CCDs at the manufacturing stage (Figure 2). One can, however, map the camera RGB values to a camera-independent, human colour space (and, under some circumstances, that of another animal) given the appropriate mapping information. Therefore the mapping, through mathematical transformation, of the camera-specific RGB values to camera-independent RGB (or other tristimulus representation) is the second crucial step in obtaining useful data from a digital image. Furthermore, and often as part of the transformation step, it will usually be desirable to 'remove' variation due to the illuminating light. The camera measures R, G and B radiance, which is the product of the reflectance of the object and the three-dimensional radiance spectrum illuminating the object (often approximated by the irradiance spectrum of the illuminant). The situation is rather more complex underwater, where the medium itself alters the radiance spectrum (Lythgoe, 1979) by wavelength-dependent attenuation. However, an object does not change colour

(much) when viewed under a blue sky, grey cloud, or in forest shade, even though the radiance spectra coming from it changes considerably. This phenomenon of 'colour constancy', whereby the visual system is largely able to discount changes in the illuminant and recover an object's reflectance spectrum, is still not fully understood (Hurlbert, 1999), but equivalent steps must be taken with digital images if it is object properties that are of interest rather than the radiance itself. Many digital cameras allow approximations of colour constancy (white-point balancing) at the point of image acquisition, for instance by selecting illuminant conditions such as skylight, cloudy and tungsten. These settings are, however, an approximation, and in practice their effects need to be eliminated, as the effect of the illuminant itself needs to be 'removed'. Removing the effect of the light source characteristics can thus be coupled to eliminating any biases inherent in the camera's image processing (such as an over-representation of some wavelengths/bands to modify the appearance of the photograph; Cardei *et al.*, 1999; Lauziére *et al*., 1999; Finlayson & Tian, 1999; Martinez-Verdú *et al.*, 2002). This is essential if accurate data representing the inherent spectral reflection characteristics of an animal's colour are to be obtained.

Many studies have used cameras to investigate animal colour patterns, but most fail to test their digital cameras to determine if all of the above assumptions are met and/or if the analysis yields reliable data (e.g. Frischknecht, 1993; Villafuerte & Negro, 1998; Wedekind *et al.*, 1998; Gerald *et al*., 2001; Kodric-Brown & Johnson, 2002; Bortolotti, Fernie & Smits, 2003; Cooper & Hosey, 2003), but see Losey (2003) for a rare exception.

We approach these problems in the sequence that a scientist would have to address them if interested in applying digital photography to a research question about biological coloration. This paper focuses on obtaining data corresponding to inherent animal colouration, such as reflection data, and of obtaining data corresponding to a

given receiver's visual system. Either of these data types may be more suitable depending upon the research question. Reflection data does not assume specific environmental conditions or a particular visual system viewing the object, and so data can be compared across different specimens easily, even when measured in different places. The lack of assumptions about the receiver's visual system, such as cone types, distributions, abundances, sensitivities, opponency mechanisms, and so on, means the data 'stand alone' and can be analysed as an inherent property of the animal or an object propagating the signal. This is useful if a researcher simply wishes to know if, for example, individual 'a' has more longwave reflection than individual 'b'. Removing illumination information coincides with evidence that many animals possess colour constancy. Conversely, simply taking reflection into account could be misleading if what one really wants to know is how a signal is viewed by a receiver. For example if an individual possesses a marking high in reflection of a specific waveband, but the environment lacks light in that part of the spectrum or the receiver is insensitive to that waveband, the region of high spectral reflection will be unimportant as a signal. Therefore, it is often necessary to include the ambient light characteristics and, if known, information concerning the receiver's visual system. However, calculated differences in quantal catches of various cone types (for example) between the different conditions do not necessarily lead to differences in perception of the signal, if colour constancy mechanisms exist. Furthermore, if reflection information is obtained, this may be converted into a visual system specific measure, either by mapping techniques, as discussed in this paper, or by calculations with illuminant spectra and cone sensitivities. Therefore, whilst this paper deals with both types of measurements, we focus more on the task of obtaining information about inherent properties of animal colouration.

We assume that images are stored to a precision of 8 bits in each colour channel, such that intensity is on a scale of 0 to 255; such 'true colour' images (2^8 cubed, or >16 million colours) are the current norm. Whilst some studies have used conventional (non-digital) cameras in studying animal coloration, we would advise against doing so. Whilst conventional film can be linearised, the corrections required from one batch of film to the next are likely to differ, even from the same manufacturer. Film processing techniques such as scanning to digitise the images, are also likely to introduce considerable spatial and chromatic artefacts, which need to be removed/prevented before analysis.

Choosing a camera

We have mentioned the non-linear response of many digital cameras and whilst we show, below, how linearisation can be accomplished, non-linearity is better avoided. Other than this, essential features to look for are (see also Table 1):

(a) The ability to disable automatic 'white-point balancing'. This is a software feature built into most cameras to achieve a more natural colour balance under different lighting conditions. The brightest pixel in any image is set to 256 for R, G and B; that is, assumed to be white. Obviously, for technical applications where the object to be photographed has no white regions, this would produce data in which the RGB values are inappropriately weighted.

(b) A high resolution. The resolution of a digital image is generally limited by the sensing array, rather than the modulation transfer function of the lens. Essentially, the number of pixels the array contains determines resolution, with higher resolution cameras able to resolve smaller colour patches allowing more detail to be measured, or the same amount of relative detail measured from a further distance from the

subject. Also important is the Nyquist frequency (half that of the highest waveform), which is the highest spatial frequency where the camera can still accurately record image spatial detail; spatial patterning above this frequency results in aliasing, which could be a problem for patterns with a very high level of spatial detail (Efford, 2000). Each pixel on a digital camera sensor contains a light sensitive photodiode, measuring the intensity of light over a broadband spectrum. A colour filter array (CFA) is positioned on top of the sensor to filter the red, green and blue components of light, leaving each pixel sensitive to one waveband of light alone. Commonly there is a mosaic of pixels, with twice as many green sensitive ones as red or blue. The two missing colour values for each individual pixel are estimated based on the values of neighbouring pixels, via so-called demosaicing algorithms, including Bayer interpolation. It is not just the number of pixels a camera produces (its geometrical accuracy) that matters, but also the quality of each pixel. Some cameras are becoming available which have 'foveon sensors', with three photodetectors per pixel, and can thus create increased colour accuracy by avoiding artefacts resulting from interpolation algorithms. However, due to the power of the latest interpolation software, colour artefacts are usually minor, especially as the number of pixels increases, and foveon sensors may have relatively low light sensitivity. Higher quality sensors have a greater dynamic range, which can be passed on to the images, and some cameras are now being produced with two photodiodes per pixel, one of which is highly sensitive to low light levels, the other of which is less sensitive and is used to estimate higher light levels without becoming saturated. A distinction should also be made between the number of overall pixels and the number of effective pixels. A conventional 5 megapixel camera actually may output 2560x1920 pixel images (4,915,200 pixels), since some of the pixels in the camera are used for various measurements in image processing (dark current measurements for instance).

(c) The ability to store images as uncompressed TIFF (Tagged Image File Format) or RAW files. Some mid-range cameras allow storage as RAW files, others do not but often allow images to be saved as Tiff files. This is something to determine before purchasing a camera. Other file types, JPEG's (Joint Photographic Experts Group) in particular, are unsuitable since information is lost in the compression process. JPEG compression is of the 'lossy' type, which changes the data coming from the CCD array, and where the lost information cannot be recovered. This is often undetectable to the human eye, but introduces both spatial and chromatic artefacts in the underlying image data, particularly if the level of compression is high (see Figures. 3 & 4 for two simple illustrations). JPEGs compress both the colour and spatial information, with the spatial detail sorted into fine and coarse detail. Fine detail is discarded first, since this is what we are less sensitive to. Gerald *et al.* (2001), for example, used digital images to investigate the scrota of adult vervet monkeys *Cercopithecus aethiops sabaeus*. They saved the images as JPEG files, but since the level of compression of the files is not stated, it is impossible to assess the degree of error introduced. Camera manuals may state the level of compression used on different settings, and image software should also state the level of compression used when saving JPEG files. However, even if the level of compression is known, the introduction of artefacts will be unpredictable and so JPEG files should be avoided. Lossy compression is different from some other types of compression, such as those involved with 'zipping' file types, where all the compressed information can be recovered. Uncompressed TIFF files are loss-less, but they can be compressed in either lossy or loss-less ways, and, like JPEGs, TIFFs can be modified before being saved other ways if the necessary camera functions are not turned off (such as white-point balancing). For most cameras, a given pixel on a CCD array has only one sensor type (R, G or B), so interpolation is required to estimate the two unknown colour values of a given pixel.

Both JPEGs and TIFF files undergo interpolation at the stage of image capture by the camera's internal firmware, which cannot be turned off, and the method is usually opaque to the user. Some cameras have the capacity to store RAW images. RAW files are those that are the direct product of the CCD array, and, unlike TIFFs or JPEGs which are nearly always 8-bit, RAW files are usually 12- or 16-bit. This means they can display a wider variety of colours and are generally linear since most CCDs are linear, and undergo none of the processing potentially affecting other file types. The RAW files from the camera in our study occupy approximately half of the memory of an uncompressed TIFF file because even though the TIFF file only retains 8-bits of information, it occupies twice the storage space because it has three 8-bit colour channels, as opposed to one 12-bit RAW channel per CCD pixel. However, before being useable as an image, RAW files must also go through interpolation steps in the computer software into which the files are read. Thumbnails of unprocessed RAW files in RGB format can be read into some software, but these are relatively useless, being only 160x120 pixels in resolution, compared to 2560x1920 for the processed images. The conversion to another file type can proceed with no modification, just as would be the case if taking photos directly as uncompressed TIFF images. One problem with RAW files is that they can differ between manufacturers and even between camera models, and so special software and/or 'plug-ins' may be needed, or the software provided by the manufacturer must be used, to convert the images to other file formats. Unfortunately, the interpolation process is rarely revealed by the manufacturer, and may introduce non-linearities into the file. It is possible to write custom programmes to read RAW files into software programmes and this has the advantage that the user can then either use the RAW data directly or decide exactly what method should be used to interpolate the RAW file into a Tiff file. Once our RAW files had been processed by software supplied by the manufacturer they had

almost identical properties to the uncompressed TIFF files (the introduction of non-linearities could be due to the software processing or a non-linear CCD). Some imaging software should allow the RAW files to be processed into TIFFs without introducing non-linearities. RAW files can be converted into 16-bit TIFF files, which show higher accuracy than 8-bit TIFFs and may highlight extra detail. These 16-bit file types occupy approximately 30 Mb, so considerable storage space is needed to keep a large number of these files. However, relatively more unprocessed RAW files can be stored than TIFFs on a memory card.

(d) The capacity for manual exposure control or, at the very least, aperture priority exposure. The calibration curve may vary with different aperture settings and focus distances so, to avoid the need for a large number of separate calibration estimates, it is more convenient to fix the aperture at which photographs are taken and work at constrained distances. If the aperture value is increased, more light from the edge of the lens is allowed through, and these rays usually do not converge on the same point as those rays coming through the centre of the lens (spherical aberration). This is especially true for colours near the edges of the human visible spectrum. By keeping the aperture constant and as small as possible (large F-numbers) this problem is unlikely to be significant.

(e) The ability to take a remote shutter release cable (manual or electronic) to facilitate photography at long integration times (slow shutter speeds) when light levels are low.

(f) Known metering characteristics. Many cameras have multiple options for light metering, such that the exposure is set dependent upon average intensity across the entire field imaged, only the intensity at the central spot, or one or more weighted intermediates. Knowing which area of the field in view determines exposure facilitates image composition.

(g) Optical zoom can be useful, particularly if the level of enlargement can be fixed manually, so it can be reproduced exactly, if needed, each time the camera is turned on. Digital zoom is of no value because it is merely equivalent to post-image-capture enlargement and so does not change the data content of the area of interest.

(h) Good quality optics. Hong *et al.* (2001) noted that in some camera lenses, light is not uniformly transmitted across its area, with the centre of the lens transmitting more light. This would result in the pixels in the centre of the image being over-represented in terms of intensity. For example, Losey (2003) found that the edges of images were slightly darker. Another problem with lenses is chromatic aberration, in which light of different wavelengths is brought to a focus in a different focal plane, thus blurring some colours in the image. This can be caused by the camera lens not focussing different wavelengths of light onto the same plane (longitudinal chromatic aberration), or by the lens magnifying different wavelengths differently (lateral chromatic aberration). Párraga, Troscianko & Tolhurst (2002) tested camera Nikon lenses of the type in our camera, by taking images in different parts of the spectrum through narrowband spectral filters and verified that the optimal focus settings did not vary significantly, meaning that the lenses did not suffer from this defect. Narrow bandpass filters selectively filter light of specific narrow wavebands, for example from 400 to 410 nm. Using a set of these filters enables images to be obtained where the only light being captured is in a specific waveband. Other lenses may not be as good, especially if they have a bigger optical zoom range. In some situations, a good macro lens is also highly desirable, since this allows close up images of complex patterns to be obtained. Without a macro lens it may not be possible to move the camera close enough to resolve complex patterns. Some cameras even come with a ‘super’ macro lens, such as the Fujifilm FinePix S7000, which allows photographs to be taken up to a centimetre from the object.

(i) The capacity to take memory cards of high capacity. TIFF files are very large (ca. 15 Mb for an image 2560 by 1920 pixels), so that a 512Mb card that can store over 200 medium-compression JPEG's will only store 34 TIFF's.

Image colour values

The colour values to be calculated and used in any analysis are stored as RGB values in Tiff files automatically when a camera saves an image or when a file is converted into a Tiff image from its RAW file format, and if 8-bit, this is on a scale of 0-255. The camera or computer conversion software may have the option to save the image as either 8-bit or 16-bit, but 8-bit is currently more standard. The steps that follow to calculate values corresponding to, for example reflection or photon catches, are spelt out below. However, one crucial point to emphasise is that many image software programmes offer the option to convert values into other colour spaces, such as HSB (three images corresponding to hue, saturation and brightness). Conversions such as HSB should be avoided and we strongly advise against this type of conversion. HSB is a subjective measurement, generally for human vision, and even in terms of human vision, it is unlikely to be accurate – a more widely used and well tested colour space for humans is CIE Lab (Commission Internationale de l'Éclairage), which may in some cases be appropriate. There are numerous pitfalls with using methodological techniques based on human vision to describe animal colours (Bennett *et al.*, 1994; Stevens & Cuthill, 2005).

Software

One of the biggest advantages of using images to analyse coloration is the existence of a huge number of flexible and powerful software programmes, coupled with the option to write custom programmes in a variety of programming languages. Some of

the programmes available to deal with image processing are standard and quite affordable, such as Paintshop Pro (Jasc Software, Minneapolis, USA) or Photoshop (Adobe Systems Inc., San Jose, USA), which can be used for a range of simple tasks. However, there are a range of other options available, including, the popular freeware Java-based (Sun Microsystems, Inc. Santa Clara, USA; Efford, 2000) imaging programme 'Image J' (Rasband, 1997-2006; Abrámoff, Magalhäes & Ram, 2004), with its huge variety of available 'plugins', written by various people for a range of tasks. This also permits custom programmes written in the language Java to accompany it. For instance, a plugin we used, called 'radial profile,' is ideal for analysing lepidopteran eyespots, and other circular features. This works by calculating the normalised intensities of concentric circles, starting at a central point, moving out along the radius. Figure 5 gives an example of this plug-in, used to analyse an eyespot of the ringlet butterfly *Aphantopus hyperantus*.

The programme MATLAB (The Mathworks Inc. Massachusetts, USA) is also an extremely useful package for writing calibrations and designing sophisticated computational models of vision. This is a relatively easy programming language to learn, is excellent for writing custom and powerful programmes, and due to its matrix manipulation capabilities is excellent for dealing with images (images are simply matrices of numbers). MATLAB can also be bought with a range of 'toolboxes' which have numerous functions already written for various tasks, including, image processing, statistics, wavelet transformations and various other disciplines. MATLAB has available an Image Processing Toolbox with a range of useful functions (Hanselman & Littlefield, 2001; Hunt *et al.*, 2003; Gonzalez, Woods & Eddins, 2004; Westland & Ripamonti, 2004).

How frequently should calibrations be undertaken?

The frequency with which calibrations should be undertaken depends upon the specific calibration required. For instance, determining the spectral sensitivity of a camera's sensors need only be done once as this should not change with time. Additionally, the calculation of the camera's response to changing light levels and the required linearisation need only be done once as this too does not change with time. However, if calculating reflection the calibration needs to be done for each session/light setup as the light setup changes the ratio between the LW, MW and SW sensors.

Calibrating a digital camera

There are several steps that should be followed when wishing to obtain values of either reflection, or data corresponding to an animal's visual system. To obtain values of reflection:

1. Obtain photographs of a set of reflectance standards used to fit a calibration curve.
2. Determine a calibration curve for the camera's response to changes in light intensity in terms of RGB values.
3. Derive a linearisation equation, if needed, to linearise the response of the camera to changes in light intensity, based on the parameters determined from step 2.
4. Determine the ratio between the camera's response in the R, G and B channels, with respect to the reflectance standards, and equalise the response of the different colour channels to remove the effects of the illuminating light and any biases inherent in the camera's processing.

If data corresponding to an animal's visual system is required (such as relative photon catches):

1. Obtain photographs of reflectance standards through a set of narrow band-pass filters, whilst also measuring the radiance with a spectrophotometer.
2. Determine the linearity of the camera's response to changing light levels and if necessary derive a linearisation. Also, using radiance data and the photographs through the band-pass filters, determine the spectral sensitivity of the camera's different sensor types.
3. Using data on the spectral sensitivity of the camera's sensors, and the sensitivity of the animal's sensors to be modelled, produce a mapping based on the response to many different radiance spectra between the two different colour spaces.

These different steps are discussed in detail below.

Linearisation

If one photographs a set of grey reflectance standards and then plots the measured RGB values against the nominal reflectance value, the naïve expectation would be of a linear relationship (Lauziére *et al.*, 1999). One might also expect the values obtained for each of the three colour channels to be the same for each standard, as greys fall on the achromatic locus of R=G=B (see e.g. Kelber *et al.*, 2003). However, as mentioned previously, many cameras do not fulfil such expectations, and they did not for the Nikon 5700 Coolpix camera that we used in our study (Figure 6; see Appendix 2). A different non-linear relationship between grey value and nominal reflection for each colour channel requires that the linearising transformation must be estimated separately for each channel. Also, it means that an image of a single grey reflection

standard is insufficient for camera calibration; instead a full calibration experiment must be performed.

We used a modification of the linearization protocols developed by Párraga (2003) and Westland and Ripamonti (2004). The first step is to photograph a range of standard greyscales of known reflectance value. Westland and Ripamonti (2004) used the greyscale of the Macbeth ColorChecker chart (Macbeth, Munsell Color Lab, New Windsor, N.Y. USA). In our study, as we required reflection standards suitable for UV photography (see later), we used a set of Spectralon^{TM} diffuse reflectance standards (Labsphere Inc., North Sutton, NH, USA). These standards, made of a Teflon micro-foam, reflect light of wavelengths between 300 nm and 800 nm (and beyond) approximately equally, and are one of the most highly Lambertian substances available over this spectral range. The standards had nominal percentage reflection values of 2%, 5%, 10%, 20%, 40%, 50% and 75%. If the object of the study, as in Westland and Ripamonti (2004; Chapter 10) and here, is to recover reflectance data from the images, then the nature of the illuminant, as long as it is stable over time (e.g. Angelopoulou, 2000) and of adequate intensity in all wavebands, is irrelevant. We used a 150 Watt Xenon arc lamp (Light Support, Berkshire, UK), allowed to warm up and stabilise for 1 hr before the calibration exercise, and then tested for stability before and after the calibration. In Párraga (2003) the goal was to recover spectral radiance, so at the same time as photographing the standards, the radiance of each greyscale patch was measured using a spot-imaging telespectroradiometer (TopCon Model SR1, Calibrated by the National Physical Laboratory, UK). After that, each sensor's grey level output was plotted against a measure of the total spectral radiance that stimulated it, at various shutter speeds. The next step is to determine the function relating the intensity values (0 to 255) for each of the RGB sensors to true reflection, or radiance, as appropriate, as measured spectrometrically. Many studies

describe power functions of the same family as those relating intensity to voltage in CRT (Cathode Ray Tube) monitors; so-called gamma functions of the type: Output = constant * (input ^ gamma), or refer to the process of 'gamma correction' in linearising data. The term gamma function means different things in film photography, digital photography and algebraic mathematics, and so is a potentially confusing term that is best avoided. Since the response of the camera's sensors is likely to be camera specific, we recommend determining the curve that best fits the data. Whilst many curves will no doubt fit the data very closely (for example, a Modified Hoerl and Weibull model, amongst others, fitted our reflection data very well), it is important to choose a function that is the same for each of the camera's three sensors; this makes producing the calibrations much easier since the calibration equation will be of the same form for each channel, with only the parameters varying. If there are several curves that all fit the data well, then choosing the simplest equation and with the lowest number of parameters makes calibration much easier. We found that the function below fitted our camera well;

$$Q_S = a*b^P \quad (1)$$

where *Qs* is the photon catch of a given sensor S (R,G or B), *P* the value of the pixel of sensor S, and *a* and *b* are constants. *Qs* is the product of the measured radiance spectrum and the sensor's spectral sensitivity, but it is rare for manufacturers to publish such data. Westland and Ripamonti (2004) mention that luminance is sometimes used as an approximation, on the assumption that for a grey standard the radiance in all three channels should be the same. But this assumes a spectrally flat light source, which no light source ever really is. Therefore the spectral sensitivity needs to be measured directly, by measuring the camera RGB values when imaging

greyscales illuminated through narrow band-pass filters. In this way one can construct spectral sensitivity curves analogous to the spectral sensitivity curves of animal photoreceptors (Figure 1). Párraga (2003) contains technical details of how to achieve this.

In Párraga's (2003) linearization exercise, the value of *b* in the equation above was found to be similar for all conditions tested (sunny, cloudy and incandescent artificial light) and all sensors. Thus, the value of *a* defined each curve. As the linearised values for R, G and B represented radiance in the wavebands specific to each sensor, the photograph's exposure was also taken into account in the calibration process (a longer exposure time representing lower radiance). Therefore, the following three equations were derived to linearise and scale each RGB value to radiance measures, where Q_S is the radiance measured by sensor S, *b* and the a_i are the coefficients estimated by ordinary least-squares regression of log-transformed values, *c* is a value to account for inherent dark current (see below) in the camera, and t is the integration time the photograph was taken on [1/shutter speed]):

$$Q_R = a_1(b^R - c_1)/\ \mathrm{t} \qquad (2)$$

$$Q_G = a_2(b^G - c_2)/\ \mathrm{t} \qquad (3)$$

$$Q_B = a_3(b^B - c_3)/\ \mathrm{t} \qquad (4)$$

If the object of the research is to obtain reflection rather than radiance measures, then t can be ignored and functions such as equation 1 could be used, provided that t is constant and measurements of known reflection standards are also

made. Because the reflection values of greyscales are, by definition, equal in all wavebands, sensor spectral sensitivity does not in principle need to be known for linearisation in relation to reflection, although in practice, one would want to know the spectral sensitivity curves for the camera's sensors in order for the data to be readily interpreted (in terms of the sensitivity corresponding to each sensor). In the case of either radiance or reflection calibration, one should check that it is valid to force the calibration curve through the origin. All digital imaging sensors have associated with them an inherent 'dark current' (due to thermal noise in the sensors) (Efford, 2000; Stokman, Gevers & Koenderink, 2000; Barnard & Funt, 2002; Martinez-Verdú *et al.*, 2002), so that a set of images with the lens cap on may not produce measurements of zero. As with spectrometry, the dark current can be estimated by taking images at the same exposure settings as calibration photos, and using the pixel values as an offset for the curve. One should also check whether increasing the integration time, or temperature changes within the range at which the camera will be used for data collection, alters these background dark current values. Figure 7 is an example of linearisation performed on the RGB values from photographs of reflectance standards (Figure 6). This shows that generally the linearization was successful. However, one should note that the values of the reflectance standards with low nominal reflection values are not accurate, since these standards were partially underexposed (i.e. there are many pixels with values equal or close to the dark current values), and for this specific set of standards some standards are slightly closer to, or further away from, the light source. This means that the calibration line will not be perfectly straight. Since the relatively darker areas (low pixel values) of images are often inaccurate in the measurements they yield, these values may be non-linear (Barnard & Funt, 2002). However, the measurement error is relatively small.

RGB equalisation

If the goal is to derive reflection data from the photographs, then greys should, by definition, have equal reflection in all three colour channels. So, if R≠G≠B in the calibration images, the next step is to equalise the three channels with respect to the images of the reflection standards, and then scale the values between 0 and 255. This, in theory, should be relatively simple: a matter of producing a ratio between the three channels and then scaling them, usually with respect to the green channel as a reference point, before multiplying the entire image by 2.55 to set the values on a scale between 0 and 255. So, for our data:

$$R' = (Rx_R)2.55 \quad (5)$$

$$G' = (Gx_G)2.55 \quad (6)$$

$$B' = (Bx_B)2.55 \quad (7)$$

where x_i is the scaling value for each channel, and *R*, *G*, and *B* are the linearised image values for each channel respectively. The equalised values were then tested for accuracy using a different set of calibration images. Figure 8 shows the result. The three channels closely match the required calibration line. Note that there is no need for 255 to represent 100% reflection; indeed, in order to obtain maximum resolution in colour discrimination within and between images, if all images to be analysed are relatively dark then it would be advisable the maximum pixel value within the dataset to 255.

An important issue is that of saturation. With regards to the above calibration results (Figures. 6-8), we maintained an integration time of 1/30 seconds and a lens aperture of 8.0. This resulted in images that were slightly under-exposed and guarded against the serious problem of saturation. Saturation (also known as 'clipping;' Lauziére *et al.*, 1999) occurs when the light levels arriving at the sensors reaches an upper limit, above which any more photons are not registered. This can be a serious problem because it prevents measurements of the true value that the pixels would have reached had saturation not occurred; a problem recognised in some studies (e.g. Hong *et al.*, 2001). The effects of saturation are easy to find, with saturated pixels in the original image yielding values of around 255, with little or no standard deviation. For example, images taken under similar circumstances, but with an integration time of 1/15 seconds produce results that at nominal reflection values of 75%, the red channel ceases to rise in pixel values. This is due to the effects of saturated pixels in the original image in the red channel, which causes the calibration to fail, since the linearisation becomes ineffective and the equalisation procedure results in the red channel grey values dropping away at higher reflection values (Figure 9). These problems can be avoided by changing the exposure/integration time (t), or altering the intensity of the light source, since these determine the flux of light reaching the camera's sensors (Hong *et al.*, 2001). However, if the exposure is to be changed between images it is important to test that the response of the camera is the same at all exposure settings, otherwise a separate calibration will need to be performed for every change in exposure. Therefore, where possible, it is recommended that the aperture value, at least, is kept constant (Hong *et al.*, 2001).

It is often the case that the red channel of a digital camera is the first to saturate (as was the case with our camera; Figure 9), possibly because the sensors in some cameras may be biased to appeal to human perceptions, with increasing red

channel values giving the perception of warmth. This may be particularly deleterious for studies investigating the content of red signals (e.g. Frischknecht, 1993; Wedekind *et al.*, 1998), which are widespread because of the abundance of carotenoid-based signals in many taxa (e.g. Grether, 2000; Pryke, Lawes & Andersson, 2001; Bourne, Breden & Allen, 2003; Blount, 2004; McGraw & Nogare, 2004; McGraw, Hill & Parker, 2005) and theories linking carotenoid signals to immune function (Koutsos *et al*., 2003; McGraw & Ardia, 2003; Navara & Hill, 2003; Grether *et al*., 2004; McGraw & Ardia, 2005). Some cameras are also biased in their representation of relatively short wavelengths, to compensate for a lack of these wavelengths in indoor lights (Lauziére *et al.*, 1999).

Selecting/controlling light conditions

The importance of selecting standardised lighting conditions and distances to some extent depends upon the calibration required. Lighting conditions should be as stable, standardised and consistent as possible for each photo if measurements of reflection are desired, especially if photographs of standards are taken only at the beginning and end of sessions. However, when photographing natural scenes and using measures of photon catch, for example, lighting conditions are likely to vary considerably. This may in fact be an important part of the study – to include information about the ambient light. Generally, it is best to avoid flashes since the output of these is difficult to measure and may be variable, however, a high-end flash with good light diffusers may be fine.

Mapping to camera-independent measures

Having used the coefficients obtained in the linearisation procedure to linearise the RGB values in one's images, the next step is to transform them to camera-independent values. This is because the R, G and B data, whether in radiance or reflectance units, are specific to the wavebands designated by the camera sensors' spectral sensitivity curves (e.g. as in Figure 1C). This may be sufficient for some research purposes; for example, if the sensitivities of the camera's sensors broadly correspond to the bandwidths of interest. However, it will often be desirable, either because a specific visual system is being modelled (e.g. human, bird), or simply to facilitate comparison of the results across studies, to transform the camera-specific RGB values to camera-independent measures. In human studies, these are frequently one of the sets of three-coordinate representations devised by the CIE for colour specification and/or matching. Different three-variable representations have been devised to approximate colour-matching for images illuminating only the M-L-cone-rich central fovea, or wider areas of the retina; for presentation of images on VDUs or printed paper; or representations that incorporate the colour balance arising from a specific illuminant, or are illumination independent (Wyszecki & Stiles, 1982; Mollon, 1999; Westland & Ripamonti, 2004). The advantage is that all these metrics are precisely defined, the formulae downloadable from the CIE website, and the values in one coordinate system can be transformed to another. Westland & Ripamonti (2004) provide formulae and downloadable MATLAB (The Mathworks Inc. Massachusetts, USA) code for such transformations.

Another possible camera-independent transformation is to map the linearised RGB values to the spectral sensitivities of the photoreceptors of either humans (Párraga *et al.*, 2002) or non-human species. In the case of RGB radiance measures this corresponds to calculating the photon catches of an animal's photoreceptors, rather than the camera's sensors, when viewing a particular scene. In the case of RGB

reflectance measures, this can be thought of as a mapping to a species-specific estimate of reflectance in the wavebands to which the animal's photoreceptors are sensitive. Both types of mapping are particularly relevant to studies involving non-human animals, where accurate psychophysical estimates of colour-matching, of the sort used to calculate human-perceived colour from camera data, are not usually available. For such mapping to be viable, it is not necessary that the species' cone spectral sensitivities match those of the camera's sensors particularly closely (for example, this is not true for humans; compare Figure 1A with 1C). However, for the transformation to produce reliable data, the species' overall spectral range has to fall within that of the camera, and the species has to have three or less photoreceptors. For example, one can map RGB data to the lower dimensional colour space of a dichromatic dog (with one short- and one medium/long-sensitive cone type; Jacobs, 1993), but a camera with sensitivities such as shown in Figure 1C can never capture the full colour world of a trichromatic bee (with UV, short- and medium-wave photoreceptors; Chittka, 1992). Mapping RGB data to a bird's colour space would seem invalid on two counts: birds have a broader spectral range than a conventional camera (extending into the UV-A) and are potentially tetrachromatic (reviewed by Cuthill *et al.*, 2000). However, if the scenes or objects of interest lack UV information, then a mapping from RGB to avian short-, medium-, long-wave cone sensitivities can be achieved. We present the method here, which can be used for any analogous trichromatic system (e.g. human) or, with simple modification, a lower-dimensional system of the type that is typical for most mammals (Jacobs, 1993). Subsequently, we consider how UV information from a separate imaging system can be combined with the RGB data to provide a complete representation of bird-perceived colour.

The goal is to predict the quantal catches, Qi, of a set of i photoreceptors (where $i \leq 3$), given a triplet of camera-sensor-estimated radiance values, Q_R, Q_G and Q_B, derived from the calibration and linearisation process described above. This amounts to solving a set of simultaneous regression equations, which are likely to be non-linear. Mappings can be done for more than three photoreceptor classes, provided that the spectral sensitivities of all types are covered by the spectral range of one or more of the camera's sensors. For example, a mapping could be produced to calculate images corresponding to the longwave, mediumwave and shortwave cones of a bird's visual system, plus a luminance image based on avian double cone sensitivity. Once mapped images have been obtained, further calculations also allow the production of images corresponding to various opponency channels. Westland and Ripamonti (2004) summarise their, and others', research on the family of equations most likely to provide a good fit to data, and conclude that linear models (with interaction terms) of the following type perform well. For ease of interpretation we use the notation R, G and B to describe the camera pixel values rather than their calibrated and linearised equivalents, Q_R, Q_G and Q_B.

$$Q_i = b_{i1}.R + b_{i2}.G + b_{i3}.B + b_{i4}.R.G + b_{i5}.R.B + b_{i6}.G.B. + b_{i7}.R.G.B \quad (8)$$

Where the b_i's are coefficients specific to receptor i, and the curve is forced through the origin (when the calibrated camera sensor value is zero, the animal's quantal catch is zero). In some cases, dependent on the camera and the nature of the visual system to which mapping is required, polynomials (i.e. including terms in R^2, G^2 and B^2, or higher orders) may provide a significantly better fit (and did in our case); this should be investigated empirically. Cheung *et al.* (2004) note that even mapping functions of unconstrained form, obtained using neural networks applied to large datasets, do not

significantly outperform polynomials. The data required to estimate the coefficients for the i photoreceptors can either be radiances directly measured using an imaging spectroradiometer (as in Párraga, 2003) or, more conveniently, radiances approximated as the product of reflectance spectra and the irradiance spectrum of the illuminant. Using equation 8, applied to a trichromat, 3×7 coeffcients need to be estimated, so the number of radiance spectra must be considerably greater than this (> 100 in our experience as a minimum, but more like a 1000 is better). Large numbers of radiance spectra can be obtained from internet databases (Parkkinen, Jaaskelainen & Kuittinen, 1988; Sumner & Mollon, 2000). The coefficients for each photoreceptor are then found by multiple regression (or, conveniently, if using MATLAB, by matrix algebra; see Westland and Ripamonti, 2004). Whilst, in principle, one could derive a mapping function (i.e. set of coefficients) for all possible natural spectra, viewed under all possible illuminants, greater precision can be achieved by determining a situation-specific mapping function for the research question at hand. For example, if the goal is to use a camera to quantify the coloration of orange to red objects under blue skies, then a very precise mapping function could be estimated by using radiance data calculated only from the reflectance spectra of orange to red objects viewed under blue sky irradiance. If one is to derive the mapping functions by calculation (i.e. calculate quantal catch for camera and desired cone sensitivities, using reflectance and irradiance data), then the sensitivity of the camera's sensors is required. However, one could also derive the mapping empirically without ever measuring camera sensor sensitivities, by measuring the response of the camera's three channels to different (known) radiance spectra, and by determining the response of the cones of the required animal's visual system. To get an accurate mapping, the camera's response would have to be measured for many hundreds of radiance spectra and this would be time-consuming, involving many stimuli.

UV imaging

In our own research, we wished to quantify lepidopteran wing patterns, with respect to avian vision, so we also needed to measure the amount of reflection in the avian-visible UV waveband. At the same time as RGB photography, images of the reflectance standards and the lepidopterans were taken with a UV sensitive video camera (see Appendix 2).

First we tested whether the camera was linear with respect to both changes in the integration time, and with respect to increases in the reflection value; being a high-specification technical camera, it was. This meant that the only calibrations needed were to scale the images to between 0 and 255; not initially as easy as it sounds, since the calibrations have to account for different gain and the integration times. Figure 10 is an example of the results for the UV calibration process. In most situations it will be simpler to maintain the same gain values, since this reduces the number of factors to consider in the calibration process.

If images are obtained from more than one camera, there is an additional consideration that must be addressed; that of 'image registration'. Images derived from one RGB camera will all be the same angle and distance from the specimens, and so the objects photographed will be on an identical scale in each of the three channels, based on the interpolations implemented. This may not be the case if obtaining images from a second camera; such as in our study, meaning that the specimens were a different size in the photographs and would not necessarily be easy to align with the RGB images. Furthermore, one camera may produce images with a lower resolution, and with less high frequency information; different cameras will have different Nyquist frequencies, meaning that whilst aligning lower spatial

frequency patterns may be relatively easy, information may be lost or poorly aligned at higher frequencies. One potential approach is to use Fourier filtering to remove the highest spatial frequency information from those images which contain it, down to the highest frequencies contained in the images from the other camera. However, this may be undesirable if the high spatial frequency information is important, as it frequently will be with complex patterns, or where edge information between pattern components is critical. The task of aligning images is made easier if: (a) different cameras are set up as closely as possible, in particular with relation to the angle of photography, since this is the hardest factor to correct and; (b) rulers are included in at least a sample of the images, so they can be rescaled so specimens occupy the same scale in different images. Including rulers in images allows for true distance measurements to be obtained and for spatial investigations to be undertaken. If images from one camera are larger than those from another, then it is the larger images that should be scaled down in size, since this avoids artifactual data, generated by interpolation, if images are rescaled upwards. Once the objects in the photographs are of the same size, it may be a relatively trivial task to take measurements from the different images that directly correspond. However, if the images are still difficult to align then an automated computational approach can be used. A variety of these are available, and users should carefully consult available manuals/information for the corresponding software to be sure of how the registration is completed, and to check what changes may occur to the image properties. However, in many cases, changes to the image structure will probably be small, especially at lower spatial frequencies, and have little influence on the results. One such plug-in, for the freeware software Image J (Rasband, 1997-2006; Abràmoff *et al.*, 2004), is 'TurboReg' (available from a link from the Image J website) (Thévenaz, Ruttimann & Unser, 1998), which comes with a variety of options to align sets of images.

How best to use colour standards

A crucial step in calibrating a digital camera is to include colour standards in some or all of the photographs taken. Including a set of colour standards in each photo allows calibrations to be derived for each individual photo, which would be highly accurate. However, this is in most cases impractical and unnecessary. For example, when the light source used is consistent, a set of reflectance standards used to fit a calibration curve need only be included in photos at the start and end of a session. Including these in each photo may leave little space for the objects of interest. In contrast, in many cases, such as when photographing natural scenes where the illuminating light may change and when wishing to calculate values such as photon catches it may be important to include at least one grey standard in the corner of each photo. Possibly the best objects to include in a photo are Spectralon^{TM} reflectance standards (Labsphere Inc., North Sutton, NH, USA), which reflect a known amount of light equally at all wavelengths in the ultraviolet and human visible spectrum. These are, however, expensive and easily damaged, and if a single standard is sufficient, a Kodak gray card (Eastman Kodak Company), which has an 18% reflectance, can be included.

Spatial measurements

Often, we do not wish to measure solely the 'colour' of a patch, but the area or shape of a region of interest. In principle, this sounds easy but has several complications. For example the colour boundary of an area visible to humans may not be exactly the same as for that of another animal. Additionally, there may be colours that we cannot see (such as ultraviolet), which have different boundaries to those visible by a human (although most colour patches generally have the same boundary for different colour

bands, such as UV, SW, MW and LW). Another problem corresponds to the acuity of the animal in question. Regions of interest with complex boundaries may be only discernable by animals with a high enough spatial acuity. Furthermore, there is a specific problem with gradual boundaries, particularly relating to defining where the actual edge of the colour region is.

There are several ways to address these issues. One method of determining the boundary of a colour patch is to produce an automated procedure to define a specific area of interest. This can be done by thresholding an 8-bit or colour image to a binary (black and white) image where each individual pixel has a value of either zero (white) or two (black) (Figure 11). This can be done by writing a custom programme where the threshold level is defined specifically by the user, perhaps based on some assumption or data. Otherwise, most imaging software has automatic thresholding algorithms, although it is not always known what the thresholding value used will be.

A different method that can be used to define an area of interest is that of edge detection. This is where an algorithm is used to determine edges in an image, corresponding to sharp changes in intensity (either luminance or in terms of individual colour channels). These edges may, for example, be found at the boundary of a colour patch (Figure 11). The useful thing about edge detection algorithms is that they can be optimised and not linked to any specific visual system, or they correspond to the way in which real visual systems work (Marr & Hildreth, 1980; Bruce, Green & Georgeson, 2003).

Once the boundary of a colour patch has been defined it is simple to measure the area of the patch. Measuring the shape of an object is more difficult, though imaging software often comes with algorithms to measure attributes such as the relative circularity of an area, and occasionally more advanced shape analyses algorithms.

Drawbacks to using digital images

The most notable drawback is that the information obtained is not wavelength specific; i.e. it is known what wavelengths contribute to each channel, but not the contribution of any specific wavelength to the RGB value of any one pixel. This drawback can be overcome by so-called multispectral imaging (or, if the number of wavebands is high, 'hyperspectral imaging'). This can involve rotating a set of filters in front of the lens, allowing the acquisition of successive images of different wavebands (e.g. Brelstaff *et al*, 1995; Lauziére *et al.*, 1999; Angelopoulou, 2000; Stokman *et al.*, 2000; Losey, 2003). This method may be particularly useful if detailed wavelength information is required, or if the visual system of the receiver the signal is aimed at is poorly matched by the sensitivity of an RGB camera. We do not cover this technique in this paper because, although it combines many of the advantages of spectrometry with photography, the technology is not practical for most behavioural and evolutionary biologists. Hyperspectral cameras are often slow because they may have to take upwards of 20 images through the specified spectral range. The equipment, and controlling software, must be constructed *de novo* and conventional photography's advantage of rapid, one-shot, image acquisition is lost. The specimens must be stationary during the procedure, since movement can cause problems with image registration. Also, as Losey (2003) acknowledges, images obtained sequentially in the field may be subject to short-term variations in environmental conditions, and thus introduce considerable noise. Averaging the values obtained from multiple frames of the same waveband may help to eliminate some of this effect (Losey, 2003).

Problems with using the automatic camera settings

Many studies of animal coloration utilising cameras apparently use the camera with its automatic settings. There are numerous problems that can arise when using the 'auto' mode. The main problem is that the settings used by the camera are adjusted according to the scene being photographed and so may be inconsistent. This need not always be an irretrievable flaw but would almost certainly need some highly complex calibration procedures to recover consistent data. One issue is that the white balance may change between photos, giving rise to different ratios between the LW, MW and SW sensor responses. Any low- to mid-range camera is likely to have some white balancing present, and most mid-range cameras will give the option to manually set the white balance. If the camera doesn't allow this option and there is no indication of this in the manual then changing the white-balance settings may not be possible. An additional problem with automatic settings is that calibration curves/settings may also change at different aperture settings, in addition to the complication that the aperture and exposure (integration) time may change significantly simultaneously, leading to complicated calibrations if values of reflection, for example, are required. One of the most serious problems with using the auto mode is that the photograph will not optimise the dynamic range of the scene photographed, meaning that some parts of the scene may be underexposed, or far more seriously, saturated.

Conclusions

One of the earliest studies to outline how digital image analysis can be used to study colour patterns is that of Windig (1991), with an investigation of lepidopteran wing patterns. Windig (1991) used a video camera, connected to a frame grabber to digitize the images for computer analysis, a similar method to that which we used to capture the UV sensitive images. Windig (1991) stated that the method was expensive, and

the programmes were highly complex, but today flexible user friendly software is available, with various freeware programmes downloadable off the internet, and the purchase of a digital camera and software is possible for a fraction of the cost of Windig's (1991) setup.

Windig (1991) argued that any image analysis procedure should meet three criteria. Firstly, completeness: a trait should be quantified with respect to all characters, such as 'colour' and area. Our procedure meets this criterion, since reflection, plus spatial measurements are attainable. Secondly, the procedure needs to be repeatable. This was also the case with our approach as the calibrations for a set of images of reflectance standards were still highly accurate for other images taken under the same conditions, but at different times. Finally, the process should be fast relative to other available methods, as was our study, with potentially hundreds of images taken in a day, quickly calibrated with a custom MATLAB programme and then analysed quickly with the range of tools in Image J.

Another advantage of capturing images with a digital camera is that there are potentially a host of other non-colour analyses. Detailed and complex measurements of traits can be undertaken rapidly, with measurements and calculations that would normally be painstakingly undertaken by hand performed almost instantaneously in imaging software, including measurements of distances, areas, analysis of shapes, plus complex investigations such as Fourier analysis (Windig 1991). This may be particularly useful if handling the specimens to take physical measurements is not possible.

The use of digital technology in studying animal coloration is a potentially highly powerful method, avoiding some of the drawbacks of other techniques. In future years, advances in technology, software and our understanding of how digital cameras work will add further advantages. In fact, it is even possible already to

extract data of a scene from behind a plane of glass (Levin & Weiss, 2004), which could become useful for studies of aquatic organisms (though most glass filters out UV wavelengths; Lauziére *et al.*, 1999). Techniques are also being developed to remove the shadows from images; shadows can make edge recognition more difficult (e.g. Finlayson, Hordley & Drew, 2002), and hinder tasks such as image registration. With the explosion in the market of digital photography products, and the relatively low cost to purchase such items, there is the temptation to launch into using such techniques to study animal signals, without prior investigation into the technicalities of using such methods. This could result in misleading results. Therefore, whilst digital photography has the potential to transform studies of coloration, caution should be implemented and suitable calibrations developed before such investigations are undertaken.

Key points/summary

Below is a list of some of the main points to consider if using cameras to study animal coloration.

1. Images used in an analysis of colour should be either RAW or Tiff files and not JPEGs.
2. Grey reflectance standards should be included in images at the start of a photography session if the light source is constant, or in each image if the ambient light changes.
3. It is crucial not to allow images to become saturated or underexposed as this prevents accurate data to be obtained.
4. Many cameras have a non-linear response to changes in light intensity, which needs linearising before usable data can be obtained.

5. To produce measurements of reflectance the response of the R, G and B colour channels needs to be equalised with respect to grey reflectance standards.
6. Measurements of cone photon catches corresponding to a specific visual system can be estimated by mapping techniques based upon sets of radiance spectra and camera/animal spectral sensitivity.
7. Digital images can be incorporated in to powerful and analyses of animal vision.
8. Do not convert image data to formats such as HSB, which are subjective and inaccurate. Instead, use reflection data, values of cone stimulation, or if working in human colour space, well tested colour spaces such as CIE Lab.
9. If using more than one camera, image registration may be a problem, especially if the different cameras have different resolutions. This problem can be minimised by setting up different cameras as close to one another as possible and ensuring that one camera does not capture significantly higher levels of spatial detail than the other.
10. Digital imaging is also a potentially highly accurate and powerful technology to study spatial patterns.

Acknowledgements

We are very grateful to D.J. Tolhurst and G. P. Lovell for help with much of the project. C.A.P. was supported by BBSRC grants S11501 to D.J. Tolhurst and T.S.T., and S18903 to ICC, TST and JCP. MS was supported by a BBSRC studentship.

References

Abràmoff MD, Magalhäes PJ, Ram SJ. 2004. Image processing with Image J. *Biophotonics International* **7**: 36-43.

Angelopoulou E. 2000. Objective colour from multispectral imaging. *International Conference on Computer Vision* **1**: 359-374.

Barnard K, Funt B. 2002. Camera characterisation for color research. *Color Research and Application* **27:** 152-163.

Bennett ATD, Cuthill IC, Norris KJ. 1994. Sexual selection and the mismeasure of color. *American Naturalist* **144:** 848-860.

Blount JD. 2004. Carotenolds and life-history evolution in animals. *Archives of Biochemistry and Biophysics* **430:** 10-15.

Bortolotti GR, Fernie KJ, Smits JE. 2003. Carotenoid concentration and coloration of American Kestrels (*Falco sparverius*) disrupted by experimental exposure to PCBs. *Functional Ecology* **17:** 651-657.

Bourne GR, Breden F, Allen TC. 2003. Females prefer carotenoid colored males as mates in the pentamorphic live bearing fish, *Poecilia parae*. *Naturwissenschaften* **90:** 402-405.

Brelstaff GJ, Párraga CA, Troscianko T, Carr D. 1995. Hyperspectral camera system: acquisition and analysis. In: Lurie JB, Pearson J & Zilioli E eds. *SPIE - Human vision, visual processing and digital displays. Geographic information systems, photogrammetry, and geological/geophysical remote sensing* Paris: Proceedings- SPIE, 150-159.

Bruce V, Green PR, Georgeson MA. 2003. *Visual Perception: Fourth Edition*. Hove: Psychology Press.

Cardei VC, Funt B, Barnard K. 1999. White point estimation for uncalibrated images. *Proceedings of the IS&T/SID Seventh Color Imaging Conference: Color Science Systems and Applications* 97-100.

Cardei VC, Funt B. 2000. Color correcting uncalibrated digital images. *The Journal of Imaging Science and Technology* **44**: 288-294.

Cheung V, Westland S, Connah D, Ripamonti C. 2004. A comparative study of the characterisation of colour cameras by means of neural networks and polynomial transforms. *Coloration Technology* **120:** 19-25.

Chittka L. 1992. The colour hexagon: a chromaticity diagram based on photoreceptor excitations as a generalised representation of colour opponency. *Journal of Comparative Physiology A* **170:** 533-543.

Cooper VJ, Hosey GR. 2003. Sexual dichromatism and female preference in *Eulemur fulvus* subspecies. *International Journal of Primatology* **24:** 1177-1188.

Cott HB. 1940. *Adaptive Colouration in Animals*. London: Methuen & Co. Ltd.

Cuthill IC, Bennett ATD, Partridge JC, Maier EJ. 1999. Plumage reflectance and the objective assessment of avian sexual dichromatism. *American Naturalist* **153:** 183-200.

Cuthill IC, Hart NS, Partridge JC, Bennett ATD, Hunt S, Church SC. 2000. Avian colour vision and avian video playback experiments. *Acta Ethologica* **3:** 29-37.

Cuthill IC, Partridge JC, Bennett ATD, Church SC, Hart NS, Hunt S. 2000. Ultraviolet vision in birds. *Advances in the Study of Behaviour* **29:** 159-214.

Dartnall HJA, Bowmaker JK, Mollon JD. 1983. Human visual pigments: microspectrophotometric results from the eyes of seven persons. *Proceedings of the Royal Society of London, Series B* **220:** 115-130.

D'Eath RB. 1998. Can video images imitate real stimuli in animal behaviour experiments? *Biological Reviews* **73:** 267-292.

Efford N. 2000. *Digital Image Processing: a practical introduction using Java*. Harlow, Essex: Pearson Education Ltd.

Endler JA. 1984. Progressive background matching in moths, and a quantitative measure of crypsis. *Biological Journal of the Linnean Society* **22:** 187-231.

Endler JA. 1990. On the measurement and classification of colour in studies of animal colour patterns. *Biological Journal of the Linnean Society* **41:** 315-352.

Endler JA, Mielke PW Jr. 2005. Comparing color patterns as birds see them. *Biological Journal of the Linnean Society* **86**: 405–431.

Finlayson GD, Tian GY. 1999. Color normalisation for color object recognition. *International Journal of Pattern Recognition and Artificial Intelligence* **13:** 1271-1285.

Finlayson GD, Hordley SD, Drew MS. 2002. Removing shadows from images. *European Conference on Computer Vision* **4:** 823-836.

Fleishman LJ, McClintock WJ, D'Eath RB, Brainard DH, Endler JA. 1998. Colour perception and the use of video playback experiments in animal behaviour. *Animal Behaviour* **56:** 1035-1040.

Fleishman LJ, Endler JA. 2000. Some comments on visual perception and the use of video playback in animal behavior studies. *Acta Ethologica* **3:** 15-27.

Frischknecht M. 1993. *The* breeding colouration of male three-spined sticklebacks (*Gasterosteus aculeatus*) as an indicator of energy investment in vigour. *Evolutionary Ecology* **7:** 439-450.

Gerald MS, Bernstein J, Hinkson R, Fosbury RAE. 2001. Formal method for objective assessment of primate color. *American Journal of Primatology* **53:** 79-85.

Goda M, Fujii R. 1998. The blue coloration of the common surgeonfish, *Paracanthurus hepatus* – II. Color revelation and color changes. *Zoological Science* **15:** 323-333.

Gonzalez RC, Woods RE, Eddins SL. 2004. *Digital Image Processing Using MATLAB*. London: Pearson Education Ltd.

Grether GF. 2000. Carotenoid limitation and mate preference evolution: A test of the indicator hypothesis in guppies (*Poecilia reticulata*). *Evolution* **54:** 1712-1724.

Grether GF, Kasahara S, Kolluru GR, Cooper EL. 2004. Sex-specific effects of carotenoid intake on the immunological response to allografts in guppies (*Poecilia reticulata*). *Proceedings of the Royal Society of London Series B* **271:** 45-49.

Hart NS, Partridge JC, Cuthill IC. 1998. Visual pigments, oil droplets and cone photoreceptor distribution in the European starling (*Sturnus vulgaris*). *Journal of Experimental Biology* **201:** 1433-1446.

Hanselman D, Littlefield B. **2001**. *Mastering MATLAB 6: A Comprehensive Tutorial and Reference.* New Jersey: Pearson Education International.

Hendrickson A, Drucker D. 1992. The development of parafoveal and mid-peripheral human retina. *Behavioural Brain Research* **49:** 21-31.

Hong G, Lou RM, Rhodes PA. 2001. A study of digital camera colorimetric characterization based on polynomial modeling. Color *Research and Application* **26:** 76-84.

Hunt BR, Lipsman RL, Rosenberg JM, Coombes HR, Osborn JE, Stuck GJ. 2003. *A Guide to MATLAB: For Beginners and Experienced Users.* Cambridge: Cambridge University Press.

Hurlbert A. 1999. Colour vision; is colour constancy real? *Current Biology* **9:** R558-R561.

Jacobs GH. 1993. The distribution and nature of colour vision among the mammals. *Biological Reviews* **68:** 413-471.

Kelber A, Vorobyev M, Osorio D. 2003. Animal colour vision – behavioural tests and physiological concepts. *Biological Reviews* **78:** 81-118.

Kodric-Brown A, Johnson SC. 2002. Ultraviolet reflectance patterns of male guppies enhance their attractiveness to females. *Animal Behaviour* **63:** 391-396.

Koutsos EA, Clifford AJ, Calvert CC, Klasing KC. 2003. Maternal carotenoid status modifies the incorporation of dietary carotenoids into immune tissues of growing chickens (*Gallus gallus domesticus*). *Journal of Nutrition* **133:** 1132-1138.

Künzler R, Bakker TCM. 1998. Computer animations as a tool in the study of mating preferences. *Behaviour* **135:** 1137-1159.

Lauziére YB, Gingras D, Ferrie FP. 1999. Color camera characterization with an application to detection under daylight. *Vision Interface, Trois-Rivières, Canada* 280-287.

Levin A, Weiss Y. 2004. User assisted separation of reflections from a single image using a sparsity prior. *European Conference on Computer Vision* **1:** 602-613.

Losey GS Jr. 2003. Crypsis and communication functions of UV-visible coloration in two coral reef damselfish, *Dascyllus aruanus* and *D. reticulates. Animal Behaviour* **66:** 299-307.

Lythgoe JN. 1979. *The Ecology of Vision.* Oxford. Clarendon Press.

Marr D, Hildreth E, 1980. Theory of edge detection. *Proceedings of the Royal Society of London, Series B* **207**: 187-217.

Marshall NJ, Jennings K, McFarland WN, Loew ER, Losey GS Jr. 2003. Visual biology of Hawaiian coral reef fishes. II. Colors of Hawaiian coral reef fish. *Copeia* **3**: 455-466.

Martinez-Verdú F, Pujol J, Capilla P. 2002. Calculation of the color matching functions of digital cameras from their complete spectral sensitivities. *Journal of Imaging Science and Technology* **46:** 15-25.

McGraw KJ, Ardia DR. 2003. Carotenoids, immunocompetence, and the information content of sexual colors: An experimental test. *American Naturalist* **162:** 704-712.

McGraw KJ, Nogare MC. 2004. Carotenoid pigments and the selectivity of psittacofulvin-based coloration systems in parrots. *Comparative Biochemistry and Physiology B* **138:** 229-233.

McGraw KJ, Ardia DR. 2005. Sex differences in carotenoid status and immune performance in zebra finches. *Evolutionary Ecology Research* **7:** 251-262.

McGraw KJ, Hill GE, Parker RS. 2005. The physiological costs of being colourful: nutritional control of carotenoid utilization in the American goldfinch, *Carduelis tristis. Animal Behaviour* **69:** 653-660.

Mollon JD. 1999. Specifying, generating and measuring colours. In: Carpenter RHS, Robson JG eds. *Vision Research. A Practical Guide to Laboratory Methods*. Oxford: Oxford University Press, 106-128.

Navara KJ, Hill GE. 2003. Dietary carotenoid pigments and immune function in a songbird with extensive carotenoid-based plumage coloration. *Behavioral Ecology* **14:** 909-916.

Newton I. 1718. *Opticks, or, A treatise of the reflections, refractions, inflections and colours of light. The second edition, with additions.* London: Printed for W. and J. Innys.

Parkkinen J, Jaaskelainen T, Kuittinen M. 1988. Spectral representation of color images. *IEEE 9th International Comference on Pattern Recognition, Rome, Italy* **2:** 933-935.

Párraga CA. 2003. Is the human visual system optimised for encoding the statistical information of natural scenes? Ph.D. Thesis, University of Bristol.

Párraga CA, Troscianko T, Tolhurst DJ. 2002. Spatiochromatic properties of natural images and human vision. *Current Biology* **12:** 483-487.

Pietrewicz AT, Kamil AC. 1979. Search image formation in the blue jay (*Cyanocitta cristata*). *Science* **204:** 1332-1333.

Pryke SR, Lawes MJ, Andersson S. 2001. Agonistic carotenoid signalling in male red-collared widowbirds: aggression related to the colour signal of both the territory owner and model intruder. *Animal Behaviour* **62:** 695-704.

Rasband WS. 1997-2006. *Image J*. National Institutes of Health, Bethesda, Maryland, USA, http:/rsb.info.nih.gov/ij/.

Rosenthal GG, Evans CS. 1998. Female preference for swords in *Xiphophorus helleri* reflects a bias for large apparent size. *Proceedings of the National Academy of Sciences USA* **95:** 4431-4436.

Samaranch R, Gonzalez LM. 2000. Changes in morphology with age in Mediterranean monk seals (*Monachus monachus*). *Marine Mammal Science* **16:** 141-157.

Stevens M, Cuthill IC. 2005. The unsuitability of html-based colour charts for estimating animal colours – a comment on Berggren & Merilä. *Frontiers in Zoology* 2: 1-14.

Stokman HMG, Gevers T, Koenderink JJ. 2000. Color measurement by imaging spectrometry. *Computer Vision and Image Understanding* **79:** 236-249.

Sumner P, Mollon JD. 2000. Catarrhine photopigments are optimised for detecting targets against a foliage background. *Journal of Experimental Biology* **203:** 1963-1986.

Sumner P, Arrese CA, Partridge JC. 2005. The ecology of visual pigment tuning in an Australian marsupial: the honey possum *Tarsipes rostratus*. *Journal of Experimental Biology* **208**: 1803-1815.

Thayer AH. 1896. The law which underlies protective coloration. *The Auk* **13:** 477-482.

Thayer GH. 1909. *Concealing-Coloration in the Animal Kingdom: An Exposition of the Laws of Disguise Through Color and Pattern: Being a Summary of Abbott H. Thayer's Discoveries*. New York: The Macmillan Co.

Tinbergen N. 1974. *Curious Naturalists: Revised Edition.* Penguin Education Books.

Thévenaz P, Ruttimann UE, Unser M. 1998. A pyramid approach to subpixel registration based on intensity. *IEEE Transactions on Image Processing* **7:** 27-41.

Villafuerte R, Negro JJ. 1998. Digital imaging for colour measurement in ecological research. *Ecology Letters* **1:** 151-154.

Wedekind C, Meyer P, Frischknecht M, Niggli UA, Pfander H. 1998. Different carotenoids and potential information content of red coloration of male three-spined stickleback. *Journal of Chemical Ecology* **24:** 787-801.

Westland S, Ripamonti C. 2004. *Computational Colour Science using MATLAB.* Chichester, West Sussex: John Wiley & Sons Ltd.

Westmoreland D, Kiltie RA. 1996. Egg crypsis and clutch survival in three species of blackbirds (Icteridae). *Biological Journal of the Linnean Society* **58:** 159-172.

Windig JJ. 1991. Quantification of Lepidoptera wing patterns using an image analyzer. *Journal of Research on the Lepidoptera* **30:** 82-94.

Wyszecki G, Stiles WS. 1982. *Color Science: Concepts and Methods, Quantitative Data and Formulae, 2nd edition.* New York: John Wiley.

Yin J, Cooperstock JR. 2004. Color correction methods with applications for digital projection environments. *Journal of the Winter School of Computer Graphics* **12:** 499-506.

Zuk M, Decruyenaere JG. 1994. Measuring individual variation in colour: a comparison of two techniques. *Biological Journal of the Linnean Society* **53:** 165-173.

FIGURES

Figure 1

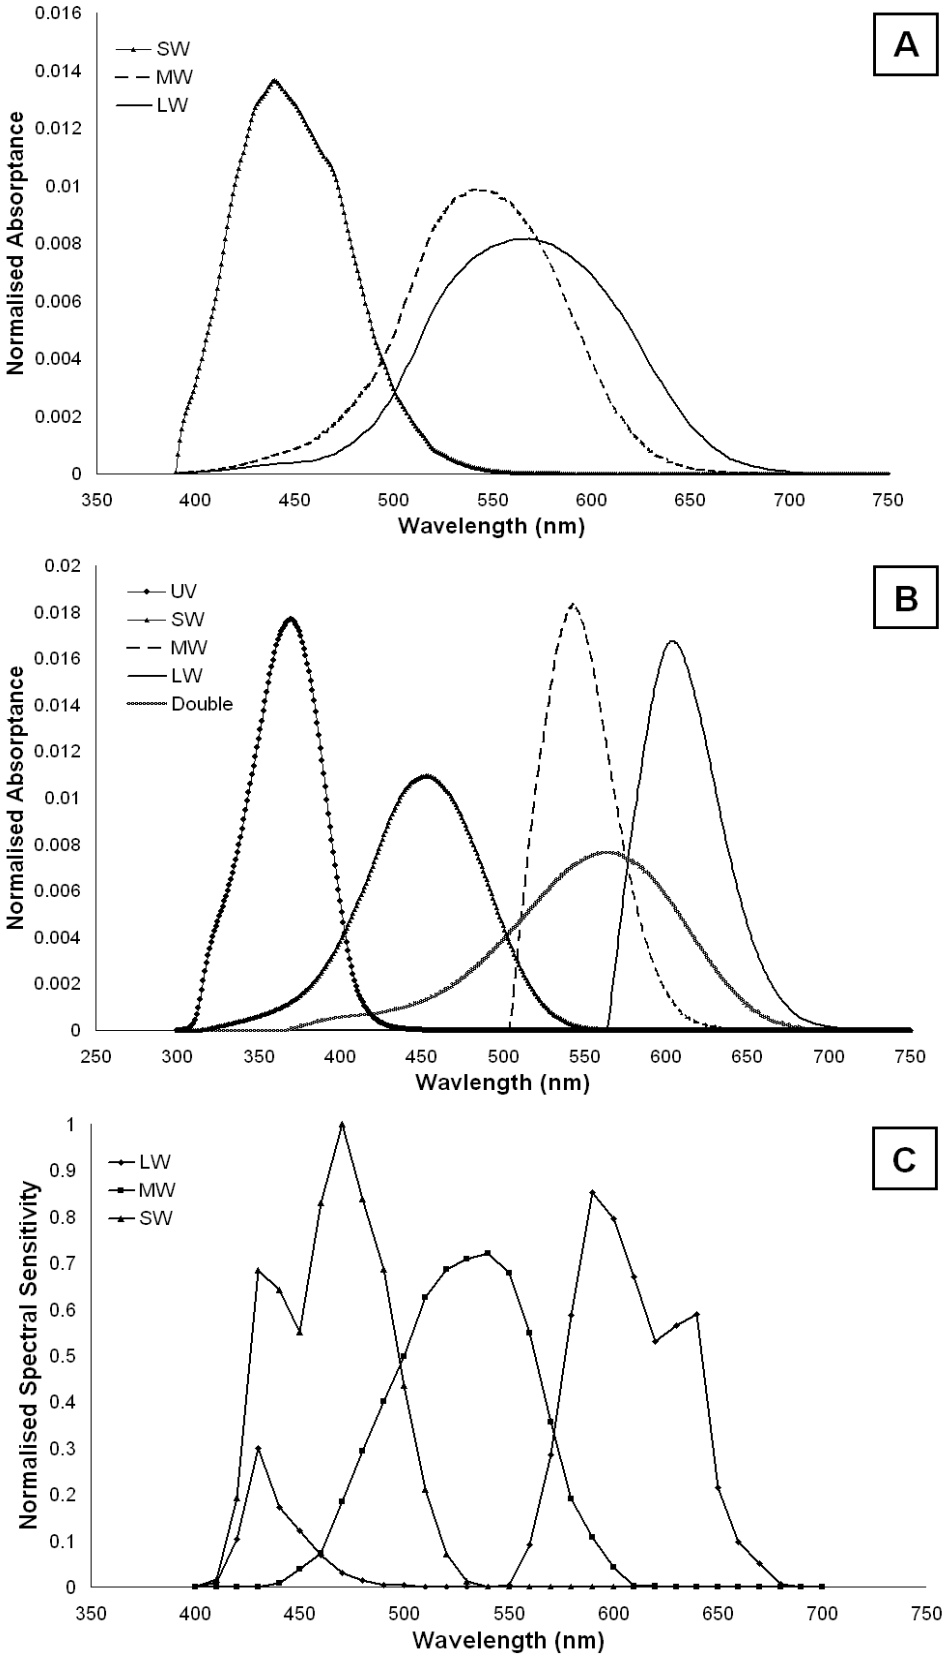

Figure 2

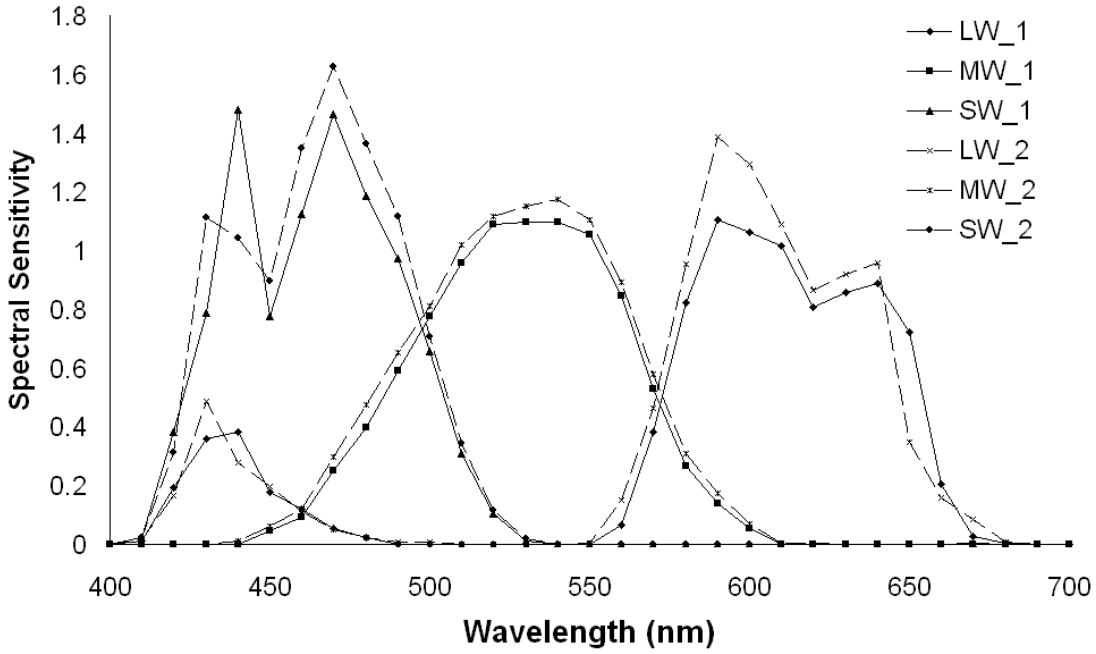
Spectral Sensitivity
1.8
1.6
1.4
1.2
1
0.8
0.6
0.4
0.2
0
400
450
500
550
600
650
700
Wavelength (nm)
LW_1
MW_1
SW_1
LW_2
MW_2
SW_2

Figure 3

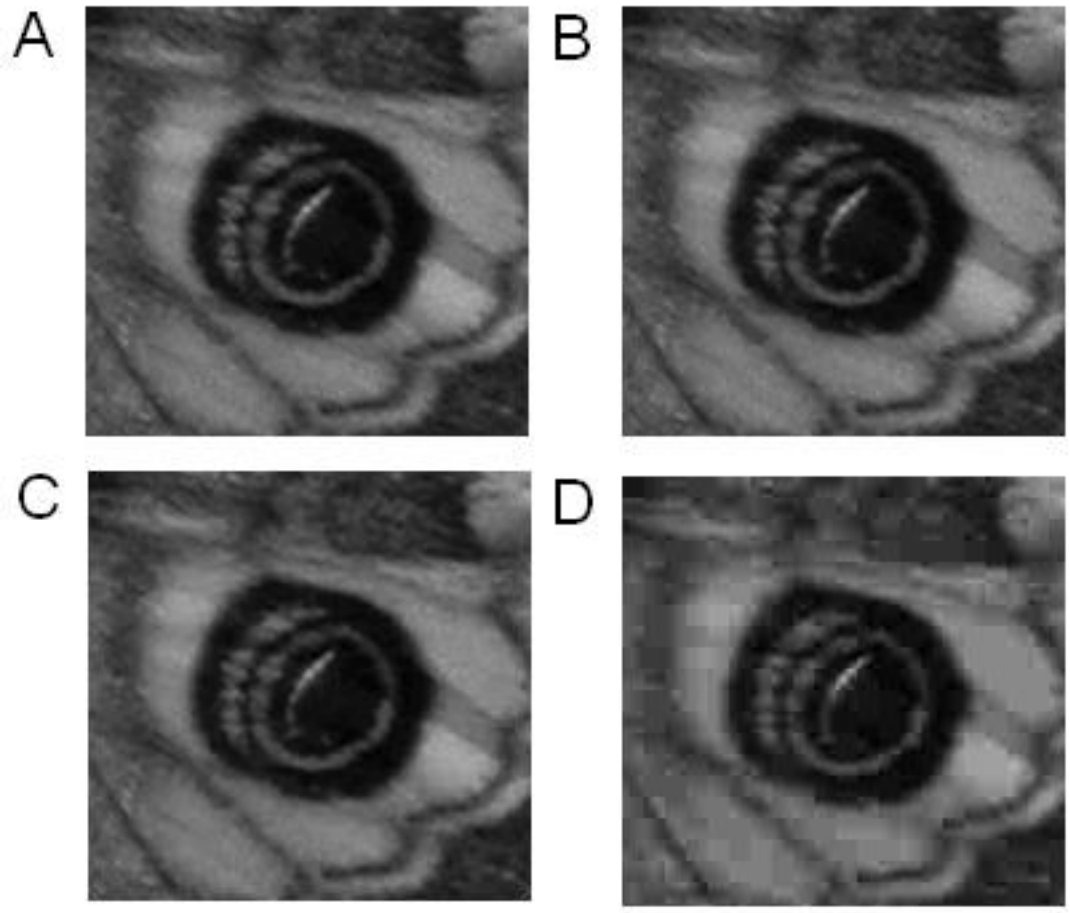

Figure 4

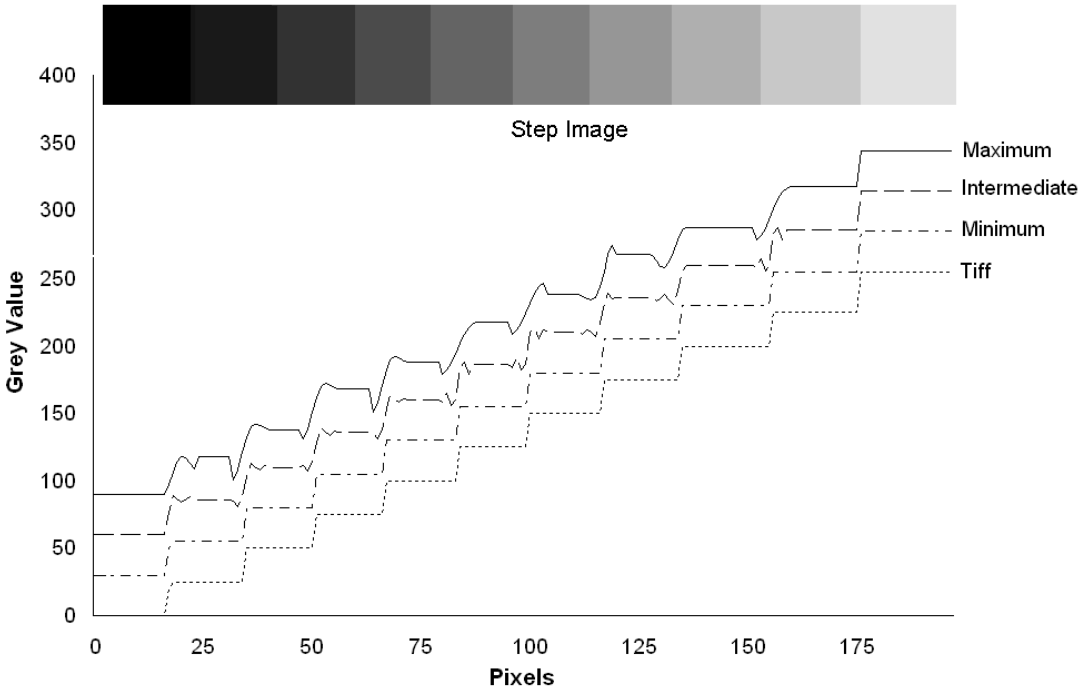
Step Image
Grey Value
400
350
300
250
200
150
100
50
0
0
25
50
75
100
125
150
175
Pixels
Maximum
Intermediate
Minimum
Tiff

Figure 5

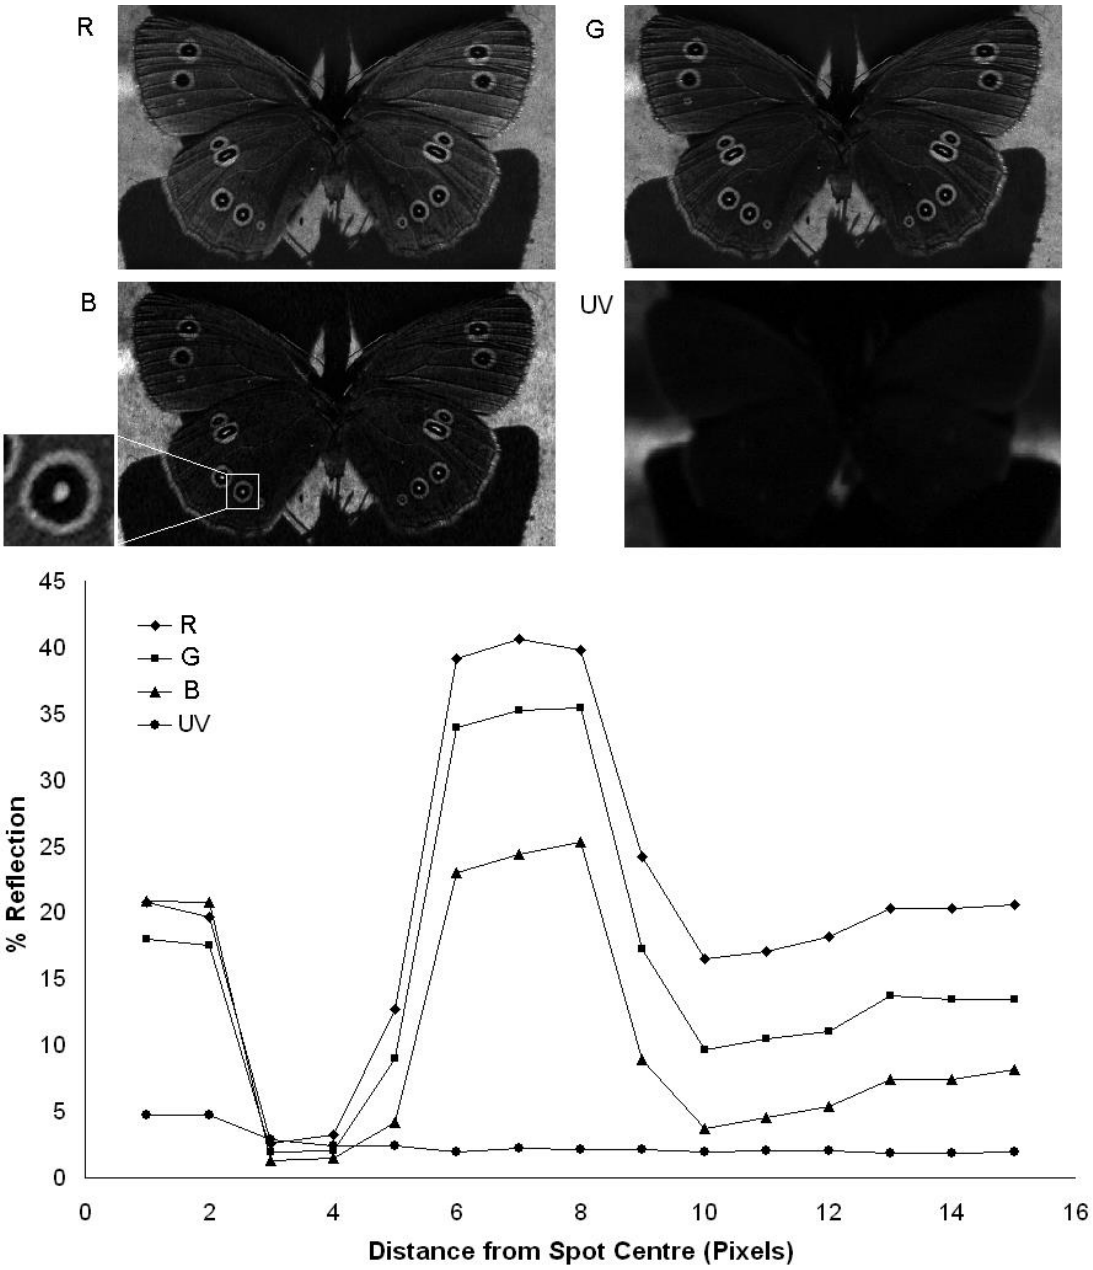
R
G
B
UV
45
40
35
30
25
20
15
10
5
0
% Reflection
R
G
B
UV
0
2
4
6
8
10
12
14
16
Distance from Spot Centre (Pixels)

Figure 6

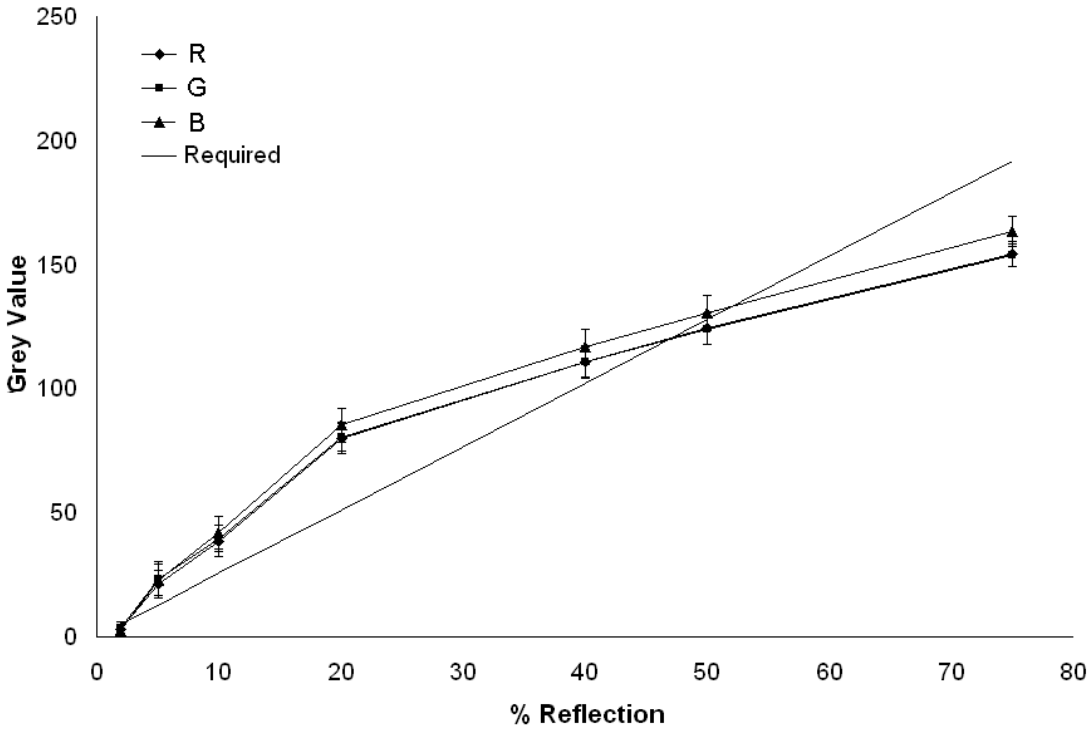
R
G
B
Required
Grey Value
% Reflection
0
50
100
150
200
250
0
10
20
30
40
50
60
70
80

Figure 7

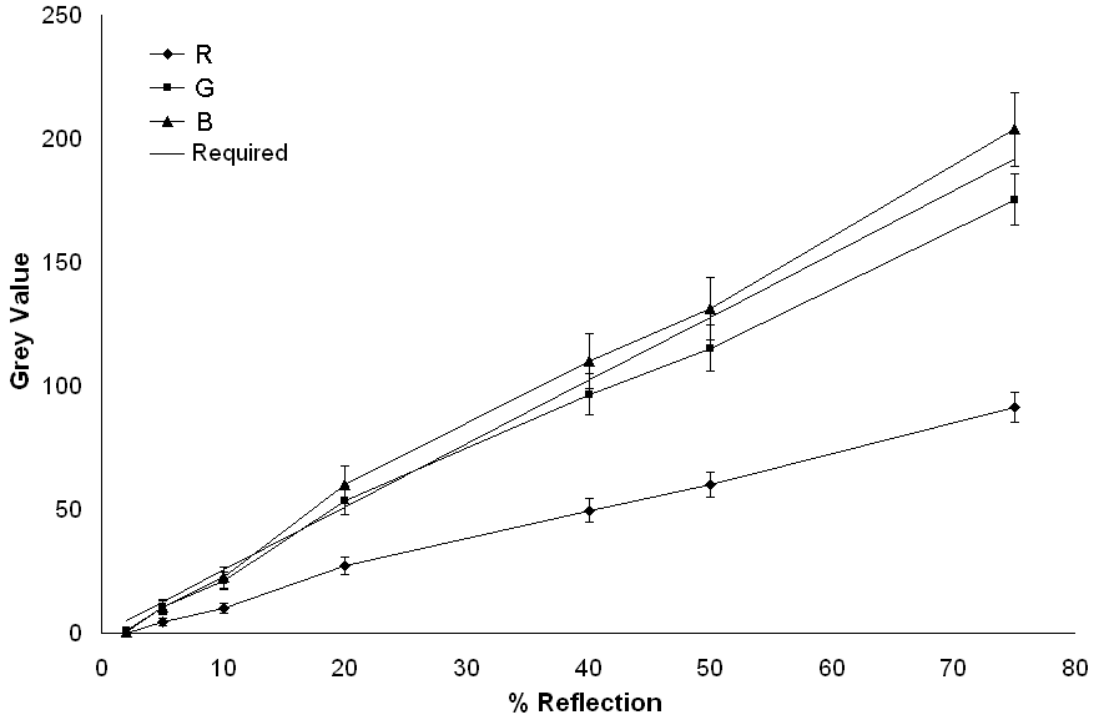
250
200
150
100
50
0
Grey Value
R
G
B
Required
0
10
20
30
40
50
60
70
80
% Reflection

Figure 8

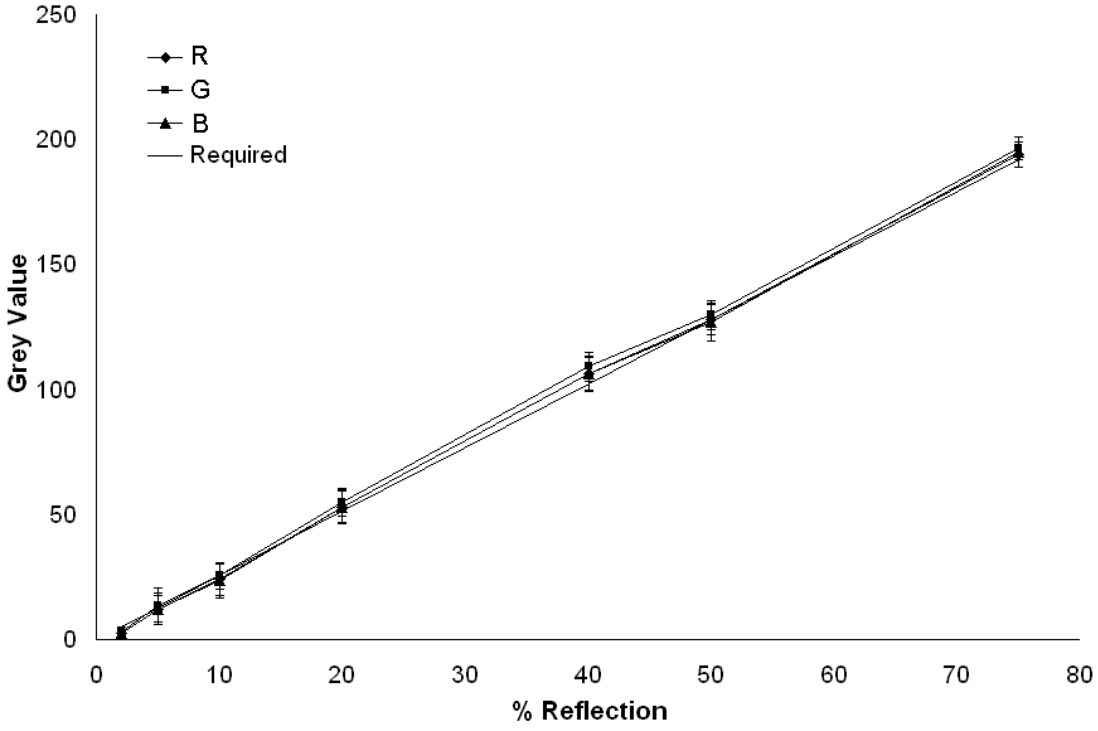
250
200
150
100
50
0
Grey Value
R
G
B
Required
0
10
20
30
40
50
60
70
80
% Reflection

Figure 9

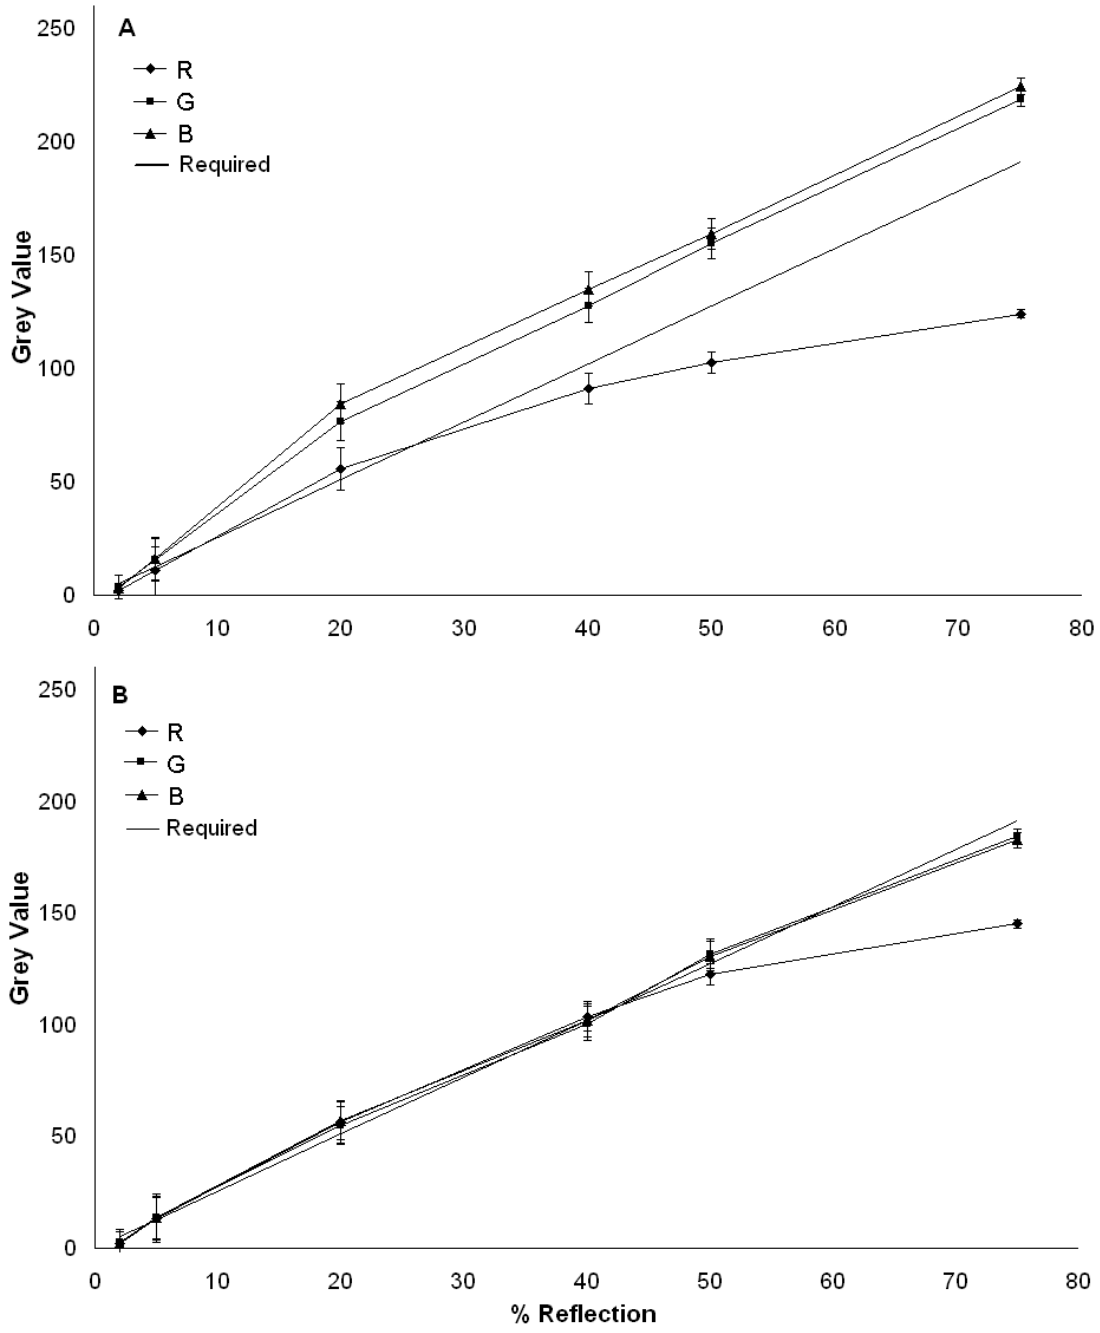
A
R
G
B
Required
Grey Value
250
200
150
100
50
0
0
10
20
30
40
50
60
70
80
B
R
G
B
Required
Grey Value
250
200
150
100
50
0
0
10
20
30
40
50
60
70
80
% Reflection

Figure 10

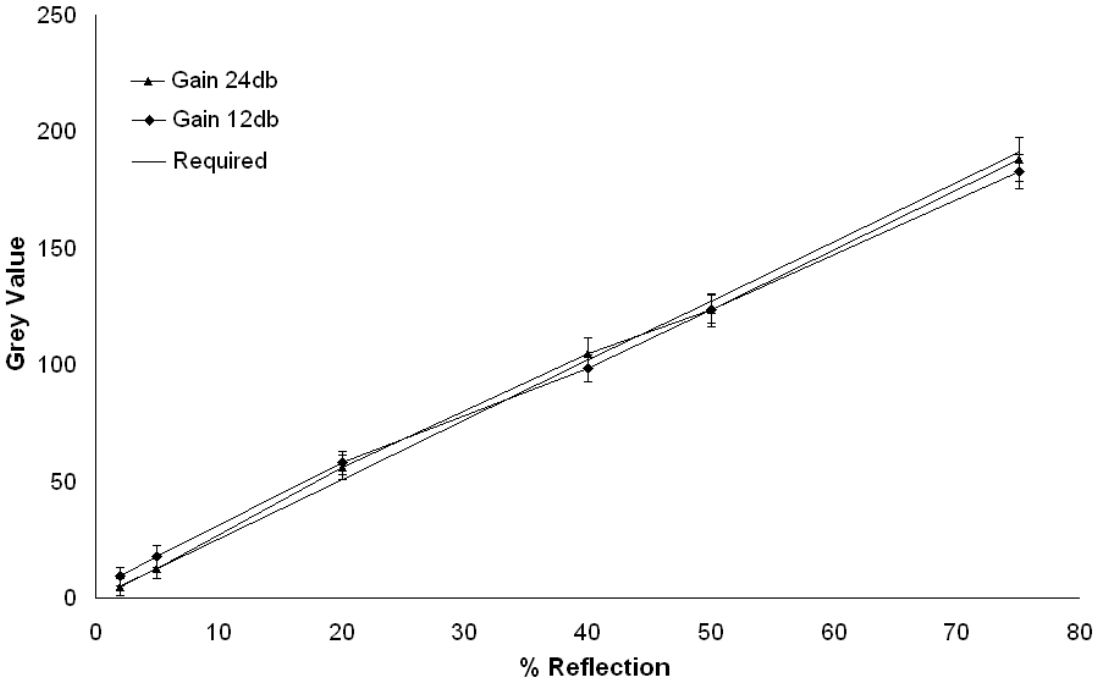
250
200
150
100
50
0
Grey Value
Gain 24db
Gain 12db
Required
0
10
20
30
40
50
60
70
80
% Reflection

Figure 11

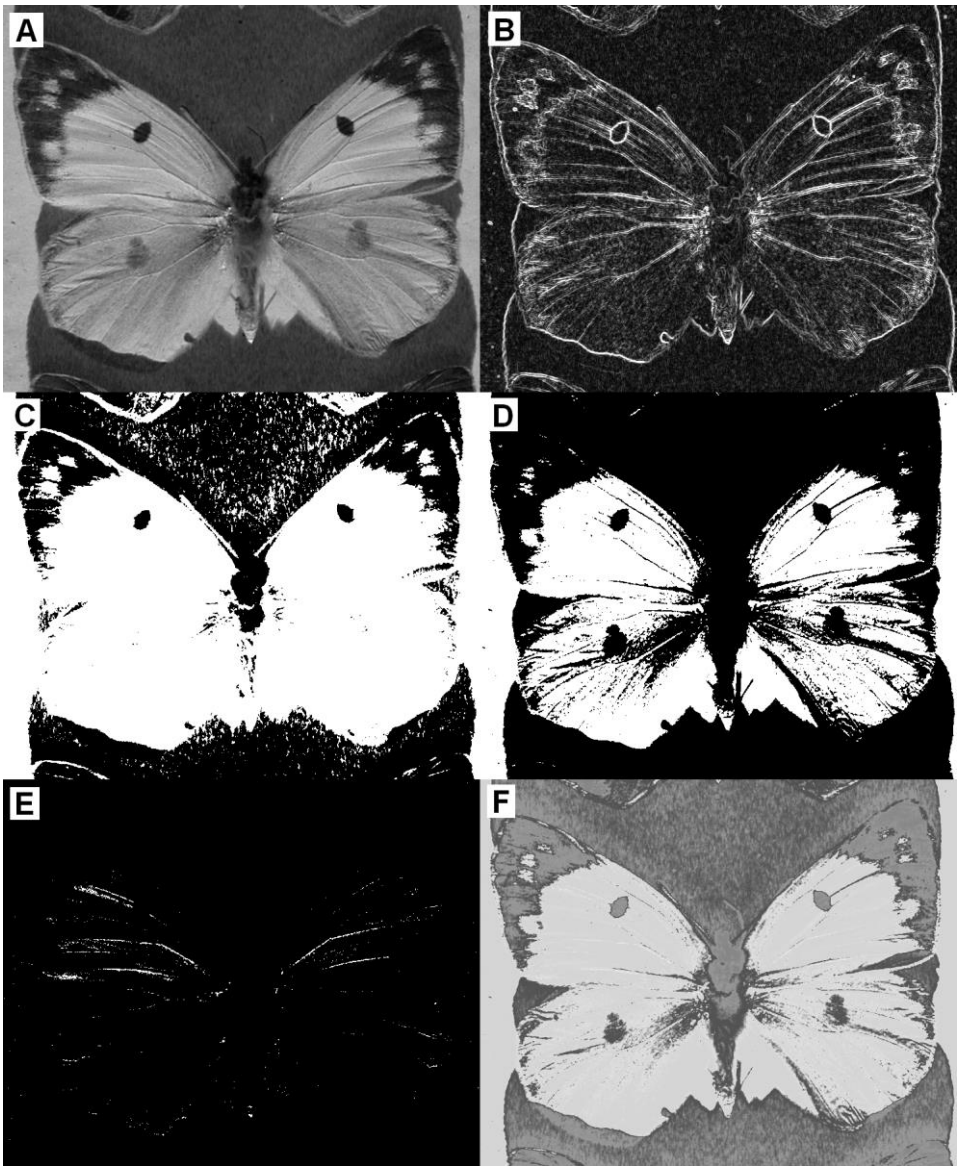

Figure Legends

Figure 1: (**A**) Normalised absorptance (equal areas under curves) of human cones. Absorbance (*N*) data from Dartnall, Bowmaker & Mollon (1983) converted to absorptance (*P*) by the equation $P = 1\text{-}10^{-1.N.L.S}$, where *L* is the length of the cone (20 μm from Hendrickson & Drucker, 1992), and *S* is specific absorbance, $0.015/\mu m^{-1}$. (**B**) Normalised absorptance (equal areas under curves) of starling cones to different wavelengths of light. From Hart, Partridge & Cuthill (1998). (**C**) Normalised spectral sensitivity (equal areas under curves) of the sensors in the Nikon 5700 Coolpix camera used in our study.

Figure 2: A plot of spectral sensitivity of two Nikon 5700 cameras for the LW, MW and SW channels. Even though the cameras are the same make and model, and were purchased simultaneously, there are some (albeit relatively small) differences in spectral sensitivity.

Figure 3: Four images of the hind left spot on the emperor moth *Saturnia pavonia* illustrating the effects of compression on image quality. (**A**) An uncompressed TIFF image of the original photograph. (**B**) A JPEG image with minimal compression (10%). (**C**) A JPEG image with intermediate compression (50%), which still appears to maintain the original structure of the image, but careful examination of the image's spatiochromatic content shows inconsistencies with the original TIFF file. (**D**) A JPEG image with maximal compression (90%) showing severe spatial and chromatic disruption.

Figure 4: Grey values measured when plotting a transect across a grey scale step image with increasing values from left to right. Grey values start at 0 on the left of the series of steps and increase in steps of 25 to reach values of 250 on the right. Plotted on the graph are the values measured for images of the steps as an uncompressed TIFF file, and JPEGs with 'minimum' (10%), 'intermediate' (50%) and 'maximum' (90%) levels of compression. Values of 30, 60 and 90 have been added to the JPEG files with minimum, intermediate and maximum levels of compression to separate the lines vertically. Notice as the level of compression increases the data measured are more severely disrupted, particularly at the boundary between changes in intensity. In the case of complex patterns, the disruption to the image structure means that measurements at any point in the image will be error prone.

Figure 5: Results from a radial profile analysis performed upon one eyespot of the ringlet butterfly *Aphantopus hyperantus*, illustrating the high percentage reflectance values obtained for the centre of the spot and the 'golden' ring further from the centre, particularly in the red and green channels, and the lack of an eyespot in the UV.

Figure 6: The relationship between the grey scale value measured for a set of seven Spectralon reflectance standards from raw digital TIFF file images and the nominal reflection value, showing a curved relationship for the R, G and B data. The 'required' line illustrates values that should be measured if the camera's response was linear and the three channels equally stimulated.

Figure 7: The relationship between measured greyscale value and nominal reflection value for the seven reflectance standards, showing the linearisation of the gamma curves.

Figure 8: The greyscale values measured for the set of reflectance standards following the process of RGB channel equalisation and scaling, showing a close fit to the required values.

Figure 9: The greyscale values measured for the set of reflectance standards following the process of linearisation (**A**) and then RGB channel equalisation (**B**) and scaling, showing that the linearisation does not produce a linear response when there are saturated pixels in the image, as is the case in the R channel in this example. Saturated pixels also result in a poor equalisation result, indicated by a dropping off of the R channel at higher values.

Figure 10: The effect of scaling the UV images obtained with the PCO Variocam camera and Nikon UV transmitting lens, showing a close fit to the required values.

Figure 11: Different images of a clouded yellow butterfly *Colias croceus*, modified to show regions of interest, such as the wing spots, identified by various techniques. **(A)** The original 8-bit grey-level image (pixel values between 0 and 255). **(B)** The image after an edge detection algorithm has been applied, clearly identifying a boundary around the two forewing spots, but not the hindwing spots. **(C)** The original image after being thresholded to a binary (black/white) image with a threshold of 64. This clearly shows the forewing spots but does not produce spots where the hindwing spots were in the original image. **(D)** The original image when converted to a binary image with a threshold of 128, now picking out both the forewing and hingwing spots (although with some 'noise' around the hindwing spots). **(E)** The original image converted to a binary image with a threshold of 192, now not showing any clear wing

spots. **(F)** The original image when first converted to a pseudocolour image, where each pixel value falling between a given range is given a specific colour. The image is then reconverted to a grey-level image and now shows the hindwing spots with marginally sharper edges than in the original image.

Table 1: Desirable characteristics when purchasing a digital camera for research

Attribute	Relative Importance
High resolution (e.g. min 5 megapixels)	High
Manual white balance control	High
Macro lens	Medium
Ability to save Tiff/RAW file formats	High
Manual exposure control	High
Remote shutter release cable capability	Low
Ability to change metering method	Medium
Optical zoom	Medium

Appendix 1: Glossary of technical terms

Aliasing

Aliasing refers to the jagged appearance of lines and shapes in an image.

Aperture

Aperture refers to the diaphragm opening inside a photographic lens. The size of the opening regulates the amount of light passing through onto the colour filter array. Aperture size is usually referred to in f-numbers. Aperture also affects the ‘depth of field’ of an image.

Charge-coupled device (CCD)

A small photoelectronic imaging device containing numerous individual light-sensitive picture elements (pixels). Each pixel is capable of storing electronic charges created by the absorption of light and producing varying amounts of charge in response to the amount of light they receive. This charge is converts light into electrons, which pass through an analogue-to-digital converter, which produces a file of encoded digital information.

Chromatic aberration

This is caused by light rays of different wavelengths coming to focus at different distances from the lens causing blurred images. Blue will focus at the shortest distance and red at the greatest distance.

Colour filter array (CFA)

Each pixel on a digital camera sensor contains a light sensitive photodiode which measures the brightness of light. These are covered with a pattern of colour filters, a colour filter array, to filter out different wavebands of light.

Demosaicing algorithms

Most digital cameras sample an image with red, green and blue sensors arranged in an array, with one type at each location. However, an image is required with an R, G and B value at each pixel location. This is produced by interpolating the missing sensor values via so called ‘demosaicing’ algorithms, which come in many types.

Exposure

The exposure is the amount of light received by the camera’s sensors and is determined by the aperture and the integration time.

Foveon sensors

Foveon sensors capture colour by using three layers of photosensors at each location. This means that no interpolation is required to obtain values of R, G, & B at each pixel.

Image resolution

The resolution of a digital image is the number of pixels it contains. A 5 megapixel image is typically 2,560 pixels wide and 1,920 pixels high and has a resolution of 4,915,200 pixels.

JPEG

JPEG (Joint Photographic Experts Group) is very common due to its small size and widespread compatibility. JPEG is a lossy compression method, designed to save storage space. The JPEG algorithm divides the image into squares, which can be seen on badly-compressed JPEGs. Then a Discrete Cosine Transformation is used to turn the square data into a set of curves, and throws away the less significant part of the data. The image information is rearranged into colour and detail information, compressing colour more than detail because changes in detail are easier to detect. It also sorts detail information into fine and coarse detail, discarding fine detail first.

Lossy compression

A data compression technique in which some data is lost. Lossy compression attempts to eliminate redundant or unnecessary information and dramatically reduces the size of a file by up to 90%. Lossy compression can generate artifacts such as false colours and blockiness. JPEG is an image format that is based on lossy compression.

Lossless compression

Lossless compression is similar to ‘zipping’ a file, whereby if a file is compressed and later extracted, the content will be identical. No information is lost in the process. TIFF images can be compressed in a lossless way.

Macro lens

A lens that provides continuous focusing from infinity to extreme close-ups.

Modulation transfer function

The modulation transfer function describes how much a piece of optical equipment, such as a lens, blurs the image of an object. Widely spaced features, such as broad

black and white stripes, do not lose much contrast, since a little blurring only affects their edges, but fine stripes may appear to be a uniform grey after being blurred by the optical apparatus. The modulation transfer function is a measure of how much bright-to-dark contrast is lost, as a function of the width of the stripes, as the light goes through the optics.

Nyquist frequency

The Nyquist frequency is the highest spatial frequency where the CCD can still correctly record image detail without aliasing.

Bit depth

This relates to image quality. A bit is the smallest unit of data, such as 1 or 0. A 2 bit image can have 2^2 = 4 grey levels (black, low grey, high grey and white). An 8 bit image can have 2^8 = 256 grey levels, ranging from 0 to 255. Colour images are often referred to as 24 bit images because they can store up to 8 bits in each of the 3 colour channels and therefore allow for 256 x 256 x 256 = 16.7 million colours.

RAW

A RAW file contains the original image information as it comes off the sensor before internal camera processing. This data is typically 12 bits per pixel. The camera's internal image processing software or computer software can interpolate the raw data to produce images with three colour channels (such as a TIFF image). RAW data is not modified by algorithms such as sharpening. RAW formats differ between camera manufacturers, and so specific software provided by the manufacturer, or self written software, has to be used to read them.

Saturation

In the context of calibrating a digital camera, we use this term to mean when a sensor reaches an upper limit of light captured and can no longer respond to additional light. This is also called ‘clipping’ as the image value cannot go above 255 (in an 8-bit image) regardless of how much additional light reaches the sensor. Saturation can also be used to refer to the apparent amount of hue in a colour, with saturated colours looking more vivid.

Sensor resolution

The number of effective non-interpolated pixels on a sensor. This is generally much lower than the image resolution as this is before interpolation has occurred.

TIFF

TIFF (Tagged Image File Format) is a very flexible file format. Tiffs can be uncompressed, lossless compressed, or can be lossy compressed. While JPEG images only support 8 bits per channel RGB images, TIFF also supports 16 bits per channel and multi-layer CMYK images in PC and Macintosh format.

White balance

Most digital cameras have an automatic white balance setting whereby the camera automatically samples the brightest part of the image to represent white. However, this automatic method is often inaccurate and is undesirable in many situations. Most digital cameras also allow you to choose a white balance manually.

Appendix 2: Technical details

In our study we used a Nikon Coolpix 5700 camera, with an effective pixel count of just under 5.0 megapixels. This does not have all the desired features described in our article (the intensity response is non-linear and the zoom cannot be precisely fixed) and we offer no specific recommendation, but it is a good mid-priced product with high quality optics and full control over metering and exposure. UV photography was with a PCO Variocam, fitted with a Nikon UV-Nikkor 105 mm lens, a Nikon FF52 UV pass filter and an Oriel 59875 "heat" filter (the CCD is sensitive to near-infra-red). The camera was connected to a Toshiba Satellite 100cs laptop and also to an Ikegami PM-931 REV.A monitor, which displayed the images that were to be saved via a PCO CRS MS-DOS based programme. With the camera remote control, the gain and the integration time of the images could be adjusted, with the gain either set to 12db or 24db and the integration time between 1 and 128 video frames (1 frame = 40 milliseconds).

Images were transferred to a PC and all measurements were taken with the (free) imaging programme 'Image J' (Rasband, 1997-2006; Abràmoff *et al.*, 2004). Measurements of standards were taken by drawing a box over the area of interest, and then using the histogram function to determine the mean grey scale value and standard deviation for each channel. All other image and data manipulations, including the linearization and transformation between coordinate systems, were performed with MATLAB (The Mathworks Inc. Massachusetts, USA), though other languages, such as Java (Sun Microsystems, Inc. Santa Clara, USA; Efford, 2000) are also useful. MATLAB has rapidly become an industry standard in vision science, on account of its efficiency at matrix mathematics and manipulation (photographic data are large

matrices). MATLAB and Image J benefit from the large number of plug-ins and toolboxes written by users for other users.