

Using digital cameras to investigate animal colouration: estimating sensor sensitivity functions

Thomas W. Pike

Received: 26 April 2010 / Revised: 20 October 2010 / Accepted: 22 October 2010 / Published online: 10 November 2010
© Springer-Verlag 2010

Abstract Spectrophotometers allow the objective measurement of colour and as a result are rapidly becoming a key piece of equipment in the study of animal colouration; however, they also have some major limitations. For example, they can only record point samples, making it difficult to reconstruct topographical information, and they generally require subjects to be inanimate during measurement. Recently, the use of digital cameras has been explored as an alternative to spectrophotometry. In particular, this allows whole scenes to be captured and objectively converted to animal colour space, providing spatial (and potentially temporal) data that would be unobtainable using spectrophotometry; however, mapping between camera and animal colour spaces requires knowledge of the spectral sensitivity functions of the camera's sensors. This information is rarely available, and making direct measures of sensor sensitivity can be prohibitively expensive, technically demanding and time-consuming. As a result, various methods have been developed in the engineering and computing sciences that allow sensor sensitivity functions to be estimated using only readily collected data on the camera's response to a limited number of colour patches of known surface reflectance. Here, I describe the practical application of one such method and demonstrate how it allows the recovery of sensor sensitivities

(including in the ultraviolet) with a high enough degree of accuracy to reconstruct whole images in terms of the quantal catches of an animal's photoreceptors, with calculated values that closely match those determined from spectrophotometric measurements. I discuss the potential for this method to advance our understanding of animal colouration.

Keywords Spectral sensitivity · Quadratic programming · Digital cameras · Animal vision · Visual ecology · Colour signals

Introduction

The evolution and functional significance of animal colouration has been a major theme within behavioural ecology and evolutionary biology for over a century (e.g. Poulton 1890; Beddard 1892). Animals employ colouration for a wide variety of different functions, including sexual signalling, social communication (e.g. Hill and McGraw 2006), warning colouration, crypsis and mimicry (Ruxton et al. 2004); however, understanding variation in animal colouration requires an effective means of quantification. Until fairly recently, studies tended to use (necessarily subjective) human comparisons to assess variation in colouration, for example by employing arbitrary categorical rankings (e.g. bright vs. dull) or by matching perceived colouration to colour 'standards' (Endler 1990). It is now generally appreciated that such comparisons are affected by a number of factors, including the illumination conditions under which the comparison takes place and variation between (and even within) observers. Even apparently less subjective comparisons performed using computer software are not necessarily immune to these limitations, and may also suffer from variation among recording devices and in

Communicated by K. McGraw

Electronic supplementary material The online version of this article (doi:10.1007/s00265-010-1097-7) contains supplementary material, which is available to authorized users.

T. W. Pike (✉)
Centre for Ecology and Conservation, School of Biosciences,
University of Exeter (Cornwall Campus),
Penryn TR10 9EZ, UK
e-mail: t.pike@exeter.ac.uk

the algorithms used in the colour analysis (e.g. Stevens and Cuthill 2005).

To address these issues, researchers are increasingly turning towards objective means of assessing colour. As a result, spectrophotometers are rapidly becoming a standard piece of equipment in studies of animal colouration, as spectrophotometric data can be used to compute objective, standardised colour metrics (Endler 1990). Increasingly, this information is being incorporated into psychophysical models of receivers' visual systems to explore how colours may be perceived by non-human animals, which is a technique that can yield significant insights into the evolution of animal colouration (Endler and Mielke 2005); however, the use of spectrophotometry has some major limitations (discussed in Stevens et al. 2007). Conventional spectrophotometers provide only point samples, so if the study is concerned with measuring the relationships between multiple colour patches this topographical information must be reconstructed from a limited number of discrete measurements (Endler and Mielke 2005). Spectrophotometry is also a fairly slow process and is really practical for use only with restrained or inanimate animals (or at the very least those that are within a short distance of the measuring equipment; e.g. Rush et al. 2003), collected samples (e.g. feathers or museum specimens) and in situations where the colours being measured are not subject to rapid change. It also generally requires a constant light source, and the necessary equipment is often expensive and impractical for field use.

To overcome some of these limitations, the use of digital photography is increasingly being explored as an alternative to spectrophotometry in the study of animal colouration. While photography has been used for some years to quantify variation in animal colouration (e.g. Frischknecht 1993), there has recently been a rapid increase in the number of studies employing this method to explore a wide variety of questions, from sexual signalling to camouflage to egg colouration (e.g. Barbosa et al. 2008; Chiao et al. 2009; Stoddard and Stevens 2009; Higham et al. 2010), principally because of the number of benefits that digital photography can bring. It is considerably faster than spectrophotometry, meaning large quantities of data can be rapidly collected. Subjects do not need to be restrained, and images can be taken of animals performing natural behaviours even when the subject is some distance away. The equipment is relatively inexpensive and often more ergonomic than spectrophotometers for field-based studies. Entire spatial patterns can be analysed, without the need to reconstruct (a necessarily depauperate) topography from point samples. Powerful image processing algorithms and software are available to analyse whole images or quickly assess spatial patterns (e.g. the size and distribution of colour markings), and video imaging can provide a

temporal dimension. Readers are referred to the excellent review by Stevens et al. (2007), which describes the necessary equipment and procedures, as well as the limitations and precautions of this technique.

Because colour is a function of the perceptual system of the intended receiver, Bennett et al. (1994) stress the importance of quantifying colour as it is perceived. One of the most exciting applications of digital cameras in the study of animal colouration is the ability to map the output of the camera's sensors to the photoreceptor quantal catches of the species under study in order to produce images that represent whole scenes in terms of an animal's visual sensitivity, which often differs markedly from our own (Endler and Mielke 2005). However, mapping between device and animal colour spaces generally requires knowledge of the spectral sensitivity curves of the camera's sensors (Stevens et al. 2007); because these data are rarely known a priori (manufacturers do not usually publish this information, and there are often differences even between cameras of the same model; Stevens et al. 2007) and empirical methods for their determination are often extremely time-consuming, technically challenging and require costly equipment (Hardeberg et al. 1998; Párraga 2003), this represents a major stumbling block for most ecologists. However, this problem has received considerable attention in the engineering and computing sciences, where a number of techniques have been developed to accurately estimate a camera's sensor sensitivities from readily obtainable digital images of a limited number of colour patches of known reflectance (e.g. Farrell and Wandell 1993; Sharma and Trussell 1993, 1996; Finlayson et al. 1998; Hardeberg et al. 1998; Thomson and Westland 2001; Alsam and Lenz 2007; Ebner 2007). Here, I focus on the quadratic programming method developed by Finlayson et al. (1998), which has been shown to allow the recovery of sensor curves that are an extremely close match to the actual curves, and demonstrate its applicability to studies of animal colouration.

Methods

Determination of camera spectral sensitivity

1. Theory

It is assumed that a camera has i sensors (this is usually three, such as the R, G and B channels in a typical colour camera, but may be only one in a monochromatic device). The output r of the i th sensor is calculated as

$$r_i = \sum_{k=1}^K C(\lambda_k) R_i(\lambda_k), \quad (1)$$

where C is the radiance of the imaged object (typically calculated as the product of the surface reflectance and

the spectral power distribution of the scene illuminant) and R is the spectral sensitivity of the i th sensor (taking into account the transmission properties of the lens and any proprietary filters, such as infrared blocking filters, placed over the sensor), over all K wavelengths of interest (for instance, $\lambda_1=400$ nm, $\lambda_2=410$ nm, $\lambda_3=420$ nm, ..., $\lambda_K=700$ nm).

It follows that, given C and r (which can be determined empirically), it should be possible to estimate R . This process is made theoretically easier because the spectral sensitivity curves of any camera's sensors are likely to exhibit a number of features that limit the range over which these sensitivities can be estimated (Sharma and Trussell 1993; Barnard 1995; Finlayson et al. 1998). Full details of these constraints are given in Finlayson et al. (1998), and a brief overview is given here. First, the sensors must be positive across all wavelengths, as it is impossible for sensors to have a negative response to a stimulus. That is

$$R_i(\lambda_k) \geq 0. \quad (2)$$

Second, the number of peaks in the sensor curve is typically limited to relatively few; assuming one, two or possibly three peaks will generally cover most sensor types. For example, a unimodal curve that has its peak at the m th sample point can be expressed as

$$\begin{aligned} R_i(\lambda_{k+1}) &\geq R_i(\lambda_k), & k = 1, \dots, m-1 \\ R_i(\lambda_{k+1}) &\leq R_i(\lambda_k), & k = m, \dots, K, \end{aligned} \quad (3)$$

and similar expressions can be written for n -modal curves. Third, sensor sensitivity curves tend to be smooth. This constraint can be incorporated by representing a sensor curve as a linear combination of a set of band-limited basis functions. For instance, a standard Fourier basis (e.g. Finlayson et al. 1998) represents the sensor curve as a linear combination of L weighted sine and cosine functions:

$$R_i(\lambda_k) = \sigma_1 B_1(\lambda_k) + \sigma_2 B_2(\lambda_k) + \dots + \sigma_L B_L(\lambda_k), \quad (4)$$

where the first few basis functions are $B_1=c$, $B_2=\sin(x)$, $B_3=\cos(x)$, $B_4=\sin(2x)$ and so on, for constant c and $x=(\lambda_k-\lambda_1)\pi/0.5(\lambda_K-\lambda_1)$, and σ_l is the vector of weights for the basis functions. Equation (1) can therefore be rewritten as

$$r_i = \sum_{k=1}^K \left(C(\lambda_k) \sum_{l=1}^L \sigma_l B_l(\lambda_k) \right), \quad (5)$$

and the problem now becomes to find σ_l .

As discussed by Finlayson et al. (1998), these three constraints naturally form a series of linear inequalities allowing the problem to be formulated as a quadratic

programming routine, subject to those constraints, in which we attempt to minimise the error of the sensor sensitivity estimate (for example, by minimising the quadratic objective function [Finlayson et al. 1998] or the sum of squares relative error between predicted and observed sensor outputs [Barnard 1995]). To estimate the sensitivity curve of a sensor, we require as input an $N \times K$ matrix C of N training spectra measured at K wavelength intervals, an associated vector of N outputs from the i th sensor r_i , the number of basis functions to use for the smoothness constraint and the estimated locations of peaks (and troughs, if any) in the resulting curve. Full details of using the quadratic programming procedure to estimate sensor sensitivities are given in Finlayson et al. (1998); the following section describes how the routine can be parameterised empirically.

2. Parameterisation of the model

The camera used in the example calibration was a Nikon D70 digital SLR fitted with a 105-mm quartz lens (UKA optics, Oyster Bay, NY) and an ultraviolet (UV)–infrared blocking filter (Baader, Mammendorf, Germany) to restrict the available range of wavelengths, and hence the range over which sensor sensitivity could be estimated, to between approximately 400 and 700 nm. This was because the camera was found to have a significant response outside these wavelengths (personal observation).

The first stage in the recovery of sensor sensitivity is to obtain the response of the camera to a number of objects of known surface reflectance, under a known illuminant. Here, I used the 24 patches of the Gretag-Macbeth ColorChecker chart (McCamy et al. 1976), a standard colour chart consisting of 24 natural colours under illumination from a constant incandescent light source (which has a smooth spectral distribution—illumination sources with sharp peaks in the spectral distribution, such as fluorescent lamps, should be avoided; Romero et al. 2003; Solli et al. 2005). The reflectance spectra of each of the 24 colour patches was measured against a Spectralon white standard using an Ocean Optics USB2000 spectrophotometer (Ocean Optics, Dunedin, FL), and the irradiance of the illuminant was measured with the same device. These values were used in subsequent calculations, although it was found that using published reflectance values (e.g. Kohonen et al. 2006) provided an adequate estimation of sensor sensitivity. The product of the reflectance and irradiance spectra provided an estimate of the radiance of each colour patch, and the resulting spectra were converted to 1-nm intervals between 400 and 700 nm. The resulting 24×301 matrix (i.e. 24 colour patches, each measured over 301 discrete wavelengths) constituted the matrix C (see above) (Note that although this

wavelength range was used for the initial sensor estimations, it was found that a better fit for the blue sensor could be obtained by extending the range down to 300 nm, even though the colour patches have very low reflectance below about 400 nm).

Following spectrophotometric measurements, the chart was photographed under the incandescent light with exposure time adjusted so that none of the colour channels became saturated (Stevens et al. 2007), and none of the pixel values fell within the top or bottom 2% of the dynamic range. Captured images were saved in Nikon's proprietary RAW format (NEF) and converted to 12-bit TIFFs (i.e. containing 12 bits in each of the three colour channels) using the open-source software DCRAW (<http://www.cybercom.net/~dcoffin/dcraw>). An approximately 30×30-pixel region in the centre of each colour patch was extracted, and the response of the R, G and B channels averaged to give three 24-element vectors of sensor outputs (\mathbf{r}_i). Many imaging devices, including the one used here, show some non-linearity in their sensors' response to variation in light intensity. Therefore, for each of the three colour channels, the relationship between the response to the six grey patches of the ColorChecker chart and the nominal reflectance values of these patches was nonlinear. Moreover, the precise relationship was different for each colour channel. Such non-linearity in sensor response can affect the estimation of their sensitivity functions, so before attempting to estimate sensor sensitivity I linearised the values of \mathbf{r}_i using the procedure described in detail by Westland and Ripamonti (2004) and Stevens et al. (2007).

In addition to radiance spectra and associated sensor outputs, the quadratic programming routine also requires an estimate of the locations of the peaks (m) and (if present) troughs in the sensitivity curve. While in theory the location of the sensor peak(s) are unknown, in practice they are likely to be restricted to particular regions of the spectrum. For example, the peak response of the long-wave sensor is likely to be in the long-range region of the spectrum (i.e. probably greater than 550 nm). By solving the equation for all plausible peak locations, it is possible to solve for the best sensor overall (i.e. the estimate with the lowest error; in this case the lowest quadratic objective function), and because the quadratic programming routine can be computed quickly, trying all plausible peak locations for putative uni-, bi- and tri-modal sensors is quite feasible (and can be readily automated). The number of Fourier basis functions used in the estimation will affect the smoothness of the resulting curve; in general, somewhere between nine and 15 basis functions have been found to be sufficient (Finlayson et al. 1998) and 15 were used in the example calibration presented here.

The quadratic programming routine was then run for each set of sensor outputs, using a custom Matlab (MathWorks, Natick, MA) function that utilised the QUADPROG function in the Optimization Toolbox (analogous functions also exist for other packages, such as in the QUADPROG library for R), to generate estimates of the spectral sensitivity curves for the R, G and B sensors. The MATLAB function (m-file) used to perform the quadratic programming procedure described in this paper is available as supplementary material (Online Resource 1), and a compiled version of the function is available on request from the author.

3. Extending sensor sensitivity estimation into the UV

In the above calibration procedure, calculation of sensor sensitivity was restricted to within the human visible range (400–700 nm); however, the sensors of many cameras will have a significant response outside of this range (although specialist lenses, such as those with quartz rather than glass optics, may be needed to transmit this light to the sensors). For example, the Nikon D70 used here has significant sensitivity in the infrared (in common with most digital cameras), but the sensors also have secondary peaks within the UV (personal observation). In principal, these sensor sensitivity peaks in the UV could be characterised by simply basing the calibration (as described above) on images of objects with significant reflectance below 400 nm as well as in the human visible range, and searching for the best-fit multi-peaked sensor response functions; however, when taking images of natural scenes, the sensors' response to the UV components of the scene are likely to be swamped by their (probably much greater) response in the human visible region. The UV components of a scene must therefore be captured through a UV pass filter (i.e. a filter that passes only UV radiation and blocks all visible light) (Stoddard and Stevens 2009). I, therefore, determined the camera's spectral sensitivity in the UV independent of sensitivity above 400 nm, using the following procedure.

Commercial colour charts typically reflect little UV radiation, so to determine the camera's UV sensitivity functions I repeated the calibration procedure described above but based on the sequential imaging of 24 custom-made colour targets that reflected light in a wavelength-specific manner between 300 and 400 nm. The targets were constructed by overlaying selected coloured filters (Lee Filters, Andover, UK) on a white diffuse reflectance standard (WS-1; Ocean Optics), which reflects strongly (>98%) and uniformly in the UV. Illumination was provided by a UVA-340 bulb (Q-Panel Lab Products, Cleveland, OH) and a Wood's glass bulb, and for the calibration the camera was fitted

with a UV pass filter (Venus U-Filter; Baader, Mammendorf, Germany).

Testing the accuracy of the calibration

1. Direct measurement of sensor sensitivities

By measuring the camera's responses to narrow-band illumination (that emits light only within a certain, narrow wavelength range) at regular wavelength intervals, it is possible to reconstruct the shape of the response curve. Here, I used the method described by Vora et al. (1997). Narrow-band stimuli were created using light from a xenon source (Ocean Optics PX-2) passed through a series of interference filters (10 nm half bandwidth, with peak transmittance ranging between 320 and 700 nm in 20-nm intervals; Knight Optical, Harrietsham, UK) and imaged onto a white diffuse reflectance standard (WS-1; Ocean Optics). Camera images of the illuminated reflectance standard were taken, as described above, at several different exposure settings to obtain the maximum non-saturated response, and therefore to minimise noise. I also measured the relative radiance of each narrow-band stimulus using the Ocean Optics USB2000 spectrophotometer. Both the camera and spectrophotometer were placed at similar geometric positions with respect to the reflectance standard. I extracted the R, G and B sensor responses from the camera images and averaged these over a 64×64-pixel rectangular section in the centre of each image. Crude sensor response curves were then estimated by calculating the sensitivity of each sensor at 20 discrete intervals across the spectrum (coinciding with the peak sensitivities of the narrow-band filters), using the 'simple' estimate of Hubel et al. (1994; see also Vora et al. 1997).

2. Comparing mathematically and empirically determined estimates of cone quantal catch

The principal aim of this work was to allow images to be manipulated so as to provide a representation of the imaged scene (typically taken under natural illumination conditions) in terms of the quantal catch of an animal's photoreceptors (Stevens et al. 2007). Therefore, to further test the accuracy of the sensor estimation procedure, I compared the empirically and mathematically determined estimates of cone quantal catch from a typical passerine bird, modelled on data available for the blue tit (*Cyanistes caeruleus*) (Hart et al. 2000). Blue tits have three single cones with peak sensitivities in the human visible range (i.e. above 400 nm), at 449 nm (S cone), 502 nm (M cone) and 563 nm (L cone). This method allows the accuracy of derived sensor sensitivity functions to be assessed

without any underlying knowledge of the true sensitivity functions of the camera's sensors.

Spectrophotometric measurements were made of each of the 264 patches of the Wolf Faust IT8 matte colour chart (Wolf Faust) and used to estimate, for each colour patch, the quantal catch, q , of the S, M and L cones in the blue tit as (Endler and Mielke 2005)

$$q_{j,n} = \sum_{\lambda=400}^{700} S_j(\lambda)I(\lambda)\rho_n(\lambda), \quad (6)$$

where λ denotes wavelength, $S_j(\lambda)$ is the spectral sensitivity of the j th cone (modelled using the rhodopsin visual pigment template of Govardovskii et al. (2000) and incorporating the transmittance spectra of the combined ocular media and relevant oil droplets (from data in Hart et al. 2000)), $I(\lambda)$ is the illumination spectrum (here, a spectrally flat white illuminant), $\rho_n(\lambda)$ is the reflectance spectrum of colour patch n , and the summation was performed between 400 and 700 nm.

The colour chart was then photographed under three illumination conditions, typical of those experienced by researchers working on animal colouration (in direct sunlight, under foliage and using a flash light; Table 1). There are several potential sources of error when using digital cameras to study colour: illumination conditions, the spectral sensitivity of the camera and the spectral sensitivity of the receiver's photoreceptors. How these sources of error were addressed is described below. In each output image, the R, G and B channels were linearised (using the observed and expected values of the 22 greyscale patches) to remove effects of non-linearity in the sensors' response, and equalised so as to make them independent of the illuminant (Stevens et al. 2007). By ensuring that pixel values are the same in each colour channel with respect to a greyscale standard, this procedure means that the spectral power distribution of the illuminant should have little effect on the resulting RGB values of a colour patch, as long as all wavelengths within the calibration range are represented. It also means that colour constancy (i.e. that the perceived colour of an object remains relatively constant under varying illumination conditions) will be observed in the animal observer (Vorobyev et al. 2001) (Note that, although here the greyscale standard was included in the same image as the colour patches under study, this need not be the case. Accurate results can be obtained from a 'sequential' approach, sensu Bergman and Bechner (2008), in which the greyscale standard is only present in the first image of a sequence taken under similar illumination conditions). These camera response data

Table 1 The degree of concordance between mathematically and empirically determined estimates of cone quantal catch for the L, M, S and U cones of a blue tit (see text for details), in terms of

the mean±SD difference between estimates, whether this difference is significantly different from zero, and the mean percentage error between estimates

Illumination condition	Cone class	Mean±SD difference	One-sample <i>t</i> test	Mean percentage error
Direct sunlight	L cone	-0.00047±0.015	<i>t</i> =0.51, <i>n</i> =264, <i>p</i> =0.61	1.52
	M cone	-0.00079±0.018	<i>t</i> =0.73, <i>n</i> =264, <i>p</i> =0.46	1.42
	S cone	0.00105±0.019	<i>t</i> =0.92, <i>n</i> =264, <i>p</i> =0.36	1.21
	U cone	-0.00084±0.021	<i>t</i> =0.28, <i>n</i> =47, <i>p</i> =0.78	1.69
Under foliage	L cone	-0.00079±0.020	<i>t</i> =0.64, <i>n</i> =264, <i>p</i> =0.52	1.60
	M cone	-0.00093±0.019	<i>t</i> =0.79, <i>n</i> =264, <i>p</i> =0.43	1.55
	S cone	0.00116±0.028	<i>t</i> =0.68, <i>n</i> =264, <i>p</i> =0.50	2.23
	U cone	-0.0082±0.073	<i>t</i> =0.76, <i>n</i> =47, <i>p</i> =0.45	5.82
With flash ^a	L cone	-0.00085±0.021	<i>t</i> =0.66, <i>n</i> =264, <i>p</i> =0.51	1.68
	M cone	-0.00072±0.017	<i>t</i> =0.70, <i>n</i> =264, <i>p</i> =0.48	1.35
	S cone	0.00116±0.030	<i>t</i> =0.62, <i>n</i> =264, <i>p</i> =0.54	2.44

^a The Nikon D70's built-in flash which does not emit UV radiation

were then used to provide further estimates of quantal catch (q') as follows. The data on estimated sensor spectral sensitivity (R_i , determined using the quadratic programming method described above) and the sensitivity of the blue tit's photoreceptors (S_j) were used to derive a mapping function (i.e. a set of coefficients) to convert between the two colour spaces (specifically, from R to S_j) using the spectral reflectances of 1250 matte Munsell colour chips obtained from an online database (Parkkinen et al. 1989) as training samples. The derived mapping function was then used on the camera response data to estimate q' for the S, M and L cones for all 264 colour patches. Detailed procedures for performing the linearization and equalisation, and for calculating the mapping function, are given in Stevens et al. (2007), and readers are referred there for full details. The accuracy of the estimation was tested by determining whether the mean difference between the mathematically and empirically determined estimates (q and q' , respectively) departed significantly from zero (using a one-sample *t* test), and by calculating the mean absolute difference as a percentage of the total possible difference (a measure of the mean percentage error of the estimate).

To perform an analogous comparison for sensor sensitivity within the UV, 47 coloured patches were photographed as described above for the determination of camera sensitivity in the UV, but using a different set of coloured filters and under two natural lighting conditions: in direct sunlight and under foliage (the flash emitted no UV radiation). Only the R and B sensors of the camera showed a significant response to UV radiation, so only the R and B channels of the resulting images were used. These two channels were

linearised (using the observed and expected values of six patches that showed approximately uniform reflectance between 300 and 400 nm) and equalised, and a mapping function derived using 2,000 randomly generated Gaussian reflectance curves to convert between the output of the R and B sensors of the camera and the blue tit's U cone, which has a peak sensitivity of 372 nm (Hart et al. 2000).

Results and discussion

In this paper, I aimed to describe how the quadratic programming method of Finlayson et al. (1998) can be used to estimate the spectral sensitivity of a camera's sensors. Estimated and measured spectral sensitivities of each of the test camera's three channels, R, G and B, are given in Fig. 1. It is clear from this figure that the estimated sensors are all positive, unimodal and reasonably smooth, to assuming bi- or tri-modality did not afford a better fit (i.e. in all other cases the value of the quadratic objective function was greater than for the given sensors). Moreover, there is a fairly good correspondence between the estimated sensor curves and those measured directly, especially for the blue and green sensors. The spectral sensitivity functions for the Nikon D70 camera determined here above 400 nm are also consistent with those reported elsewhere (e.g. Alsam and Lenz 2007; Huynh and Robles-Kelly 2007), and identical curves were obtained using another lens (Nikkor 55–200-mm glass lens). It is important to note that the estimates given here will only hold for the specific lens and filter combination used during the calibration process; however, given the speed at which calibrations can

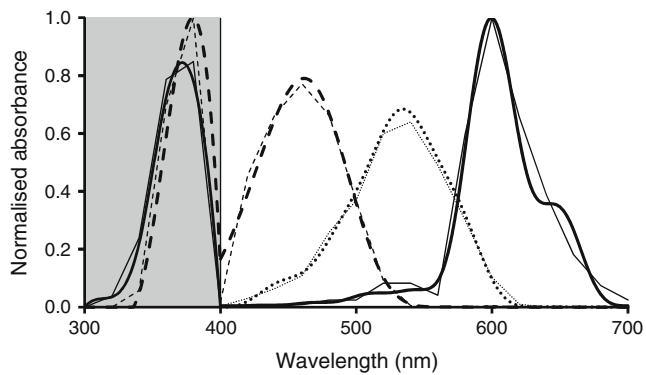


Fig. 1 Estimated (*thick lines*) and measured (*thin lines*) spectral sensitivity functions for the calibrated Nikon D70 camera's three sensors, R (*solid line*), G (*dotted line*) and B (*dashed line*). Calibrations >400 nm were performed using the UV–infrared blocking filter; calibrations <400 nm (*shaded region*) used the UV pass filter. In each calibration region the area under the curves is equalised so that each sensor produces an equal response to a flat spectrum 'white' light within the given wavelength range. Note that the G sensor had an almost zero response below 400 nm and so is not shown

be performed, repeating for a variety of lens and filters would not be too arduous.

While the estimated red sensor deviates slightly from the measured curve (see also Huynh and Robles-Kelly 2007), most notably by smoothing out the secondary peak at ~530 nm and introducing a hump on the right shoulder of the main curve (Fig. 1), the mathematically and empirically determined quantal catches of the three cone types showed an extremely close correspondence under all illumination conditions: the mean error was typically <3% and in no case did the difference between estimates differ significantly from zero (Table 1). This demonstrates that, despite these slight inconsistencies, the estimated sensor sensitivities allow sufficiently accurate mapping functions to be derived to map between device and animal colour spaces with a high degree of precision (with the only caveat being that the number of cones must be less than or equal to the number of camera sensors; Stevens et al. 2007). The limited number of training spectra used here to calculate sensor sensitivity between 400 and 700 nm (24) seems not to be a limiting factor in the accuracy of the resulting sensitivity functions, since almost identical functions were derived using the 264 patches of the

Fig. 2 a Image of a peacock butterfly (*Inachis io*) taken with the calibrated camera. How the butterfly would appear to a passerine predator (modelled using data for the blue tit) in terms of the quantal catch of **b** the L cone ($\lambda_{\max}=563$ nm), **c** the M cone ($\lambda_{\max}=502$ nm), **d** the S cone ($\lambda_{\max}=449$ nm) and **e** the U cone ($\lambda_{\max}=372$ nm), where *dark and light areas* indicate regions of low and high cone stimulation, respectively. **f** The outcome of a putative opponent-processing mechanism, termed (LM)S, comparing equally weighted inputs from the mean of the L and M cones to the S cone (Osorio et al. 1999), where $(LM)S=(q_{LM}-q_S)/(q_{LM}+q_S)$ (see Eq. 6; Endler and Mielke 2005) with $q_{LM}=(q_L+q_M)/2$. *Light regions* indicate high relative stimulation of the L and M cones, *dark regions* high relative stimulation of the S cone, and *medium grey regions* approximately equal stimulation of both

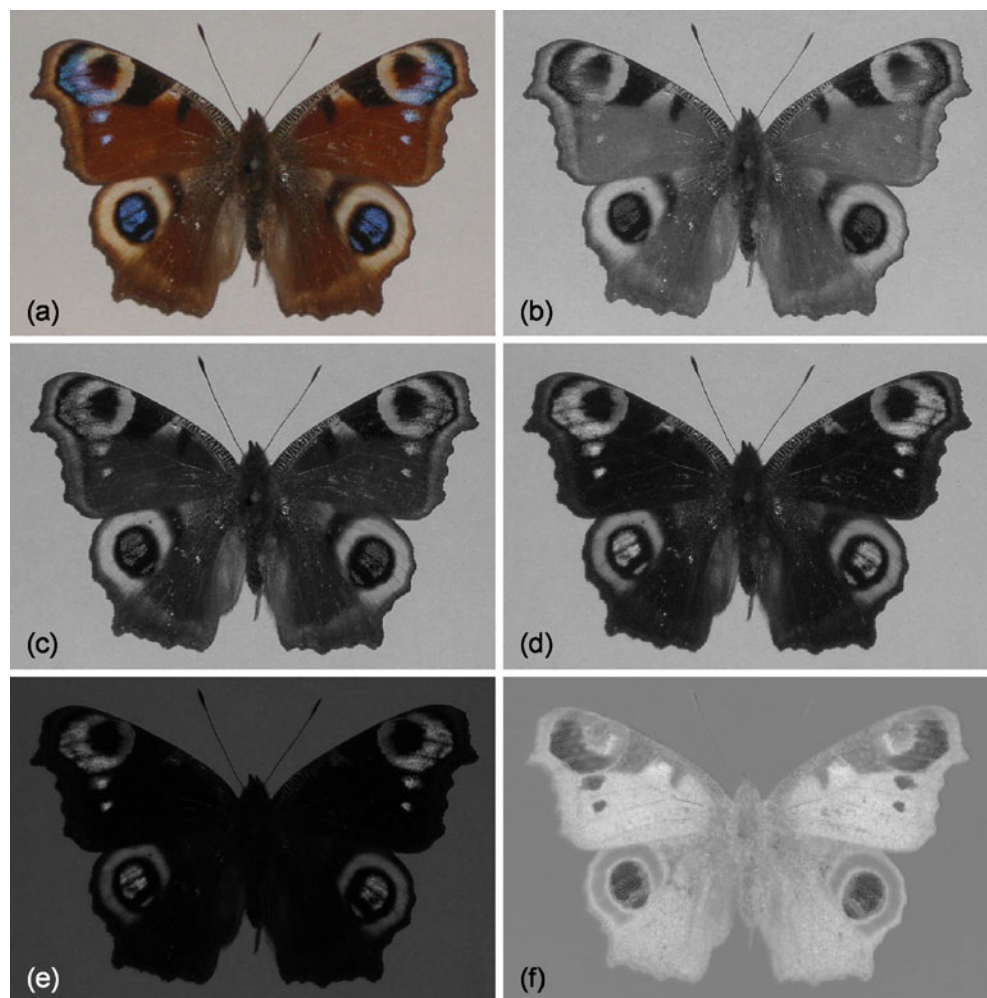


Table 2 A summary of the steps needed to utilise knowledge of camera sensor sensitivity functions to create images that represent scenes as viewed in terms of the quantum catch of an animal's photoreceptors

1. Determine the camera sensors' spectral sensitivity functions (e.g. using the method presented here or one described elsewhere; see references herein).^a
2. Using data on the spectral sensitivity of the camera's sensors (obtained in step 1) and the sensitivity of the animal's photoreceptors to be modelled (typically obtained from published data), derive mapping functions that allow conversion between camera and animal colour spaces (Stevens et al. 2007).^a
3. Obtain photographs of the colours under study using the calibrated camera and at least one image of a greyscale standard, under similar illumination conditions.
4. Linearise and equalise each image (Stevens et al. 2007).
5. Apply the mapping functions derived in step 2 to convert each image into animal colour space.

^a Steps 1 and 2 need only be performed once for each camera/lens/filter combination

Wolf Faust colour chart (data not shown); however, the number of bits in the camera image does appear to be important. Here, I used images with 12 bits per colour channel (i.e. $2^{12}=4,096$ possible values for each pixel) obtained from the RAW output of the camera and, although others have used similar techniques to derive sensitivity functions from images with 8 bits per channel (as found in standard JPEG images, for example; e.g. Finlayson et al. 1998), I found that the optimisation function often failed to converge with 8-bit input data and when it did the resulting curves differed markedly from those obtained with 12-bit data.

Imaging devices that are sufficiently sensitive to UV radiation potentially allow whole-image comparisons to be made within the full gamut of animal colour vision, which for many species extends beyond that of humans (e.g. Bennett et al. 1996). The Nikon D70 used here was found to have significant secondary sensitivity below 400 nm, at least for two of the camera's sensors (R and B). Normally, this would be undesirable as it would cause noise in an image, but it can be used with appropriate filters to accurately estimate the response of an animal's photoreceptors with peak responses below 400 nm (at least when the illuminant contains sufficient UV radiation). Cameras in Nikon's DX range (including the D70) are only fortuitously sensitive to UV, although specialist cameras do exist that are designed for UV sensitivity (such as the Fujifilm IS Pro). Even if a particular camera shows no UV sensitivity, a standard digital camera could be used in tandem with a dedicated UV-sensitive camera, such as described in Stevens et al. (2007), to extend the wavelength range over which measurements can be taken.

Although I used the algorithm developed by Finlayson et al. (1998) in the calibration presented here, a number of

different algorithms have been proposed for this task (e.g. Farrell and Wandell 1993; Sharma and Trussell 1993, 1996; Finlayson et al. 1998; Hardeberg et al. 1998; Thomson and Westland 2001; Alsam and Lenz 2007; Ebner 2007), all of which claim to achieve a similarly high degree of concordance between predicted and observed values.

If mapping functions derived from the sensor sensitivity data are applied to whole images (e.g. Fig. 2a), on a pixel-by-pixel basis, they allow estimates to be obtained on how the imaged scene would appear to an animal observer in terms of relative quantal catches for each cone type (Fig. 2b, c, d, e). Once a camera has been calibrated, performing this conversion is as straightforward as any other image manipulation; full details of the steps that need to be taken are given in the relevant references, the "Methods" of this paper and in Stevens et al. (2007). A summary is given in Table 2. Subsequent image-processing techniques allow a host of comparisons to be made, including combining images of quantal catches to explore putative opponent-processing mechanisms comparing the inputs of particular cone combinations (e.g. Stevens and Cuthill 2006; Fig. 2f), modifying images to show the

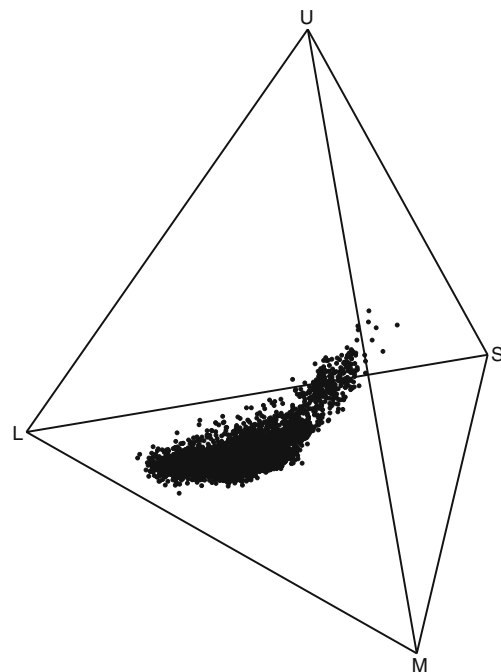


Fig. 3 The peacock butterfly from Fig. 2 depicted in a tetrahedron representing blue tit colour space (Endler and Mielke 2005), in which each data point corresponds to a single pixel of the image (for simplicity, only 10,000 randomly selected pixels have been plotted). The apices of the tetrahedron represent stimulation of a single one of the four cone classes (labelled *L*, *M*, *S* or *U*, accordingly), and the location of a data point is determined by the relative stimulation of the four cone types. The high density of data points near the *L* apex arises from the predominantly red-orange body of the butterfly while the cluster further towards the *S* apex is from the blue-UV-white eyespot regions

effects of different illumination conditions, combining images to highlight regions that differ most between two hypothetical observers (e.g. with different cone sensitivities or a different number of cones) and so on. Image analysis also allows statistics to be performed on the whole image or between images, and for the application of techniques that mimic neural processing mechanisms for tasks such as edge detection (Stevens and Cuthill 2006) and pattern analysis (e.g. Barbosa et al. 2008; Chiao et al. 2009). Quantum catch data can also be extracted from the modified image and used in exactly the same way as if the data were obtained via visual system modelling on spectrophotometric data (an increasingly common form of analysis in animal colour research; reviewed in Endler and Mielke 2005), except that the spatial resolution and quantity of data available is considerably greater. For example, the chromatic information from the butterfly in Fig. 2 can be visualised using a suitable colour space diagram, in which each data point depicts one pixel from the image (Fig. 3). By obtaining similar data for multiple individuals, relevant ecological questions can be readily addressed using existing statistical techniques; for instance, one could look for whole-body variation in colouration in response to an experimental manipulation using the multivariate test described by Endler and Mielke (2005), which is specifically designed for ‘clouds’ of data points such as those generated by this approach.

In fact, the range of possibilities that arise through understanding the sensitivity of a camera’s sensors can far exceed those performed using non-calibrated cameras (e.g. Bergman and Beehner 2008), and several recent studies have already exploited such techniques to study animal colouration. Specific applications have included the analysis of patterns in visual signals (e.g. egg maculation patterns, Stoddard and Stevens 2009); the application of image-processing techniques (e.g. edge detection to delineate pattern markings, Stevens and Cuthill 2006; Fourier transforms to assess image attributes, Stoddard and Stevens 2009) and the use of images to analyse entire spatial scenes (e.g. to assess the detectability of red fruit against green foliage, Lovell et al. 2005). Knowledge of spectral sensitivity can also be used to estimate irradiance spectra (e.g. Chiao et al. 2000) or reflectance spectra directly from digital images, a technique that has been used to explore sexual signalling in primates (Higham et al. 2010); although it should be noted that because the illumination hitting different parts of an object will not necessarily be constant, it may be difficult to separate reflectance from illuminance variation. Other, simpler methods exist for studying animal colouration; for example the method described by Bergman and Beehner (2008) provides a straightforward, accurate and repeatable method for assessing variation in colour; however, they do not take into account the spectral sensitivity of potential receivers, which

may be crucial for determining the proportion of variation that is important to the species in question (Bennett et al. 1994; Higham et al. 2010). The potential applications of the technique described above, combined with the benefits digital photography offers over spectrophotometry outlined in the Introduction, provide a range of exciting possibilities for the study of animal colouration, especially in the wild.

Acknowledgements I thank Stuart Townley and Laurel Fogarty for the discussion on the mathematical aspects of this paper and three anonymous referees for their extremely helpful comments on an earlier version. This work was supported by a NERC fellowship (NE/F016514/1).

References

- Alsam A, Lenz R (2007) Calibrating color cameras using metameric blacks. *J Opt Soc Am A* 24:11–17
- Barbosa A, Mäthger LM, Buresch KC, Kelly J, Chubb C, Chiao C-C, Hanlon RT (2008) Cuttlefish camouflage: the effects of substrate contrast and size in evoking uniform, mottle or disruptive body patterns. *Vis Res* 48:1242–1253
- Barnard K (1995) Computational color constancy: taking theory into practice. Master’s thesis, Simon Fraser University
- Beddard FE (1892) Animal coloration: an account of the principle facts and theories relating to the colours and markings of animals. Swan Sonnenschein, London
- Bennett A, Cuthill I, Norris K (1994) Sexual selection and the mismeasure of color. *Am Nat* 144:848–860
- Bennett ATD, Cuthill IC, Partridge JC, Maier EJ (1996) Ultraviolet vision and mate choice in zebra finches. *Nature* 380:433–435
- Bergman TJ, Beehner JC (2008) A simple method for measuring colour in wild animals: validation and use on chest patch colour in geladas (*Theropithecus gelada*). *Biol J Linn Soc* 94:231–240
- Chiao C-C, Osorio D, Vorobyev M, Cronin TW (2000) Characterisation of natural illuminants in forests and the use of digital video data to reconstruct illuminant spectra. *J Opt Soc Am A* 17:1713–1721
- Chiao C-C, Chubb C, Buresch KC, Siemann L, Hanlon RT (2009) The scaling effects of substrate texture on camouflage patterning in cuttlefish. *Vis Res* 49:1647–1653
- Ebner M (2007) Estimating the spectral sensitivity of a digital sensor using calibration targets. Proceedings of the 9th annual conference on genetic and evolutionary computation, London, 7–11 July 2007, ACM, New York, pp 642–649A
- Endler JA (1990) On the measurement and classification of colour in studies of animal colour patterns. *Biol J Linn Soc* 41:315–352
- Endler J, Mielke P (2005) Comparing entire colour patterns as birds see them. *Biol J Linn Soc* 86:405–431
- Farrell JE, Wandell BA (1993) Scanner linearity. *J Electron Imaging* 23:225–230
- Finlayson GD, Hordley S, Hubel PM (1998) Recovering device sensitivities with quadratic programming. IS&T/SID sixth color imaging conference: color science, systems and applications, Scottsdale, Arizona, Nov 1998, pp 90–95
- Frischknecht M (1993) The breeding coloration of male threespined sticklebacks (*Gasterosteus aculeatus*) as an indicator of energy investment in vigour. *Evol Ecol* 7:439–450
- Govardovskii VI, Fyhrquist N, Reuter T, Kuzmin DG, Donner K (2000) In search of the visual pigment template. *Vis Neurosci* 17:509–528

- Hardeberg JY, Brettel H, Schmitt F (1998) Spectral characterization of electronic cameras. *Proc SPIE* 3409:100–109
- Hart NS, Partridge JC, Cuthill IC, Bennett ATD (2000) Visual pigments, oil droplets, ocular media and cone photoreceptor distribution in two species of passerine bird: the blue tit (*Parus caeruleus* L.) and the blackbird (*Turdus merula* L.). *J Comp Physiol A* 186:375–387
- Higham JP, Brent LNJ, Dubuc C, Accamando AK, Engelhardt A, Gerald MS, Heistermann M, Stevens M (2010) Color signal information content and the eye of the beholder: a case study in the rhesus macaque. *Behav Ecol* 21:739–746
- Hill G, McGraw KJ (2006) *Bird coloration: function and evolution*. Harvard University Press, Cambridge
- Hubel PM, Sherman D, Farrell JE (1994) A comparison of methods of sensor spectral sensitivity estimation. *Proc. IS&T/SID. 2nd color imaging conference: color science, systems and applications*, pp 45–48
- Huynh CP, Robles-Kelly A (2007) Comparative colorimetric simulation and evaluation of digital cameras using spectroscopy data. *Digital image computing techniques and applications, 9th biennial conference of the Australian pattern recognition society*, pp 309–316
- Kohonen O, Parkkinen J, Jääskeläinen T (2006) Databases for spectral color science. *Col Res Appl* 31:381–390
- Lovell PG, Tolhurst DJ, Párraga CA, Baddeley R, Leonards U, Troscianko J, Troscianko T (2005) Stability of the color-opponent signals under changes of illuminant in natural scenes. *J Opt Soc Am A* 22:2060–2071
- McCamy CS, Marcus H, Davidson JG (1976) A color rendition chart. *J Appl Photogr Eng* 2:95–99
- Osorio D, Vorobyev M, Jones CD (1999) Colour vision of domestic chicks. *J Exp Biol* 202:2951–2959
- Parkkinen JPS, Hallikainen J, Jääskeläinen T (1989) Characteristic spectra of Munsell colors. *J Opt Soc Am* 6:318–322
- Párraga CA (2003) *Is the human visual system optimised for encoding the statistical information of natural scenes?* PhD Thesis, University of Bristol
- Poulton EB (1890) *The colours of animals: their meaning and use, especially considered in the case of insects*. Kegan Paul, London
- Romero J, Valero E, Hernández-Andrés J, Nieves JL (2003) Color-signal filtering in the Fourier-frequency domain. *J Opt Soc Am* 20:1714–1724
- Rush VN, McKinnon JS, Abney MA, Sargent RC (2003) Reflectance spectra from free-swimming sticklebacks (*Gasterosteus aculeatus*): social context and eye–jaw contrast. *Behaviour* 140:1003–1019
- Ruxton GD, Speed MP, Sherratt TN (2004) *Avoiding attack. The evolutionary ecology of crypsis, warning signals and mimicry*. Oxford University Press, Oxford
- Sharma G, Trussell HJ (1993) Characterization of scanner sensitivity. *Proc. IS&T/SID color imaging conference: transforms and portability of color*, pp 103–107
- Sharma G, Trussell HJ (1996) Set theoretic estimation in color scanner characterization. *J Electron Imaging* 5:479–489
- Solli M, Andersson M, Lenz R, Kruse B (2005) Color measurements with a consumer digital camera using spectral estimation techniques. *Lect Notes Comput Sci* 3540:105–114
- Stevens M, Cuthill IC (2005) The unsuitability of HTML-based colour charts for estimating animal colours—a comment on Berggren and Merilä (2004). *Front Zool* 2:14
- Stevens M, Cuthill IC (2006) Disruptive coloration, crypsis and edge detection in early visual processing. *Proc R Soc B* 273:2141–2147
- Stevens M, Párraga CA, Cuthill IC, Partridge JC, Troscianko TS (2007) Using digital photography to study animal coloration. *Biol J Linn Soc* 90:211–237
- Stoddard MC, Stevens M (2009) Pattern mimicry of host eggs by the common cuckoo, as seen through a bird's eye. *Proc R Soc B* 277:1387–1393
- Thomson M, Westland S (2001) Color-imager characterisation by parametric fitting of sensor responses. *Col Res Appl* 26:442–449
- Vora P, Farrell J, Tietz J, Brainard D (1997) Digital color cameras 2—spectral response. Technical report HPL-97-54, HP Laboratories
- Vorobyev M, Marshall J, Osorio D, Hempel de Ibarra N, Menzel R (2001) Colourful objects through animal eyes. *Col Res Appl* 6:214–217
- Westland S, Ripamonti C (2004) *Computation colour science using MATLAB*. John Wiley & Sons, Chichester