

A Unifying Colour Model for Natural Visual Systems

Thomas Bangert*, Ebroul Izquierdo ☆*

Queen Mary University of London, Mile End Road, London, E1 4JN, England

Abstract

Although the ancestors of mammals had the same visual sensor arrangement as the ancestors of birds, colour vision in most mammals is very limited. Humans belong to a small subset of mammals that have re-developed some of the colour ability lost by their immediate ancestors. They have done this by diverging one of the two colour sensors available to mammals into a third semi-independent colour sensor. The most common sensor arrangement for animals with complex visual systems is, however, four colour sensors. Animals such as birds not only have four fully independent colour sensors but each sensor has a coloured oil droplet that acts to spectrally restrict the light that falls onto the sensor. Because of this difference in the sensor arrangement, it has been suggested that it is difficult (or “impossible”) for us to know the perception of colour in animals such as birds. We propose a general colour model that aims to unify our understanding of colour vision in a wide variety of organisms, thereby allowing us to have some insight into how animals such as birds perceive colour.

Keywords: colour theory; colour model; opponent colour

1. Introduction

An organism’s ability to gather information about its environment by measuring the temporal flux of electromagnetic radiation is so vitally important that it developed in the earliest of our ancestors (see Goldsmith (2006) and Vorobyev (2006)). The sensors these early organisms developed were simple neurons with the ability to be activated by photon absorption. Once developed, these sensors have seen little change in hundreds of millions of years of evolution (see Goldsmith & Butler (2005)). The sensors differ primarily in the use of the material used to absorb photons (the *pigment*), which determines the wavelength the sensor responds best to. Absorption is a probabilistic process, and therefore the pigment determines how probable it is that a photon of a specific wavelength is absorbed, and hence counted (Lennie (2000); Solomon & Lennie (2007)). Most pigments used by these sensors absorb photons over a broad spectral range, with the probability of absorption being naturally distributed. The primary function of these sensors is therefore to measure how bright (or *luminant*) a given stimuli is at any point in time. The visual systems of many simpler organisms use sensor arrays that rely only on this type of luminance information. Sensor density is typically an important factor in the overall design of such sensor arrays. Most natural organism that rely on a visual system to navigate their environment maintain a dedicated two dimensional sensor

space onto which the photons of stimuli are projected. As the dimensions of this space are commonly very limited, the visual acuity of an organism depends on the number of sensors that may be set into this space. The number of sensors that have a presence on this space must therefore be maximized, but without significant loss in sensitivity (determined by the amount of photosensitive material used). An optimal solution to these essentially conflicting requirements is a shape that is long and thin and as a result sensor neurons of this type are commonly referred to as *rods*.

For organisms with more complex visual systems, sensor neurons dedicated to measuring luminance continue to find specialized application for conditions of extremely low luminance. However, for common diurnal levels of luminance many animals have developed a more advanced sensor. The photon absorption structure of these neurons has a triangular shape and as a result these sensor neurons are commonly referred to as *cones*. The principal feature of these neurons is that the spectral region to which they are sensitive is restricted. This spectral restriction varies, and is determined primarily by the pigment used and by any material covering the respective sensor, causing it to be partially obscured. Many organisms use a special purpose coloured oil droplet in the outward facing segment of the sensor neuron itself, which selectively absorbs photons of undesired wavelengths thereby preventing these photons from reaching the actual sensor mechanism. As a result these sensors come in a number of varieties. The sensor arrays of natural visual systems that use this more advanced mechanism commonly consist of four different

☆Multi-Media and Vision Research Group

*Corresponding author

Email address: thomas.bangert@qmul.ac.uk (Thomas Bangert)

types of such sensors, each responding most strongly to photons of different wavelengths. For this reason, these sensors are associated with the perception of colour and are often referred to as *colour sensors*. Aside from this spectral restriction, cones do nothing more than provide a measure of photons absorbed and therefore may be seen as a variation on the basic design of luminance sensor neurons (for brief review see Goldsmith (2006)).

The nature of the information that visual sensor neurons provide is therefore well understood. Luminance information is well understood, and indeed a *luminance model* is trivial to define. However, despite the fact that colour is known to be derived from simple spectrally restricted luminance measurements, there is no theory of colour that provides a general framework from which sensor values may be mapped to perception. Even for organisms that have relatively simple visual systems. In fact it has been suggested that for organisms such as birds,

... it is difficult – impossible in fact – for humans to know what their perception of colours is actually like ... They not only see in the near ultraviolet, but they also can see colours that we cannot even envision ... a variety of colours that lie beyond direct human experience. Goldsmith (2006)

We suggest in attempting to resolve this unknown that the underlying model for colour was set by our distant common ancestors and has seen little if any change in the wide variety of organisms that descended from them. The nature of colour information and how it is processed at the earliest stages of visual processing we suggest is therefore universal (in higher animals such as birds, mammals and lizards). We propose a general theory of colour information (a colour model), which will include the processing steps by which raw input produced by the sensors is transformed into the information that underlies perception.

2. Colour

Natural visual systems have developed through the process of evolution, by which complex systems develop from the simple by small gradual incremental change. To understand the development of colour in natural visual systems one must therefore have regard to how visual systems that use only luminance might develop step by step into visual systems that supports colour.

Natural visual systems do not simply measure luminance using a single sensor; sensors are always organized into pairs Hurvich & Jameson (1957). Moreover, a sensor pair with identical luminance values (produced by an equi-luminous stimuli) does not send any information to the visual system. Information is sent only if there is a luminance difference (a contrast) between neighbouring sensors, and it is this difference information that is used for higher level visual processing. Abstractly, the fundamental unit of the visual sensor array is therefore a sensor pair that

produces a difference measurement (an opponent measurement). Sensors arranged in this way may be referred to as *opponent pairs*. The measurement that an opponent pair produces is in effect a relative measure of luminance rather than an absolute measurement that a single sensor would provide.

Visual sensor systems may be seen as systems that produce a simple photon count. Photons are only absorbed in discrete units (quanta), and therefore counting photons is inherently a positive integer function. By contrast, a difference measurement derived from two or more separate photon counts is a signed integer value. Preserving the sign of these contrast values is important for any visual system (for a review see Bangert (2008)). Most neurons communicate only by producing discrete unary events known as *action potentials*. Action potential event counts are inherently positive integers and therefore negative values cannot be coded directly. For neural systems that code information so that the sign is preserved when it is sent from one area of information processing to another the mechanism that facilitates this transfer may be referred to as an information *channel* (Conway (2009)).

Once an organism (the class or species) has developed a luminance based visual system where the relative luminance measurements produced by opponent pairs are sent for visual processing via opponent channels and the organism has come to rely upon this system to navigate its environment it may find that it is often very beneficial to be able to discriminate between stimuli by a difference in wavelength (that is, individuals of the given class that have the ability are at a competitive advantage over those that do not). A food source may for example reflect only short wavelengths, but be equiluminous to its surroundings. An organism that is able to discriminate short wavelengths would have a significant advantage over rivals lacking this ability (see Conway (2009)). Much as the earliest luminance sensor arrays would have been little more than ‘bug detectors’, it is likely that early wavelength discrimination was of a fixed function that assisted with perceptual segregation (Mollon (1989)). Such a visual system would detect *colour* only if there was a difference between a dedicated *colour* sensor and the luminance sensor. Perceptually, such a visual system might be seen to present a visual representation to the organism akin to a child’s colouring book – in luminance contrast outline only, but with elements considered important being *coloured in*. As it would operate with a single colour only, such a visual system could be called *unichromatic*.

A *unichromatic colour sensor* may be implemented by modifying an opponent sensor pair where the spectral sensitivity of one of the luminance sensors is modified. This may be achieved by changing the composition of the sensor pigment that captures the photons (Solomon & Lennie (2007)), something which can be the result of a single point mutation. It may also be achieved by the development of a spectral filter, which may have its origins in something that obscures or covers the sensor. A small change therefore has

the potential to modify an existing luminance channel into a *colour* channel. Modifying the spectral sensitivity of a single sensor allows a luminance *channel* to be modified from a unit that measures spatial luminance difference to a unit that measures spectral difference. The possibility of fundamental change from opponent luminance channel to opponent *colour* channel may therefore be seen to have a potential origin in simple natural variation within a class of organism; with development then being a function of the selective pressure of the environment. Once this small change has been established then further small incremental modification would be able to extend the very rudimentary initial *colour* functionality.

A unichromatic colour system requires only unidirectional (positive) colour information. However, as it may have developed through modification of a luminance contrast channel it would inherently be capable of bi-directionality. An organism that relies on a well-developed unichromatic visual system might at some point need to diversify its food sources. It might well have been that a food source which reflected predominantly short wavelength light had been becoming increasingly rare while an alternative food source which reflected predominantly long wavelength light became increasingly more common. An organism (as a general class) might switch its *colour* pigment from short wavelengths to long wavelengths, but such a solution would make it dependant on one or the other food sources. It would be evolutionarily advantageous for the organism to be able to detect both food sources at the same time. One solution would be to simply re-use the inherent bi-directionality of the existing colour channel. Negative values might represent long wavelengths and positive values could represent short wavelengths. This solution could be obtained by a modification of the spectral sensitivity of the second sensor of the *colour* opponent pair, to move its sensitivity peak in the opposite direction of its counterpart. With this further small modification, the colour channel may now be seen as sending fully opponent colour information to the visual system. If both food sources were equal in value then this may be presented perceptually as before, with a single colour. However, it might well have been that one of the food sources subsequently developed a defence against predation, perhaps a toxin. It would therefore then have become very important for the organism to perceptually discriminate between objects that aside from spectral difference appear identical: one an essential food source and the other a deadly poison. This discrimination could have been achieved by employing two different percepts for the information being received by the colour opponent channel. Perceptually, continuing with the child's colouring book analogy, this may be seen as giving the visual system a double sided colouring crayon. Objects may be filled in with one colour or the other, but not both at the same time. Such a system could be called *dichromatic*.

Once an organism has developed dichromacy, it would have to hand the underlying opponent mechanism required

for colour. An opponent mechanism for colour was initially proposed on the basis of psychological observations by Hering and physiological evidence for this in humans was shown by Hurvich & Jameson (1957). Once an organism has developed the ability to highlight visual stimuli by the use of one opponent channel, were it to require additional colour discrimination then it could simply develop further opponent colour channels by divergence of existing colour sensor pairs. A system with two sensor pairs and two opponent channels would be called tetrachromatic and a system with three sensor pairs would be called hexachromatic. It is known that our distant ancestors, the ancestors of birds (dinosaurs), and indeed the ancestors we have in common with birds all possessed tetrachromatic visual systems (see Goldsmith (2006) for a general overview). While some of the animals for whom vision was of lesser importance have subsequently lost tetrachromacy and in some cases reverted to dichromacy (Neitz et al. (1989)), it is important to note that organisms with the most developed visual systems such as birds employ two colour channels. This suggests that the information from a single colour channel is insufficient in some way to represent the environment effectively; but that the information from two opponent channels is sufficient and no further evolutionary advantage is to be gained by the addition of further colour channels.

3. An Ideal Colour Model

The primary input to any natural visual system is luminance contrast information (areas of luminance difference, which in terms of the colouring book analogy are the outlines). The sensor array therefore presents to the visual system a two dimensional luminance contrast mosaic. A simple visual system might have developed colour initially purely to perceptually label areas of important luminance contrast on an ad-hoc basis in order to make them distinguishable from areas of unimportant luminance contrast. This ad-hoc fixed function, however, is inflexible as it leaves the early visual system to decide which elements of the information being received from the sensor array are important and which are not, leaving the organism vulnerable when the environment contradicts the implicit assumptions made by the fixed function of early visual processing. The first hurdle faced by a natural visual system driven by evolutionary pressure to make more general use of the colour information which it discerns from its sensors is therefore the development of a more general *colour model* into which the simple spectral frequency-differentiated luminance measurements provided by the sensors can be systematically mapped. As colour does not map directly from the sensor input, a colour model requires a general paradigm that establishes the essential nature of the information being derived from the sensor measurements.

The spectral power distribution of the light that the sensor array is exposed to is often complex. However, a

simplification that is inherent to unichromatic and dichromatic opponent visual systems (as outlined above) is that the light to which the sensors are exposed to is either equi-luminant across the spectral range of the sensors or it is monochromatic (light of a single wavelength). This is for two reasons: (1) under opponency, equi-luminance (or any other spectral distribution which stimulates the colour sensors to an equal degree) produces no colour information (and the light may be referred to as *achromatic*); (2) an opponent colour channel will produce the strongest opponent value when one member of the sensor pair is not stimulated at all and the other is stimulated fully. The strongest *colour* response is therefore best achieved by light that is approximately monochromatic, with any achromatic element inherently weakening the colour response. Moreover, light that is dichromatic at wavelengths that stimulate the opponent sensors equally will produce no colour information and is therefore equivalent to achromatic light. More generally, as more chromatic elements over the visible spectrum are added to the light the sensor is exposed to the more closely the light will appear to approximate achromatic light. This is due to the fact that the information provided by the sensors is incapable supporting true spectral analysis. Such a sensory arrangement therefore inherently leads to an assumption that stimuli from which colour information may be obtained must be monochromatic. Monochromatic stimuli may be represented fully by coding their luminance level and their wavelength. Therefore, if it can be practicably assumed that all colour stimuli are monochromatic, then it is likely that successive iterations of a natural visual system that initially rely on simple spectrally restricted luminance measurements to assist with visual discrimination but which are driven by evolutionary pressure to extend this ability to more generally represent stimuli will likely proceed implicitly on the assumption that monochromatic wavelength is the information that must be resolved from the raw luminance difference measurements obtained from the sensors.

If an opponent sensor pair is arranged on the spectrum as shown by figure 1 then monochromatic stimuli within the range of the sensor pair may have their wavelength determined by simple subtraction of the opponent pair, if one member of the pair is stimulated fully (to the reference level). Opponency is defined inherently as subtraction, and therefore if one of the opponent pairs is stimulated to the reference level the function of the opponent pair is equivalent to a relative measurement of wavelength. Wavelength may be determined as a function of distance from the reference position, with the sensor stimulated to the reference level acting as the reference position (of known wavelength).

4. A Practical Colour System

Natural visual stimuli will only rarely meet the requirement of ideal monochromacy (see Wyszecki & Stiles (1967)




Figure 1: Spectrally adjacent sensor pair Red (R) and Green (G), joined at wavelength λ (when $R = G$) with two spectral ranges of δ where one sensor exhibits a linear response to wavelength while the second sensor maintains a reference level. Note that the full range of the sensors cannot be used for measuring wavelength with one opponent pair, only that region where one of the two sensors is able to maintain a reference level.

for an overview of the spectrum of a wide range of natural visual stimuli). Most stimuli will either fail to stimulate the reference sensor fully (or stimulate the reference sensor too much) or the stimuli may be either not fully monochromatic or not fully equi-luminant. A visual system that functions on the basis of the monochromacy assumption and which uses spectrally restricted opponent sensor pairs to determine wavelength, would when presented with sub-optimal sensor values attempt (if possible) to place the sensor values into the standard reference form, from which wavelength may be determined.

The problem of stimuli that are monochromatic but which fail to stimulate the reference sensor optimally may be solved by simply normalizing the sensor input by scaling to the reference level. On the other hand, for stimuli that are not monochromatic simple luminance sensors are unable to determine multiple wavelength components (the information is not sufficient for a true spectral analysis). However, if it is assumed that non-monochromatic stimuli are simply monochromatic stimuli with an equiluminance component then the equiluminance component (if known) can simply be subtracted – leaving the monochromatic component to be scaled to the reference level to determine wavelength. If luminance is calculated indirectly at a further stage then any normalization and equiluminance information must be preserved – and therefore a *channel* for both a normalizing component and an equiluminance component must be allocated.

With equiluminance and normalization, any colour stimuli (no matter how complex) may be reduced to an equiluminance component and a monochromatic component. The monochromatic element varies by its position on the spectrum (the wavelength) and its amplitude (its *monoluminance*) and the equiluminance varies only by amplitude (see figure 3). As monoluminance and equiluminance are inversely proportional in respect of total luminance, only the equiluminance information needs to be kept, once the overall luminance is known.




Figure 2: Idealized sensor response across the visible spectrum of a dual channel tetrachromatic visual system. The sensors may be labelled blue (B), green (G), yellow (Y) and red (R). The opponent pairs are R-G and Y-B. To increase the spectral range, circularity is employed. The active range in this case over which wavelength of monochromatic stimuli may be determined is 430-670nm.

A simplified colour paradigm may not accurately reflect the underlying physics of the electromagnetic spectrum, as its function is to provide information that is useful to an organism only within the environment that provides the evolutionary impetus. Natural visual stimuli that fall outside the visual system's colour paradigm may therefore resolve to code words of visual stimuli that are consistent with the paradigm. This many to one mapping of colour may be referred to as metamerism (see Wyszecki & Stiles (1967) and Hunt (1995) for an overview of metamerism in the human visual system).

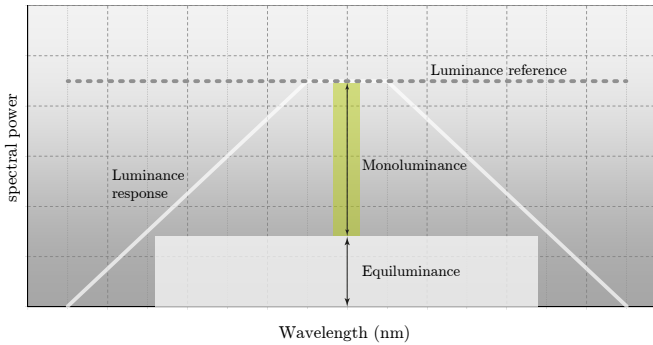


Figure 3: A schematic representation of practical opponent colour, with mono-luminance being the single (monochromatic) wavelength element which is superimposed upon a base equiluminance element. The monochromatic wavelength is to be determined and coded, which once in standard form represents *colour*. Equi-luminance is the degree to which a colour is im-pure – mixed with all wavelengths (which would otherwise be coded as achromatic – *white*). Overall luminance simply expressed is monoluminance + equiluminance. Luminance may decline at the sensor peripheries due to sensor limitations.

When wavelength is measured as difference from a known reference position, one sensor is required for the reference position and a second sensor measures the distance from the reference. While these values may be scaled, which preserves the ratio between reference and distance

from reference, they do not differentiate between luminance and equiluminance. This must therefore be determined from an independent sensor. An opponent system that relies on a single opponent pair to measure wavelength must therefore add a third sensor to be able to measure equiluminance. In natural visual systems sensors are always arranged into opponent pairs and therefore the requirement for a third sensor would lead to the development of a second opponent pair. If equiluminance can be measured not from a dedicated sensor but simply taken as the lowest sensor value of the four sensors from two opponent pairs then this would allow both opponent pairs to be used for determining monochromatic wavelength.

As shown by figure 2, if two opponent sensor pairs are available, they may be arranged so as to be interleaved on the spectrum allowing a sensor of one pair to be at the reference level and a sensor of the second pair to determine wavelength. This design allows a much greater proportion of the spectral range to be used to measure wavelength. This form of opponent design also increases the number of reference positions – enabling a reference position at the peak response of each sensor. What is measured is the distance from one of these known reference positions. When more than one reference position is used, it is necessary to code which reference position the difference measurement is linked to.

To allow the distance from the reference position to be determined, any equiluminance (the minimum value) must first be subtracted from the sensor values. Then these values must be normalized to the value closest to the reference level (the maximum value). This amounts to a degree of pre-processing that must be carried out on the raw sensor values before the wavelength of the monochromatic component can be determined. A system that uses two opponent channels and an element of pre-processing to place the visual information into a standard form that allows determination of wavelength may be referred to as *dual channel reference opponency*.

5. Coding Dual Channel Opponency

Under the opponent colour model wavelength measurement is not absolute but always in relation to a known reference position. For the opponent system shown by figure 2, once the opponent values are put into the standard form, wavelength may be determined by simple subtraction of the opponent values. This will produce two values, one being the reference and the second indicating the distance from the reference. If the reference position is known then wavelength may be determined by distance from the reference. Dual channel opponency has four reference positions, each with a dual transition space. When the measured distance from a reference position is precisely null this position may be referred to as a *primary* position. A dual channel opponent system would have four primaries. The information that must be coded is the active primary and the distance from that primary. This may be coded directly as two opponent values, with one opponent value held at a constant (indicating which primary is active) and the second being the measured distance from the active primary. The opponent coding is given by equation 1 (lower case indicates sensor values after pre-processing):

$$\lambda = \langle r - g, y - b \rangle \quad (1)$$

If natural visual systems were to use dual channel opponency with colour coded as suggested and if this opponent code were to be preserved throughout the system (that is, it is used as a ‘native’ code) then it would be predicted that the opponent positions would be reflected directly in colour perception. If the reference positions are preserved throughout and are reflected in perception then they may be referred to as the *perceptual primaries*. A dual channel reference opponent system would code for four primary colours, with all other *pure* perceptual colours being the product of the transition between adjacent primaries. It would be predicted that transition colours would be perceived as a *mixture* of the primaries and as an extension of this, that transition colours could be reproduced by stimuli consisting a *mix* of monochromatic light at the primaries. Reference opponency would predict that only adjacent primaries could be mixed to produce colour stimuli; a mix of non-adjacent primaries would be lost to opponency (leaving only luminance).

6. The Problem of Yellow

Natural visual stimuli are rarely monochromatic (see figure 4). Many displays of colour as found in the natural environment are inconsistent with the hypothesis that *colour* is monochromatic. Specifically, organisms that purposefully display colour do so in a way that is inconsistent with the monochromatic hypothesis. *Yellow* is a display colour commonly found in organisms such as plants which use colour to advertise for a specific service (such as a flowering plant advertising a reward of nectar for the service of

pollen delivery). This is directed at organisms which have visual systems able to discriminate colour and which have the ability to deliver the service. A *natural yellow* is commonly produced by absorbing the short (*blue*) wavelengths and reflecting a very broad area of the spectrum from green to red (approximately 530nm-700nm). For a dual channel reference opponent system, a *pure* natural yellow would therefore have no equiluminance component (as $b = 0$), leaving all the remaining sensor values (r , y , and g) at the reference positions. While this contradicts the assumption of monochromacy (where only one sensor may be at the reference level); in practice, as opponent pairs subtract, the $r - g$ opponent pair will subtract to null, whereas $y - b$ will leave y at the reference level (as $b = 0$). This will result in an opponent code of $\langle 0, 1 \rangle$, which is *pure* yellow (indicating a monochromatic wavelength exactly on the yellow primary). Under opponency, a *natural yellow* stimuli will therefore be coded identically to a monochromatic yellow stimuli. The stimuli will, however, differ significantly in their luminance; a spectrally broad natural yellow will result in a higher luminance value than the equivalent monochromatic yellow. This then could explain why organisms commonly advertise with a natural yellow rather than a monochromatic yellow.

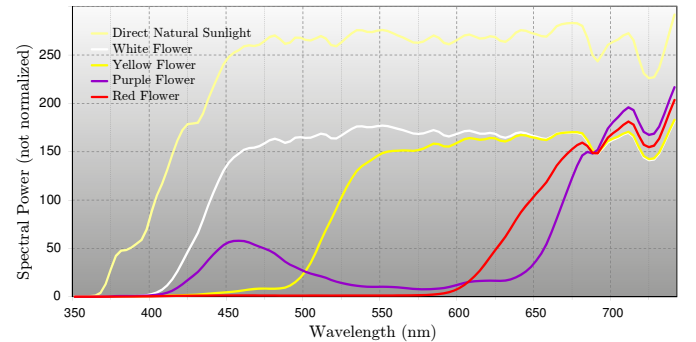


Figure 4: Spectral reflectance of flowers produced by a common decorative flowering plant; colours are produced by differing varieties of the same plant.

7. Conclusion

The proposed opponent colour model allows the term *colour* to be not only defined (see Fairchild (2004) for a review of the difficulties a definition entails) but defined objectively as part of a formal system of information processing. The fundamental element of this definition is that colour is inherently a measure of monochromatic wavelength. Colour is the means by which a luminance based visual system extends itself to represent visual stimuli with greater fidelity. This requires colour to be *backwards compatible* and coexist with luminance; which leads necessarily to a third item of information that must be measured and stored – the level of equiluminance in ratio to monochromatic luminance. This third value may be seen

as the bridge between pure achromatic luminance and *pure* monochromatic colour (sometimes referred to as fully saturated colour).

By convention colour is defined in terms of perception. If the colour model being proposed accurately reflects the function of natural visual systems then perceptual colour may be defined simply as the perceptual by-product of an underlying visual information processing and coding system. This system reduces any visual stimuli within the range of its sensors to the standard three values of luminance, equiluminance and wavelength. It is proposed that it is the role of the early visual system to process the input from the visual sensors; transforming raw sensor input into a standard form suited for further processing. The information may then be transmitted in that form to a variety of distant higher level visual processing areas. If any further processing preserves the underlying information format then this will ultimately be reflected in perception.

If natural visual systems operate on the basis proposed then the underlying paradigm that must be assumed is that visual stimuli are partly pure monochromatic and partly pure equiluminant. Given this assumption, the dual channel reference opponent system shows how wavelength may be effectively calculated without performing a true analysis of the spectrum; wavelength may be determined by the use of pairs of modified luminance sensors and the application of simple arithmetic. Well-formed stimuli (under the proposed paradigm) may be coded without any loss of information, and all non well-formed stimuli have well-formed equivalents (metamers).

What is proposed is a theoretical information processing model that shows how *colour* information may be effectively organized given certain simplifying assumptions. It will be left to future work to show empirical evidence that natural visual systems use this underlying colour model, but it may be noted that the visual sensor arrangement of birds, reptiles and some arthropods appear very similar to that predicted by dual channel reference opponency (see figure 5 in relation to figure 2; in respect of the use of filters to shape sensor response, see Cronin & Caldwell (2002) for a brief review). It may also be noted that the dual opponent colour model provides a theoretical basis that allows us to explain why colour perception in humans is divided into four perceptually unique primaries, two of which appear to be mutually exclusive. The primaries simply reflect the reference positions used by the opponent system, where each primary must be identified by a unique code. Colour therefore is categorized by these codes which identify the specific wavelength region, and if the format of this code remains unchanged throughout then ultimately that format will be reflected in perception.

If natural visual systems do use the underlying colour model proposed then colour perception becomes not only something that we can envision but which we can readily predict. We can know what the perception of colours for animals such as birds is like because we ourselves have the same underlying system. Indeed, humans form a spe-

cial case of a small class of animals which have 4-sensor dual opponent channel ancestors which due to their nocturnal lifestyle lost the sensors driving one of their opponent channels. When the descendants of these animals re-emerged into the diurnal world some of these animals began the process of re-developing these sensors. This process remains in the early stages, with only limited divergence of one of the sensors for the second opponent pair. We suggest that this explains why some animals (including humans) have three sensors rather than the four sensors required by dual opponency. It will be left for further work to show that a fourth sensor may be effectively emulated from three physical sensors to drive the second opponent channel. If the visual system of humans and similar animals perform such emulation then it would be predicted that this function must be carried out by early visual processing. Moreover, it would be predicted that a three sensor solution is an interim solution (an evolutionary crutch or bridge) until a full complement of colour sensors has developed.




Figure 5: Spectral sensitivity of visual sensors found in birds – four sensors each with narrowed spectral frequency distribution by coloured oil droplet which acts as spectral filter. The dotted lines show the distribution prior to filtering. Adapted from Goldsmith (2006).

8. Acknowledgements

Thanks to support from EU Seventh Framework Programme and the EcoPix Project .

References

- Bangert, T. (2008). Trianglevision: A toy visual system. In V. Kurkov, R. Neruda, & J. Koutnk (Eds.), *Artificial Neural Networks - ICANN 2008* (pp. 937–950). Berlin/Heidelberg: Springer volume 5163 of *Lecture Notes in Computer Science (LNCS)*. URL: <http://www.springerlink.com/content/y5651272h2u1uq21>. doi:10.1007/978-3-540-87536-9_96.
- Conway, B. R. (2009). Color vision, cones, and color-coding in the cortex. *The Neuroscientist*, 15, 274–290. URL: <http://nro.sagepub.com/content/15/3/274>. doi:10.1177/1073858408331369.
- Cronin, T. W., & Caldwell, R. L. (2002). Tuning of photoreceptor function in three mantis shrimp species that inhabit a range of depths. ii. filter pigments. *Journal of Comparative Physiology*

- A, 188, 187–197. URL: <http://link.springer.com/article/10.1007/s00359-002-0292-z>. doi:10.1007/s00359-002-0292-z.
- Fairchild, M. D. (2004). Color Appearance Models: CIECAM02 and Beyond. In *The 12th Color Imaging Conference*. IS&T. URL: <http://www.cis.rit.edu/fairchild/PDFs/AppearanceLec.pdf>.
- Goldsmith, T., & Butler, B. (2005). Color vision of the budgerigar (*Melopsittacus undulatus*): hue matches, tetrachromacy, and intensity discrimination. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, 191, 933–951. URL: <http://dx.doi.org/10.1007/s00359-005-0024-2>. doi:10.1007/s00359-005-0024-2.
- Goldsmith, T. H. (2006). What Birds See. *Scientific American*, 295, 68–75. doi:10.1038/scientificamerican0706-68.
- Hunt, R. W. (1995). *The Reproduction of Colour*. Kingston-upon-Thames, England: Fountain Press.
- Hurvich, L. M., & Jameson, D. (1957). An opponent-process theory of color vision. *Psychological Review*, 64, 384–404. URL: <http://psycnet.apa.org/psycinfo/1959-02846-001>. doi:10.1037/h0041403.
- Lennie, P. (2000). Color vision: Putting it together. *Current Biology*, 10, R589–R591. URL: <http://www.sciencedirect.com/science/article/pii/S0960982200006321>. doi:10.1016/S0960-9822(00)00632-1.
- Mollon, J. (1989). “Tho’ she kneel’d in that place where they grew...” The uses and origins of primate colour vision. *J Exp Biol*, 146, 21–38. URL: <http://jeb.biologists.org/content/146/1/21.short>.
- Neitz, J., Geist, T., & Jacobs, G. H. (1989). Color vision in the dog. *Visual Neuroscience*, 3, 119–125. doi:10.1017/S0952523800004430.
- Solomon, S. G., & Lennie, P. (2007). The machinery of colour vision. *Nat Rev Neurosci*, 8, 276–286. URL: <http://www.nature.com/nrn/journal/v8/n4/full/nrn2094.html>. doi:10.1038/nrn2094.
- Vorobyev, M. (2006). Evolution of colour vision: The story of lost visual pigments. *Perception*, 35, 168. URL: <http://www.perceptionweb.com/abstract.cgi?id=v060037>.
- Wyszecki, G., & Stiles, W. S. (1967). *Color Science: Concepts and Methods, Quantitative Data and Formulas*. New York: John Wiley & Sons.