Colour:
an algorithmic approach

Thomas Bangert

School of Electronic Engineering and Computer Science

Queen Mary, University of London

A thesis submitted for the degree of

Master of Science

August 2012

I would like to dedicate this thesis to my loving and patient partner ...

Acknowledgements

Abstract

The fundamentals of colour vision were set out in the mid-19th century but have been split between the empirical observation that the underlying hardware responsible for vision was based upon three classes of physical sensors and the perceptual finding that colour consisted of variations of four underlying indivisible primaries, organized into two opponent pairs (blue-yellow and red-green). One of the great advances in the understanding of colour vision was developing an understanding of the mechanism of opponency that makes up the first layer of the neural circuitry that resides directly behind the sensor array of the human visual system. Two opponent colour channels were found, precisely as predicted by the study of perception. Despite the fact that the neural processing circuitry of the visual sensor array consists of only two or three layers of neurons, little further progress has been made to decipher the functionality of subsequent layers. As a result there is little agreement on the nature of the information that is produced by the neural systems that lie directly behind the sensors (at the front of the brain) which is sent to the visual system at the rear of the brain. In this thesis it is proposed that the failure to understand the nature of this information stems mainly from two factors: (1) a need to compensate for an inherent deficiency in the sensor array specific to our evolutionary history (2) the success of the paradigm under which colour is a property of perception rather than information structured by underlying function. In this thesis a paradigm of colour as functional information of an artificial computational visual system is proposed, a simplified artificial colour sensor processing system is presented and parallels are drawn between how this system processes information and how the human visual system is known to process information. It is suggested that understanding the computational requirements of functional colour processing might be helpful in understanding the complex functionality that resides directly behind the sensor array of the human visual system.

Contents

List of Figures

Chapter 1

Introduction

The understanding of colour vision has long been divided between theorists that seek to associate the subjective study of human colour perception with neural function and practitioners that seek to find a systematic means to accurately reproduce colours that match perception. This division centres on two separate empirical findings: (1) that human colour perception is organized into precisely four primary colours, divided into two opponent pairs and (2) that the human visual system relies on precisely three types of light sensors for colour vision Lennie [2000]. These findings still remain central to the unresolved differences faced by those seeking to develop a unified theory of colour vision.

Great progress has also been made in the ability to systematically produce light stimuli that match colour perception Demarsh and Giorgianni [1989]. However, there continue to exist colours which are outside the range of these colour matching systems Hunt [1995], which suggests that these systems do not fully reflect the functionality of the human visual system.

It is almost universally accepted that the human visual system (HVS) with respect to colour vision employs a sensor array with only three types of sensor [1], each tuned to respond best to a different part of the visible spectrum of electromagnetic radiation (light) Lennie [2000]. Those seeking to reproduce colour take it as axiomatic that any perceptual colour may be reproduced with three monochromatic light stimuli (often called the primaries), with the assumption that the peak frequencies of the primaries are the frequencies the respective sensors are most sensitive to. Colour reproduction therefore seeks to understand the relationship between a given set of primaries and the perception of colour. The basis of this approach are subjective studies which seek to determine this relationship by mapping the perception of specific colours with

---

[1]Studies by Jordan and Mollon [1992] indicate that a small subset of female humans might as a result of being carriers of a gene that encodes a photo-pigment with a spectral sensitivity somewhere between red and green (which in male humans leads to anomalous trichromacy) have a fourth colour sensor Greenwood [2012].

Figure 1.1: Normalized RGB colour matching functions for monochromatic beams of light of a specified wavelength (λ) in CIE 1931 colour space CIE [1932].

relative brightness levels of a given set of primaries. The relationships established by these studies are often called colour matching functions. When these primaries are presented on their own they are generally perceived approximately as the three primary colours traditionally associated with each of the three sensors: red, green and blue (RGB). The difficulty this approach faces is that as one systematically attempts to match the perceptual response to stimuli of a single wavelength (monochromatic light) throughout the visual spectrum (approximately 400 nm to 700 nm, see Figure 1.1) one finds that there is no set of three monochromatic primaries with which one is able to match all the monochromatic light stimuli within the visible spectrum Demarsh and Giorgianni [1989]; Thornton [1999]. In fact approximately one third of the visible spectrum cannot be matched, specifically the central areas of the visible spectrum that are most important to the perception of how bright (luminant) light stimuli are (approximately 450 nm to 545 nm). A match can only be achieved by adding *negative light*, which in practice means adding one of the primaries to the monochromatic light to be matched. To avoid negative values a convention has been adopted that employs a set of *imaginary primaries* Demarsh and Giorgianni [1989] that contain negative energy at certain wavelengths. This system is effective as a colour matching tool but it is clear that as a theoretical model it fundamentally fails to reflect the function of the human visual system.

Colour matching operates on the premise that colour may be fully described by three physical sensors being matched to three primaries. The study of perception, on the other hand, indicates that there are four primary colours rather than three. This has led to much dispute between the proponents of those who seek to reproduce perceptual colour and those seeking to understand the underlying neural mechanisms of colour vision. One of the triumphs of neurophysiology in the last half century has been to show that while the human visual system relies on three sensors for colour vision, early processing immediately behind the sensor array translates the input of the three sensors into a two channel opponent code, coding for yellow-blue and red-green opponent values exactly as predicted by colour theory set out more than a century before. The HVS is therefore said to employ a two-stage process in processing colour information DeValois and DeValois [1993]; Hurvich and Jameson [1957]; Mancuso et al. [2010]. RGB sensor values are determined in the first processing stage and these values are then translated into two opponent values in the second stage.

An opponent value is mutually exclusive; a value may be one or the other, but not both. Colour opponency therefore allows bluish-green or reddish-yellow as valid colours but not reddish-green or bluish-yellow. Opponency is essentially a way of coding information more compactly, by taking a set of two natural numbers and coding them as a single integer. With colour opponency any colour may be defined by just two integer values. The two stage colour opponency model has been widely accepted by those studying the HVS from a neurophysiological point of view, and is often said to have reconciled colour opponency theory with the colour models that have arisen from the practice of colour reproduction Lennie [2000].

The apparent triumph of colour opponency has, however, not dislodged the view that the HVS is inherently a *tri-chromatic* system. It has also not led to improved means of colour reproduction and, with some minor exceptions, colour reproduction has remained steadfastly a three sensor/three primaries approach. One reason for this is that there are well known perceptual inconsistencies between the predictions of colour opponency and perception Neitz and Neitz [2008]. Figure 1.2 shows the relationship between the findings of colour matching and the predictions of colour opponency to be discrepant. These difficulties as well as a failure to make further progress in the understanding of the neural circuitry responsible for opponency has stymied progress in the development of a more general theory of colour Conway [2003]; Neitz and Neitz [2008]; Solomon and Lennie [2007].

Where both proponents of colour opponency and colour reproduction agree is that colour is a property of perception. The aim of the latter is to understand how one might combine the three primaries at different ratios to reproduce colour perception while the aim of the former is to understand the neural circuitry responsible for producing colour perception. Both approaches have led to unresolved difficulties Conway [2003]; Demarsh and Giorgianni [1989]. It is proposed here that the principal obstacle in the development of a general theory of colour is the success of the paradigm under which a measure of subjective perception is used as the basis for a general theory. Colour perception must in part reflect the processing performed by the underlying mechanisms of the visual system but it does not necessarily define that underlying system, and it is therefore unsound to rely purely on perception to build a theory of colour. The central aim of any theory of colour should be to build a general understanding of the underlying system. The fundamental function performed by any visual system is the processing of the information which has its origin with the measurements of light stimuli taken from the external environment. Any general theory of colour must therefore be set out primarily in terms of an information processing system. The first step performed by any information processing system is to code information, and as with any code the first step in deciphering it is to understand the nature of the information being coded.

1.1 Structure of Thesis

This thesis is presented as a theoretical investigation of the principles of colour vision. The investigation is presented from an information processing (computational) point of view and is restricted to the early processing elements of colour vision: the sensors themselves, the nature of the information produced, and how that information is transformed by the initial stages of visual information processing. Colour vision in natural organisms is assumed to be an evolutionary adaptation of pre-existing monochromatic vision. From what is known to be true about the early visual system in a variety of organisms, an attempt will be made to follow the engineering design choices available

to the process of evolution. Evolution is a process of incremental modification of an existing design under the constraint that each modification must be useful in itself. As an initial step, it will therefore be shown how the sensors available to visual systems (which gathers only luminance information) might be modified to adapt to environmental conditions where luminance information fails to provide sufficient information about a light stimulus. From this modification, a working model of colour information will be developed. First, the nature of colour information that the modification implies will be set out. Second, the engineering design issues that will allow the modification to become a more general solution will be addressed. It will then be shown how this might be applied to the constraints of the human visual system. Lastly, it will be discussed how the proposed design might find practical application with image reproduction and how it might be applied to the attempt to understand the visual systems of animals such as birds. As this work is an attempt to design the first stages of a visual system from first principles, it does not primarily rely on existing work on colour theory. A review of the literature on colour theory is therefore presented in the appendices (Appendix B provides a review of the CIE colour model).

1.2 Motivation

It is often the case that practice precedes theory. The quest to construct machines that performed useful work preceded a theory of thermodynamics (which was developed on the basis of these machines) and the quest to build machines that could fly preceded a correct understanding of the principles of aerodynamics. It was known for example that a broad class of natural flying machines (birds) existed and this was something that had been studied for almost two and a half millennia, with little progress. Only once it was found that practical flying machines could be built and could be useful was it that the foundations of a correct understanding of aerodynamics began to be set out. It follows from this and many related examples that the best way to understand the visual system found in humans and other higher animals might be to devise artificial visual systems; that is to build systems that are able to take visual stimuli as input and from this perform some useful function. Highly evolved systems are often built on simple principles but suffer from ornate complexity that detracts from the underlying principles. With birds for example, the complexity of the feather and the irregular oscillations of wing movements hid the underlying principles of aerodynamics. The complexity in that case was simply a by-product of the evolutionary process which started with a machine inherently unsuited to flight and adapted it step by step into a virtuoso of the air. Modifying a pig for flight will in time lead to a hawk, but it will not provide a simple and elegant demonstration of the underlying principles of aerodynamics. Here it is proposed that natural visual systems are in the same way also built on inherently simple underlying principles, but that they suffer from unnecessary complexity as a result of the evolutionary process. In the subsequent sections a practical working sensor array is presented that is based on the known principles of natural

visual systems but which has been designed with the needs of a practical artificial visual system in mind Bangert [2008]. It is expected that the underlying functionality of this system will reflect that of natural visual systems.
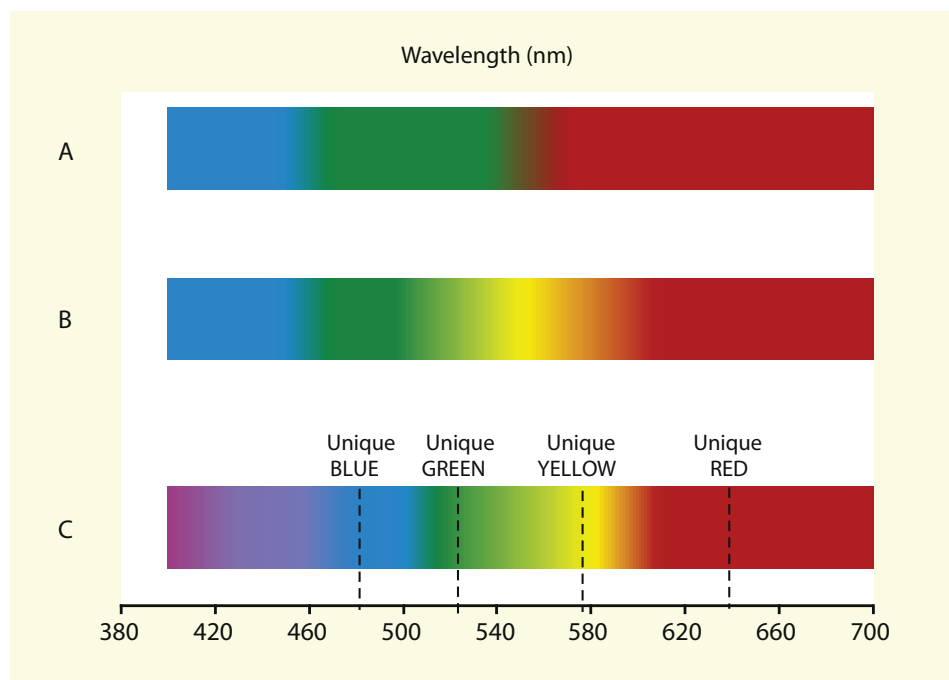
Figure 1.2: Given a monochromatic light stimulus of a given wavelength: (A) shows predicted perception from a 3 colour sensor system (red, green, blue) (B) shows predicted perception from dual opponency (yellow-blue, red-green) (C) shows findings of human perceptual studies. Adapted from Neitz and Neitz [2008].

Chapter 2

A Brief History of Colour

An organism's ability to gather information about its environment by measuring the flux of electromagnetic radiation is so vitally important that it developed in the earliest of our ancestors Goldsmith [2006]; Vorobyev [2006]. The sensors these early organisms developed were simple photon counters; they employed a photo-sensitive material (or *pigment*) that selectively absorbs photons within a certain range of wavelengths and linked it to a mechanism that measures the effects of photon absorption. Once developed, these sensors have seen little change in hundreds of millions of years of evolution Goldsmith and Butler [2005]. The fundamentals are identical in all higher organisms that possess visual systems. The sensors differ primarily by the pigment used, 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 and Lennie [2007]. 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 system of many simple organisms is a sensor array which provides only luminance information. Indeed, the HVS has such a luminance-only sensor array that is reserved for conditions of extremely low luminance. The sensor neurons of this array (called rods) sit alongside the sensor neurons that operate at normal luminance levels (called cones).

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

It is known that luminance based visual systems do not simply measure luminance, but a relative measure of luminance (luminance contrast). A sensor array that is exposed to a uniformly luminous visual field does not send any information to the visual system. Information is gathered only if there is a luminance difference (a contrast) between

neighbouring sensors, and it is this difference information that is used for further visual processing. Abstractly, the fundamental unit of the visual sensor array is therefore a sensor pair that measures contrast. What it provides is in effect a relative measure of luminance rather than an absolute measurement.

Visual sensors produce a simple photon count, which is always a positive integer value. Once we take the difference between sensor values the resulting value can be positive or negative. Preserving the sign of these contrast values is important for any visual system (for a review see Bangert [2008]). The sensor array of a luminance based visual system may therefore be seen as producing a sparse array of signed numbers. Neurons generally represent information by means of the rate by which they *fire* action potentials, which is inherently a positive value. Therefore, to represent signed (or bi-directional) values the number must be coded and transmitted as an *opponent channel* Conway [2009]. This is simply a means by which signed integer values are coded for transmission from the sensor array at the front of the brain to the visual system at the rear of the brain.

Once an organism has developed a luminance based visual system and relies upon it to navigate its environment it is often very beneficial to be able to discriminate between the different wavelengths of stimuli the organism is presented with. A food source may for example reflect only the short frequencies, but be equiluminous to its surroundings. An organism that is able to discriminate the short wavelengths would have a significant advantage over rivals without this ability Conway [2009]. Much as the earliest luminance sensor arrays would have been little more than 'bug detectors', it is likely that early frequency 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 *colour sensor* may therefore be seen as a modified luminance contrast sensor pair where the spectral sensitivity of one of the luminance sensors has been changed. This may be achieved by changing the composition of the sensor's pigment that captures the photons Solomon and Lennie [2007], which can be accomplished by a single point mutation. This small change would therefore allow an existing luminance channel (to which the sensor pair is attached) to be modified into a *colour* channel. Modifying the pigment of one of the sensors allows a sensor *channel* to be modified from a unit that measures spatial luminance difference to a unit that measures spectral difference.

A unichromatic colour system requires only unidirectional (positive) colour information. However, as it has developed by modifying a luminance contrast channel it is capable of bi-directionality. An organism with 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 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 improved by modifying 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 if consumed. 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 has a system with fully functional colour opponency. If an organism were 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 Goldsmith [2006]. While some of the higher 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 none of the higher animals are known to have developed more than two colour channels.

It follows from this very early development of tetrachromacy that the information from a single colour channel is insufficient in some way to represent the environment effectively. Conversely, the relative stability of dual channel colour indicates 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.

Chapter 3

The Nature of Colour Information

Any organism that possesses a chromatic visual system uses the colour information to assist it to effectively navigate the environment in which it lives. This information is not necessarily perceptual and no non-human organism gathers colour information for the purpose of re-presenting it to a secondary visual system (that is, there are no cameras in nature). While many cephalopods for example make dramatic use of coloured pigments to the end of disguising themselves they do not seek to accurately reflect functionality of their visual system or re-present (play back) visual stimuli to other visual systems accurately. The primary function of all sensor arrays of natural visual systems is to gather information that is helpful to the organism's ability to navigate its environment effectively.

The primary input to any natural visual system is luminance contrast information. The sensor array therefore presents to the visual system a two dimensional luminance contrast mosaic. A simple visual system might wish only to use colour to label important luminance contrast in order to distinguish it from unimportant contrast. This fixed function, however, is very 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 fixed function of early visual processing. Nevertheless, more complex visual systems which have a more generalized colour model have all been derived from simple fixed function visual systems. All *colour* functionality is based on simple frequency-differentiated luminance contrast. The first hurdle faced by a natural visual system that seeks to make more general use of colour information is therefore the development of a general colour model into which simple spectral frequency-differentiated luminance information provided by the sensors can be mapped.

The spectral distribution of the light that the sensor array is exposed to is often com-

plex. However, a simplification that is inherent to monochromatic and dichromatic opponent visual systems is that the light which a sensor is exposed to is either equi-luminant across the spectral range of the sensor array or it is monochromatic. Oppo-nency means that equiluminance (or any other spectral distribution which stimulates the colour sensors to an equal degree) produces no colour information. This leads to an assumption that if there is colour information that it must be monochromatic. And if a light stimuli that is being presented to a sensor is monochromatic, then to represent that light stimuli more generally we need to know only how luminant that stimuli is and what the dominant spectral frequency is. Therefore if the principle of monochromaticity of colour is assumed, then it is likely that a natural visual system that is seeking to develop a more general colour model would choose the principal element of colour information as spectral frequency. Early dichromatic visual systems may be seen as very crude variants of this, dividing the spectrum up into two regions (short and long).

If a visual system receives luminance and spectral frequency information from its sen-sor array then it must employ a system to effectively represent spectral frequency. Any sensor array, particularly neural sensors, produce noisy, contradictory and unreliable information. Spectral frequency information therefore cannot simply be measured but must be computed. It is known from the study of the human visual system that per-ceived colour is often adjusted from the colour presented to the sensor or even inferred from achromatic stimuli Lotto and Purves [2000]. With reflected light stimuli for exam-ple, the quality of the light source is of considerable importance and can vary dramat-ically, producing a wide variety of sensor values all of which must be resolved to the same colour (colour constancy). This requires complex processing. Any visual system that uses colour to represent spectral frequency therefore must have a systematic rep-resentation of spectral frequency that lends itself to efficient colour calculations. What is required is not only to measure spectral frequency accurately, but to also to develop a model of the spectrum that lends itself to efficient colour computation.

Given a model of the spectrum, we may therefore define colour not as subjective perception (an *ability* of the visual system, or possibly as a direct imprint of physical phenomena onto perception) Conway [2009] but in terms of the representational prop-erties of a coding system, a system that is designed to systematically code spectral frequency in a way that allows the computations suggested by colour constancy to be readily performed.

Chapter 4

Monochromatic Colour

If it is assumed that light stimuli are monochromatic then any stimulus may be fully described by the two parameters of luminance and spectral frequency. A visual system that is capable of representing luminance and is seeking to develop the ability to represent monochromatic light stimuli more accurately would therefore develop an additional visual parameter to represent the spectral frequency of the light stimuli being measured. The sensors and the processing of the early visual system would therefore have to develop the means to support spectral frequency determination. This must be achieved with sensors that only measure luminance and provide no direct information on spectral frequency. The simplest modification to a sensor that measures luminance is to change the spectral frequency it responds best to, which in natural visual systems is achieved by a physical change to the pigment that is employed to absorb light. Photon absorption and consequently detection is a probabilistic process and therefore sensors generally function within a spectral range, with a declining sensor response with increasing spectral distance. If the distribution of the response is too broad, it may be modified by the use of a spectral filter. Therefore, any system that seeks to determine spectral frequency using the sensors of the type generally available to natural visual systems must make design choices on the following three sensor parameters: (1) the choice of pigment which determines the wavelength the sensor responds best to, (2) the choice of filter which determines the distribution of response to non-optimal wavelengths, and (3) the number of pigment types (number of different types of sensor).

4.1   A Partial Solution to Spectral Frequency

Natural visual systems rely exclusively on sensors that measure luminance within specific bounds of the electromagnetic spectrum. They use sensors which respond most strongly to a specific spectral frequency and whose response declines with increas-
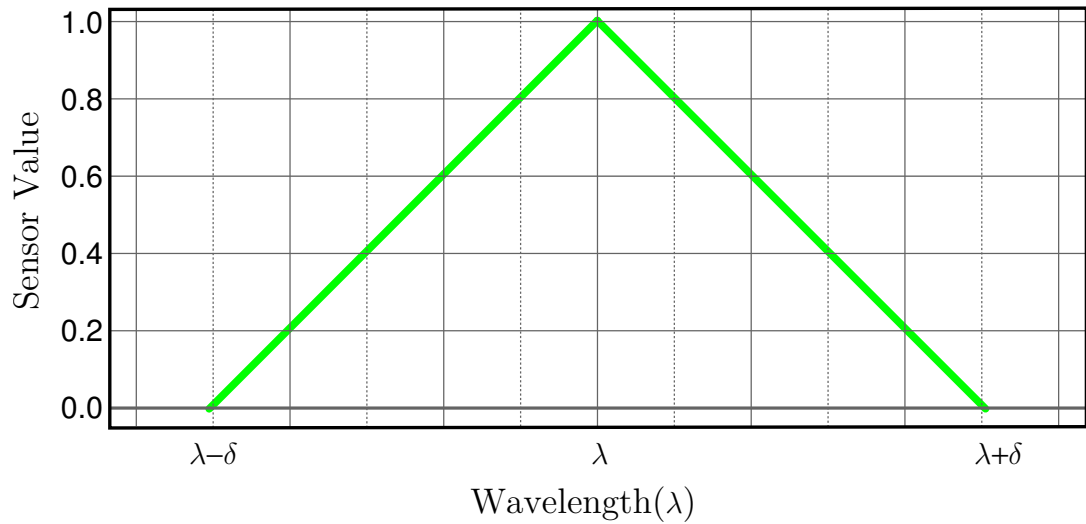
Figure 4.1: Single Sensor Green (G), with a spectral response which is linear within its range. The sensor has a peak response at wavelength λ and a spectral range of 2δ.
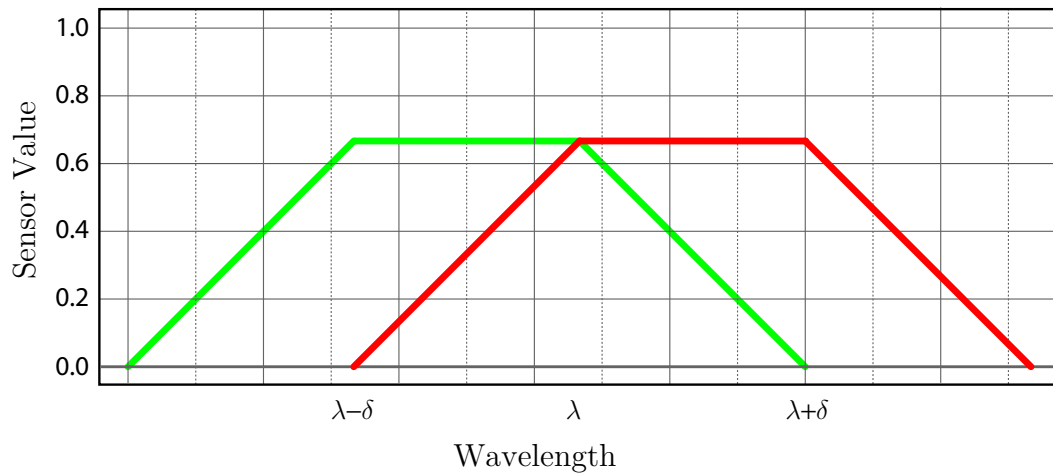

Figure 4.2: Spectrally adjacent sensor pair Red (R) and Green (G), joined at wavelength λ (when $R = G$) with two spectral ranges of δ where one sensor exhibits a linear response to wavelength while the second sensor pair maintains a reference level.

ing spectral distance. An ideal form of this sensor is shown by Figure 4.1, which has a maximum response to a light stimuli of frequency λ, and a response which decreases linearly with increasing or decreasing spectral frequency. For convenience, the response of the sensor may be put in normalized form, between 0 and 1. When optimally stimulated by a light stimulus whose dominant wavelength is λ the sensor will have a value of 1 and beyond the spectral distance from this of δ the sensor no longer responds and will have a value of 0. Assuming the sensor is optimally stimulated by a monochromatic stimulus of a constant luminance and within the spectral range, spectral frequency may be determined with this single sensor by equations 4.1 and 4.2 :

$$f = (\lambda - \delta) + \delta G \tag{4.1}$$

$$f = \lambda + \delta(1 - G) \tag{4.2}$$

This solution suffers from the limitation that any sensor response will lead to two solutions for spectral frequency. A single sensor provides no further information to determine the spectral direction from the sensor's peak response wavelength (λ). Frequency may therefore be determined only if it is known that the wavelength of the stimulus is greater or less than the peak response wavelength of the sensor.

The problem of determining the stimulus direction in relation to the sensor may be solved by introducing an additional sensor that provides a point of reference. A more general solution to determining spectral frequency is therefore to use two sensors with adjacent spectral responses in pairs, with one sensor being used as the point of reference (the *reference* sensor) and the response of the second sensor (the *active* sensor) being used to measure the spectral distance from the point of reference. As one sensor is used as a point of reference with respect to the second sensor, the peak response of this sensor must be broadened to equal the area of response of the second sensor (see Figure 4.2). This allows spectral frequency to be measured accurately between the peak responses of the sensor pair. For an adjacent sensor pair (R and G) which have a joint response at wavelength λ and with an overlapping spectral range of 2δ, frequency may be determined between $\lambda - \delta$ and $\lambda + \delta$ by equation 4.3.

$$f = \lambda + \delta(R - G) \tag{4.3}$$

When a sensor pair is employed in this way the fundamental operation in the process of frequency determination is simple subtraction. This is a form of opponency, which may be referred to as *reference opponency*, and is equivalent to single channel opponency found in the sensor arrays of natural visual systems.

4.2 Normalization

Natural light stimuli vary across a broad range of luminous intensity. This range may be narrowed by the physical design of the sensory system to accomodate optimal sensor stimulation, but it cannot be expected that a luminance sensor can always be optimally stimulated. Reference opponency requires that the reference sensor be fully stimulated for the active sensor value to represent the spectral distance from the reference sensor.

Reference opponency relies upon the linear relationship between the reference sensor and the active sensor. This relationship is preserved within the available range of sensor response. If the sensor used as a reference (S_{max}) differs from the optimal response by a factor γ (so that $\frac{1}{\gamma}S_{max} = 1$) then frequency may be determined by equation 4.4. The sensor used as the reference sensor is the sensor with the maximum response.

$$f = \lambda + \delta(\frac{1}{\gamma}R - \frac{1}{\gamma}G) \tag{4.4}$$

The value of γ may therefore be seen as the value which preserves the original sensor value after compensating for non-optimal levels of luminance. This value is therefore a spectrally restricted (colour related) measure of luminance, a value sometimes referred to as *lightness*. The fundamental operation of normalization is multiplication to determine the normalized value of the measured sensor.

4.3 De-Saturation

Natural light stimuli are rarely monochromatic. Typically, a natural light stimulus will have all the frequencies of the spectrum and therefore will to a certain degree stimulate all luminance sensors. Therefore, for the underlying assumption of monochromacy to be useful, a further assumption must be made. The rigid distinction between equiluminance and monochromacy will be modified. Rather than assuming that a light stimulus is either equiluminous (white) or monochromatic (coloured) a graded distinction between the two categories will be introduced. It may be assumed that any stimulus will consist of a degree of equiluminance and a degree of monochromacy. With this assumption, to accurately describe a natural light stimulus requires the additional parameter of equiluminance, and therefore it is one of the tasks of the chromatic early visual system to separate and measure the degree of the equiluminant component in addition to determining the frequency of the mochromatic component. As an equiluminant light stimulus will stimulate all sensors equally, it may simply be subtracted from the initial sensor values (white subtraction) as an initial processing step. The

proportion subtracted when compared to full equiluminance is usually referred to as *saturation* Lennie and D'Zmura [1988].

Equiluminance stimulates all sensors equally. Therefore, the degree of equiluminance may be determined by the sensor with the lowest value (S_{min}). If all sensors (normalized by γ) are equally stimulated by the proportion ι, with $\iota = S_{min}$ then spectral frequency may be determined by equation 4.5.

$$f = \lambda + \delta \left(\left(\frac{1}{\gamma - \iota} \right) (R - \iota) - \left(\frac{1}{\gamma - \iota} \right) (G - \iota) \right)$$ (4.5)

Equation 4.5 may be simplified by expressing the preprocessing functions of normalization and desaturation separately for each sensor. If $r = \left(\frac{1}{\gamma - \iota} \right) (R - \iota)$ and $g = \left(\frac{1}{\gamma - \iota} \right) (G - \iota)$ then spectral frequency may be determined by equation 4.6.

$$f = \lambda + \delta(r - g)$$ (4.6)

The fundamental operation of desaturation is subtraction. Determining saturation requires an independent sensor (S_{min}) in addition to the reference pair used to determine spectral frequency.

4.4   A Full Solution to Spectral Frequency

As shown above, reference opponency with white subtraction requires at least three sensors. If it is assumed that sensors should always be organized into opponent pairs, then the requirement of an additional sensor would lead to a design with a second sensor pair. If this is used as a second opponent channel it may be used to cover an additional area of the spectrum. As shown by Figure 4.2, a single channel opponent design allows ⅔ of the available spectral range to be used in resolving spectral frequency. Figure 4.3 shows a dual opponent pair design, which allows 60% of the available spectral range to be used. A dual opponent design therefore approximately doubles the spectral range that a single opponent pair design can cover.

If four normalized and desaturated sensors (r, g, b, y) are arranged spectrally adjacent as in Figure 4.3, with the wavelength λ set to the lower boundary wavelength (in the case of Figure 4.3, $\lambda = 430$ and $\delta = 30$) then frequency may be determined either by equation 4.7 or equation 4.8:

$$f = \lambda + \delta(2(y - b) + (r - g) + 3)$$ (4.7)

$$f = \lambda + \delta(2(r - g) + (y - b) + 5)$$ (4.8)

17

If only relative frequency is required then $\lambda = 0$ and $\delta = 1$. If frequency is coded as a single value then a conditional must be evaluated to select the equation to be used. The sensor with the largest value is used as the reference and therefore the condition that determines whether equation 4.7 or equation 4.8 is used is $(y + b) > (r + g)$. Alternatively, if frequency is coded as two signed values then this condition can be avoided. Relative frequency (f') may be determined by directly coding the two opponent values, given by equation 4.9:

$$f' = <r-g\,,\,b-y> \tag{4.9}$$

Figure 4.3: Two pairs of spectrally adjacent sensors red(R)-green(G) and blue(B)-yellow(Y̊). Note: the symbol Y is often used to represent linearly coded luminance (Y' is used to represent non-linearly coded luminance – *luma*). When normalized and desaturated lower case symbols r, g, b and y are used.

Chapter 5

Components of Colour

If spectral frequency of monochromatic natural light stimuli is to be determined from spectrally disparate luminance measures, then the information required to represent the stimuli must include not only the dominant wavelength (hue) but also the degree of luminance normalization (*lightness*) and equiluminant saturation needed. These two latter factors may be considered as additional *channels* of information. Colour information therefore consists of three separate dimensions: *hue*, *lightness* and *saturation*. Taken together, these values represent the information inherent to colour, and the individual steps of sensor pre-processing must reflect the computation inherent in determining these values.

While luminance requires minimal processing for each element of the sensor array to be presented as an information *channel* to the visual system, colour information requires a greater degree of processing for the information to be put into a standard format. The processing required to determine saturation and lightness is only simple (non-repeating) arithmetic but needs to be performed for each element of the sensor array. As visual sensor arrays are typically very large and any processing delay must be minimal these functions are best performed as a pre-processing element of the sensor array itself.

5.1 Number of Sensors and Spectral Distance

Natural visual systems typically organize their sensors into opponent pairs (channels), the simplest form of which is dichromacy. An analysis of dichromacy shows that dichromats suffer from at least one achromatic area of the spectrum Shepard and Cooper [1992]. Single channel reference opponency does not inherently suffer from achromatic areas, but being restricted to two sensors it is unable to determine equiluminance. Any dichromatic visual system that seeks to accurately determine spectral fre-

quency of stimuli that have an element of equiluminance using reference opponency must therefore compromise between determining saturation accurately and determining hue accurately. This compromise necessarily leads to achromatic areas of the visible spectrum. Dichromatic reference opponency is therefore not a viable solution for a visual system that wishes to determine the spectral frequency of natural light stimuli accurately.

The visual systems of most higher organisms capable of general colour vision employ two colour channels. Each opponent channel is fed by two sensors and therefore visual systems that employ dual channel opponency generally require four sensors. Indeed, many animals that possess acute colour vision (such as birds) have been found to use four colour sensors Goldsmith [2006].

The spectral sensitivities of the colour sensors of many species of birds have been precisely measured Goldsmith [2006]. Figure 5.1 shows the relative spectral absorption of the pigments of the four colour sensors commonly found in birds. The spectral response of the sensors that this implies has two important characteristics. The sensors are (with the exception of the long sensor) equally spaced on the spectrum (approximately 60 nm apart) and the spectral sensitivity of the sensors is narrowed (by means of an oil droplet which covers the sensor, acting as a light filter) so that the response of spectrally non-adjacent sensors do not overlap.

5.2   A Normal Sensor

As shown above, reference opponency requires that the sensor used as a reference has a flat response for the reference sensor and a linear response for the active sensor. A linear response allows a simple translation from photon count to relative spectral frequency by the identity function. Spectral frequency may then be determined by simple addition or subtraction from the reference.

Measuring natural phenomena often relies on processes that are inherently probabilistic in nature and therefore such measurements are not absolute but distributed by the inherent randomness of the underlying phenomena being measured. Absorption of photons by a pigment is a cumulative sum of a probabilistic process and as a result any measurement of absorption will tend to be normally distributed. Figure 5.1 shows the spectral absorption distribution of the pigments commonly found in birds Goldsmith [2006]. Figure 5.2 shows an ideal normal distribution. Any system that seeks to determine relative spectral frequency accurately must solve for the respective near-gaussian function. A full solution therefore involves a square root and a natural logarithm. The complexity of these functions precludes their use in simple pre-processing circuitry. An alternative is to transform the normally distributed sensor response into an approximation of a linear response.

A normal distribution will have an approximately linear response for approximately

half of its distribution. An approximate linear response may therefore be achieved by a simple pre-processing step. The non-linear tail must be truncated and the peak response must be capped, as shown by Figure 5.3. When the luminance level for the reference sensor is optimal then the reference sensor value may simply be truncated by a cut-off level, however, when the reference sensor response must be normalized this becomes non-trivial. The difference between the ideal linear sensor response and a response with a normal distribution may be seen as an error which increases in inverse proportion to the measured sensor response. This may then be corrected for by subtracting the normally distributed peak response from the reference sensor value prior to determining γ. If k is the proportion of the sensor response dedicated to reference use (from the point $S_{max} == S$), then an approximate solution to γ is given by equation 5.1:

$$\gamma = \frac{1}{S_{max} - kS_{max}\left(\frac{(S_{max}-S)}{S_{max}}\right)} \tag{5.1}$$

5.3   Dual Channel Opponency

The assumption of desaturated monochromacy means that after white subtraction only two adjacent sensor values will be non-zero. These sensor values will either be equal or one sensor value will be greater than the other. If the sensor values are not equal then the dominant sensor may be used to determine the reference spectral area, and once normalized, the second sensor value may be used to determine the spectral distance from the reference point of the dominant sensor. As the primaries are known values, this allows the spectral frequency to be determined by simple addition. When the sensor values are equal, then this is a special case when the spectral frequency is precisely at the mid-point between the primaries. In this case either sensor value may be used as the primary.

As non-adjacent sensors cannot have a value they may be seen as mutually exclusive. It is therefore useful to place non-adjacent sensors into opponent pairs which may then be used to produce the signed opponent values required to determine spectral frequency.

Tetrachromatic (dual opponent channel) reference opponency therefore allows a general complete solution to the problem of determining spectral frequency. Spectral frequency may be accurately determined for any area of the spectrum covered by adjacent sensor pairs.
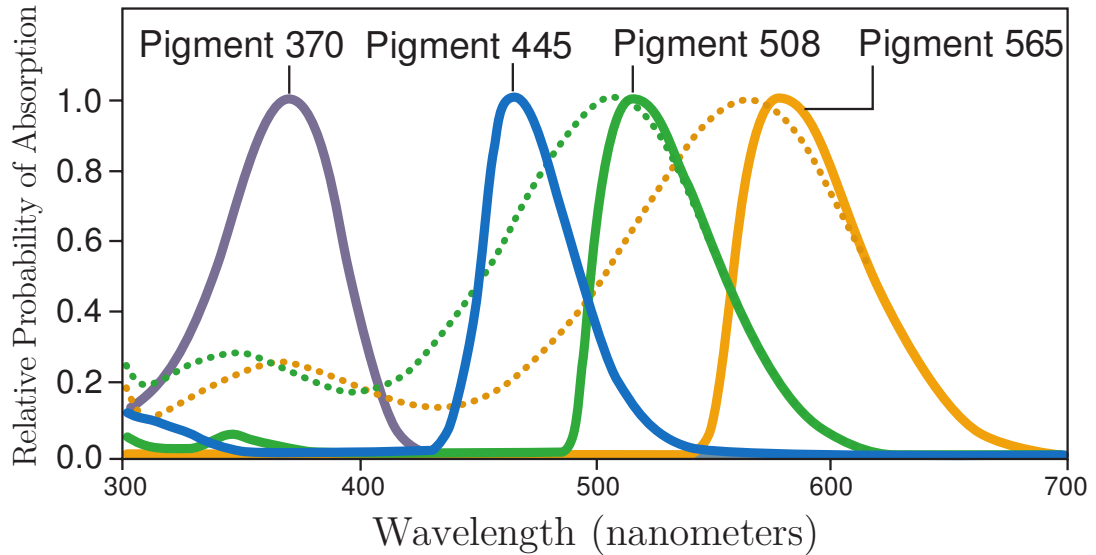
Figure 5.1: Sensor system of birds, with four sensors, demonstrating narrowed spectral frequency distribution (by use of coloured oil droplet) and equal spectral spacing throughout the visible spectrum. The dotted lines show the distribution prior to filtering. Adapted from Goldsmith [2006].
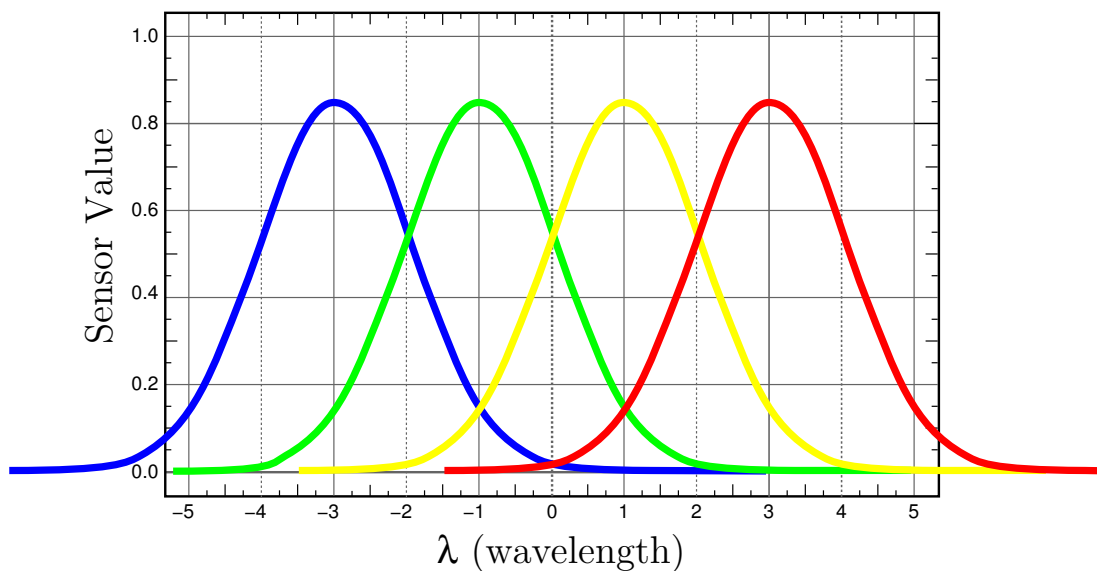
Figure 5.2: Four ideal sensors with peak responses spread equally along the visible spectrum, and with each sensor having an approximately normal distribution in respect of its response to spectral frequency.
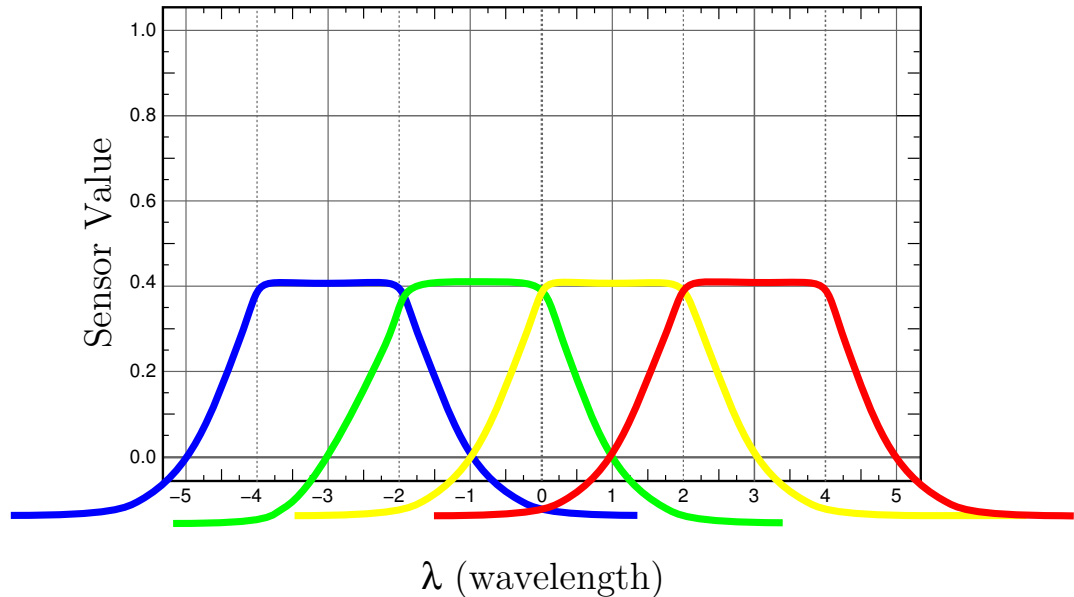
Figure 5.3: Four ideal sensors with an equal spread and with peak response capped and non-linear tail truncated.

5.4 Colour Pre-Processing

With dual channel opponency and a sensor arrangement as set out by Figure 4.3 it is not the individual opponent value itself that determines colour but its relative value in relation to the adjacent sensor value of the other opponent sensor pair. One opponent value is always held at the maximum or minimum reference value while the other opponent pair increases or decreases. Carrying out white subtraction and normalization sets the reference value. This allows colour information to be sent to the visual system in a standardized format. To avoid processing a conditional, hue may be represented not as a single value but as two opponent values, one of which is always at maximum or minimum. In the simplest case, the calculation of hue is simple subtraction. White subtraction and normalization may be achieved (for the most part) by adaptation and other feedback mechanisms that may be built into the sensor array itself. Computationally, normalization requires only division and multiplication. A tetrachromatic dual channel colour opponency sensor array with colour sensor arranged optimally therefore allows accurate spectrum determination with a minimum of computational complexity.

The processing steps that must be performed to determine colour information are:

1. White subtraction – Saturation

23

2. Normalization of sensor values – Lightness

3. Calculate opponent values – Hue

Lightness and saturation should not be seen as discrete values but as proportions which reflect normalization and white subtraction, respectively. Taken together, the values of hue, lightness and saturation preserve the information carried by the original sensor values.

Chapter 6

Colour Processing

Initial colour processing is carried out as part of the sensor array. The information this produces is put into a standard format and then passed to the visual system. All values other than hue may be represented by individual signed values. Hue, however, is derived from two sensor contrast values and to avoid evaluating a conditional it is desirable to continue to code hue as two separate but linked values. Coding hue with dual opponent values requires a specific mechanism for colour calculation.

6.1 Colour Arithmetic

To link the two hue opponent values the arithmetic operations must be defined. This will be shown in terms of the primitives of computation: increment, decrement and conditional. Let C_B refer to the blue-yellow opponent channel (chroma blue), C_R to the red-green opponent channel (chroma red), the constant Max to the maximum/minimum chroma values and the operation $++ / --$ as the increment / decrement operator. Incrementing hue may then be implemented by the following ordered set of rules (expressed using conventions of the C programming language):

1. $if\ (C_R == +Max)\quad C_B --$
2. $if\ (C_B == -Max)\quad C_R --$
3. $if\ (C_R == -Max)\quad C_B ++$
4. $if\ (C_B == +Max)\quad C_R ++$

This algorithm has the important property of being stateless and complete. No information needs to be stored for the algorithm to correctly increment an arbitrary hue value, and an increment of any arbitrary valid hue value will always result in a valid hue value. This algorithm also has the important property of circularity. Incrementing

'red' ($C_R = Max, C_B = 0$) will lead to 'blue' ($C_R = 0, C_B = -Max$), whereas decrementing blue will lead to red. A decrement may be implemented simply as the inverse of increment. All other arithmetic operations can subsequently be defined as a finite sequence of increments and decrements.

Practical implementations of the increment required by this algorithm may avoid the evaluation of a conditional to determine a maximum or minimum state by the use of overflow or underflow. An increment or decrement of a maximum or minimum value can simply leave the value unchanged.

This mechanism for coding dual channel opponency is arbitrarily extensible. Each opponent channel may be seen as numerical digits of a finite circular system to represent numbers, of which dual channel opponency is the simplest case. At each stage only two channels would be active (one marking the condition and the second counting), but once the right active channel reaches the +Max state and the left channel reaches 0, the two active channels shift one channel to the right. Linear increments or decrements for *n*-channel opponency would therefore flow two channels at a time across the available channel space and by circularity reach the initial state once all the available channels have been used.

This therefore demonstrates that coding spectral frequency by means of multiple opponent values is equivalent to coding it by a single value. Using one channel as a placeholder, however, involves redundancy and this allows additional information to be stored. A valid hue requires one of the opponent channels to be at a maximum or minimum. A zero state could therefore be used as a special code-word designated for stimuli which have no colour information. Multiple channel opponent coding also inherently divides the code space into discrete regions (or *primaries*) and this information would be lost if hue is conflated into a single value.

6.2   Coding the Spectrum

The physical electromagnetic spectrum is a one dimensional continuum from the very long wavelengths to the very short. Dual channel opponency is able to accurately code spectral frequency for monochromatic stimuli within the inner bounds of the sensor range. If the sensor output is coded directly into the opponent channels (aside from white subtraction and normalization) then any stimuli that are out of range would be coded either as short (*blue*) or long (*red*). It would be expected that as the wavelength of the light stimulus is increased or decreased the hue would remain constant, with luminance gradually declining as the wavelength of the stimuli moves out of range of the sensors. Perceptual studies in relation to colour opponency, however, show that this prediction fails Neitz and Neitz [2008]. Figure 1.2B shows how monochromatic wavelength stimuli would be expected to be perceived under sensor direct dual channel colour opponency and Figure 1.2C shows how perceptual studies show them to

be perceived by human subjects. With respect to light stimuli on the short end of the spectrum, perceptual studies show that as spectral wavelength decreases beyond the opponent boundary the hue continues to change systematically.

For human perception the boundary colours are blue for short wavelengths and red for long wavelengths. If the wavelength of a blue light stimulus is increased then it will perceptually appear to be increasingly green; by the degree that spectral wavelength is increased. This is linked to the fact that the *green* sensor is spectrally adjacent to the *blue* sensor. The proportion of green will increase until it is precisely equal to blue, at which point the stimulus will be perceived as blue-green (the named colour cyan). If wavelength is increased beyond this point then the proportion of blue will decline until the stimulus appears a pure green. Colour perception therefore suggests an inherent relationship between the primary colours of perception. Any primary colour may be said to be adjacent to one primary of higher wavelength and to one of lower wavelength, and perceptually there is a systematic transition between any two adjacent primaries.

Any hue other than a primary is therefore an intermediary colours between two adjacent primaries. This adjacency is a reflection of the spectral arrangements of the sensors themselves. This relationship however fails at the sensor boundaries. As the wavelength of a blue light stimulus is decreased it will go out of the range of the opponent sensors and therefore no transition would be expected. Studies of human perception, however, show that as wavelength of a blue light stimulus is decreased it will appear to be increasingly *violet*, which is a named colour indicating a mixture of blue and red. As wavelength is decreased from blue, the proportion of red increases. This indicates a transition from blue to red and therefore indicates red and blue to be adjacent primaries.

Although the *red* sensor is often portrayed as not being able to provide useful information at short wavelengths (due to the spectral distance), direct spectral absorption studies of the *red* sensor pigments (as shown by Figure 6.2) indicate an increase in absorption into all the short wavelengths. After decreasing as expected with decreasing wavelength, there is shift to a small but measurable increase in absorption beyond approximately 430 nm. While this increase does not affect the measurement of luminance, it contrasts with a continued decrease for the *green* pigment. Significantly, the absorption for *red* goes above *green* at wavelengths lower than approximately 430 nm. Opponency is particularly sensitive to this kind of change in direction, with the crossover point being the point of maximum sensitivity for this type of measurement (see Section 8.2). Nevertheless, the red-green opponent value taken directly from the sensors will be very small in comparison to the the blue-yellow value. Despite this, monochromatic stimuli of these wavelengths are perceived to have an increasing red component as the wavelength is decreased. This indicates that the perception of the colour red is not purely related to its associated sensor value but must in part be a product of the underlying processing. The perceptual response in this case indicates conditional processing which amplifies the red-green value. The purpose of this pro-

Figure 6.1: Idealized sensor response across the visible spectrum of a dual channel tetrachromatic visual system. The sensors are blue (B), green (G), yellow (Y̊) and red (R), with peak sensitivities at 460 nm, 520 nm, 580 nm and 640 nm. With full circularity.

cessing in this case is to impose an adjacency between the blue and red primaries, a property that is not supported either by the physical spectrum or the way the sensors respond to the spectrum. If short wavelengths are perceived to be adjacent to long wavelengths this must therefore be a product of the underlying coding mechanism, and indicates that the code must be circular.

One of the features of dual channel opponency coding is the incremental transition between opponent channel maxima or minima, where one channel is held at a maximum or minimum value while the other codes for a transition value. Dual channel opponency therefore has four poles (or *primary* positions) and by virtue of the proposed rules of dual opponency colour arithmetic each pole is linked by an incremental transition. Increments will lead to the adjacent primary of a higher wavelength and decrements will lead to the adjacent primary of a lower wavelength. As any primary may transition by increment or decrement, this code is inherently circular (see Figure 6.1). Employing a circular code to represent a linear physical spectrum requires that the outer extremes of the perceptible spectrum must be joined by an artificial meridian.

It would be expected that if circular opponency coding is implemented without a transition sensor that the perception of a monochromatic stimuli which falls on the meridian could be simulated by stimulating the outer sensors independently (with blue and red light stimuli). If the human visual system uses dual channel opponency with circular coding it would be expected that stimulating the short and long colour sensors in equal proportions would produce a percept mid-way between blue and red, that is a colour equal in proportions of red and blue. Perceptual studies show there exists a named colour magenta which is a transition colour equal in proprotions of red and blue and which can reliably be produced by dichromatic stimuli of blue and red light in equal proportions. However, this colour cannot (in most humans) be produced by a monochromatic stimuli (and is therefore referred to as *extra-spectral*). Other colours within the meridian that are closer to the primary can be produced by a monochromatic light stimuli but all have perceptually identical metamers produced by dichromatic stim-

uli. All triphosphor display technology relies on this equivalency. This indicates that the human visual system uses a circular colour code for hue. Transitional metamers indicate that no sophisticated processing is employed to distinguish dichromatic stimuli from monochromatic stimuli, and therefore this suggests that the coding is performed by the early stages of visual processing. Furthermore, the existence of extra-spectral colours indicates that the bridging function which links the short wavelength region to the long wavelength region has only limited functionality and as a result leaves a meridional spectral gap. That this gap can be closed by artificially produced dichromatic stimuli to produce previously unknown percepts provides further evidence in support of the hypothesis that the underlying code for colour is fully circular.
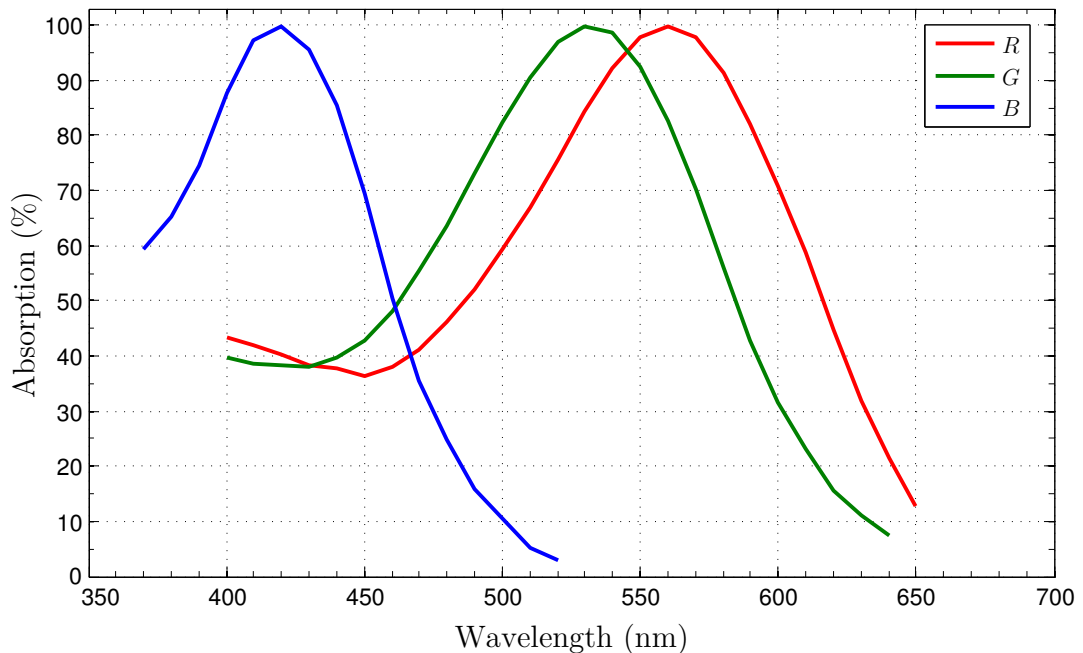

Figure 6.2: Human sensor pigment absorption. Short sensor (blue), medium sensor (green) and long sensor (red). Note slight increase in absorption of the long (red) sensor in the short wavelength area of the spectrum. Adapted from Dartnall et al. [1983].

Chapter 7

Dual Opponency with Three Sensors

Dual channel opponency requires four sensors, one pair for each opponent channel. Despite the ubiquity of natural visual systems which have a full complement of four colour sensors there is one important class of animal which due to its nocturnal past has lost important elements of its visual system that relate to colour. Some mammals such as the star-nosed mole (Condylura cristata) have lost their visual systems entirely, but most of those that re-emerged into the luminant world from their underground burrows and nocturnal hunting grounds to take advantage of the KT-Extinction event are limited to single channel dichromacy Goldsmith [2006]; Neitz et al. [1989]. It is only in apes and old world monkeys that dual channel colour vision has re-emerged Neitz et al. [2001], a development that is relatively recent and remains as yet rather tenuous Nathans [1989]. Many individuals within these populations (including the human population) remain dichromats or suffer from chromatic anomalies DeMarco et al. [1992]; Mancuso et al. [2010]. The difficulty these animals have faced is that although the dual channel opponency mechanism remained functional Mancuso et al. [2010], the sensors to drive one of these channels were in the millions of years of darkness [Vorobyev, 2006, 168] and ensuing genetic drift irretrievably lost Goldsmith [2006]. This has meant that these sensors have had to be redeveloped anew Solomon and Lennie [2007]. This is a process of gradual step by step differentiation, and it is a process that in mammals still remains in the early stages Nathans [1989]. In mammals where dual channel opponency has been restored, only one sensor has thus far been differentiated. Nevertheless, the underlying mechanisms of dual channel opponency have not been lost and as a result these animals have found a way to simulate a four sensors system with only three physical sensors. This simulation is not trivial and involves significant additional complexity in the early stages of visual processing.

If we assume the simplest case of simulating the inputs of dual opponency with only three sensors, then one solution is illustrated by Figure 7.1. We assume a standard

dual channel opponency sensor arrangement (Figure 6.1) in which the mid-long (*yellow*) sensor has been lost. This sensor value would then need to be simulated by use of the remaining three sensors. The solution given by Figure 7.1 broadens the distribution of the neighbouring sensors to adequately cover the spectrum. This solution has two difficulties: accuracy is reduced by half and the missing sensor leads to a normalization gap. Loss of colour accuracy is an acceptable side effect, but normalization is an essential element of opponency. Simple normalization requires the redundancy of sensor overlap, which cannot be achieved with only three sensors. It may be observed, however, that the sensor responses may be arranged so that they are proportional when they do not overlap. It is possible therefore to normalize on the linearized difference ratio of adjacent sensors; with an approximate solution being given by equation 7.1 (if $G > R$) and equation 7.2 (if $R > G$). Calculating the difference ratio requires a division and linearization (conventionally an exponent).

$$\left(\frac{2(G-R)}{3G} \right)^{1.35691544885674}$$
(7.1)

$$\left(\frac{3(R-G)}{2R} \right)^{1.35691544885674}$$
(7.2)

Once normalized and desaturated, raw sensor values must be translated into the four sensor format. As the sensor responses have had to be broadened these raw sensor values are no longer suitable for determining opponent values directly. Therefore, in addition to computing the fourth sensor value the remaining sensor values must also be transformed into the original form. For the solution shown by Figure 7.1 this may be done by the following algorithm:

1. $r = (R - 0.8) + (2B > G) : (1.6 - G)|0$

2. $g = (G - 0.8)$

3. $b = B$

4. $y = (R > G) : G|R$

The missing yellow sensor value is computed by a conditional. It is assigned either the green or the red sensor value depending on the condition of whether green is greater than red. One side-effect of broadening the distribution of the red and green sensors is that the overlap with the blue sensor allows circularity to be implemented by simply subtracting the green sensor value from the blue when green is greater than blue. For a detailed example see Appendix A.

Figure 7.1: Abstract representation of method used to interpolate the fourth sensor value (*yellow*) from three physical sensors (RGB).

Chapter 8

A Colour Model

8.1 Luminance Information

Determining the spectral frequency from a non-optimal luminance measure requires placing the sensor information into a standard form. The sensor values must be normalized (based on the maximum sensor value) and have any element of equiluminance removed (based on the minimum sensor value). The normalization value is therefore a spectrally restricted measure of luminance, and may be referred to as colour luminance or *lightness*. The degree of equiluminance may be referred to as *saturation*. As these elements are removed from the raw sensor values prior to the opponent values being determined, lightness and saturation information must be preserved to allow an arbitrary light stimulus to be fully represented. Colour information for a visual system that employs opponency must therefore include lightness and saturation. However, colour luminance must be kept distinct from general luminance information. Under opponency coding a stimulus may be coded as having no colour information, and in this case the only information that is sent to the visual system is luminance information. Colour luminance codes for a spectrally restricted measure of luminance and therefore preserves spectral information whereas general luminance is a weighted sum across the entire visible spectrum and as a result any spectral information is lost. Furthermore, natural visual systems do not measure luminance as an absolute measure, but as local luminance difference between adjacent sensors (or adjacent groups of sensors). Information is sent to the visual system only if a stimulus triggers sufficient luminance difference. This reduces the amount of information that must be sent to the visual system, which is in most natural visual systems necessitated by the fact that the link between the sensor array and the visual system is only able to carry a fraction of the information the sensor array is able to produce. Colour information by contrast, even though it is calculated by taking sensor differences, is spatially an absolute measure and therefore ubiquitous wherever there is a spectral difference even if that spectral difference is uniform. To limit redundant colour information being

sent to the visual system it is possible to link luminance and colour information so that colour information is sent to the visual system only when luminance information is sent. Colour information may, therefore, be seen as additional information that may be attached to luminance information if available.

8.2 Information Measurement

Information produced for any visual system starts with measurements of light stimuli. When the sensors that take the measurements are of poor quality there is a technique often used in physics called the *null* or *difference* method Hurvich [1997]. The simplest example of the difference method is the balance scale, where the weight of an unknown object may be determined by balancing it (that is, nullifying its weight) with objects of known weights (on the opposing side of the balance). Once an equilibrium is reached, the unknown quantity can be precisely determined in terms of a known quantity. This principle finds its widest application in electrical measures where quantities such as voltage or resistance can be precisely measured by *balancing* the current flow through a circuit to zero.

The same principles of the null method may be employed in the measurements required by opponency (which are inherently a measurement of the electrical quantities when electro-magnetic wave quanta are absorbed by a pigment material). By this means a poor quality sensor that is unsuited to making reliable absolute measurements may be used to make precise relative measurements. Natural visual systems use opponency both for luminance and colour measurements. The zero state for natural visual systems is therefore not the absence of a stimuli but the presence of two stimuli that balance the opponent sensor pairs. Under this design, luminance information is a measure of spatial imbalance on the sensor array and colour information is a measure of spectral imbalance.

8.2.1 Accuracy of Measurement

The measurement of the degree of imbalance in an electrical system whose natural state is equality is inherently non-linear Anderson [1995]. A measurement of an imbalance is a measurement of the ratio of difference. The accuracy of a measurement is limited by an imbalance which is at the borderline of detectability (j), and if this is constant throughout the range of measurement then the discriminable difference (ΔY) within the range (Y) may be given by Equation 8.1.

$$\Delta Y = j^{Y+j} - j^Y \qquad (8.1)$$

The ability of a difference method to measure a stimulus accurately therefore declines as a power function of the magnitudes being compared. The human visual system is able to discriminate approximately a 1% luminance difference (a *just noticeable difference*) over a contrast range of magnitude 100 Poynton [1993]. A system that is capable of direct measurement with a constant limit of discrimination throughout its range would produce 10,000 unique equidistantly spaced measurements. The method of difference has the same initial ability to discriminate but this ability decreases as a power function of the magnitude. Solving Equation 8.1 for $\Delta Y = 1$ and $j = 0.01$ gives a maximum of approximately 463 unique measurements. To represent such measurements would require code words 9 (binary) bits in length. An 8 bit code would be able to represent an order of magnitude less and a 10 bit code would be able to support two orders of magnitude more (with a $\Delta Y > 200$).

Direct measurement of luminance is ideally linear. A modern interpretation of Grassman's law shows that human luminance (and colour) perception is also linear. By coincidence, the way that one of the most common display technologies (the CRT) produces light almost exactly mirrors the non-linear way the human visual system measures light Poynton [1993]. The distinction must therefore be drawn between linear measurements of luminance and the perceptual linearity placed onto an inherently non-linear system of measurement. The human visual system codes luminance measurements in an inherently non-linear way but treats the resulting code as if it were linear. This is mirrored by the CRT, which when given a non-linear control signal displays levels of luminance that appear (to the human visual system as) linear. On this basis it is often assumed that the human visual system measures and represents luminance (and colour) linearly, and this has led to much error and confusion in respect of the terminology used to refer to luminance Poynton [2003b]. As a consequence the convention of referring to a non-linear coding of luminance as luma (Y′) has been established to distinguish it from linear luminance (Y). The mapping between linear and non-linear luminance is generally referred to as *gamma*. Luma may be mapped to luminance by the use of a power law with the inverse of *gamma* as the exponent. A *gamma* of between 2 and 3 is typically used with respect to transmissive colour reproduction Poynton [2003b].

The difference method can be used to measure very accurately at low magnitudes. The early visual system in humans has been found under certain (dark adapted) conditions to be capable of discriminating at a level of accuracy up to that of individual photons Rieke and Baylor [1998]. Photons are irreducible discrete quanta of electromagnetic radiation which can only be measured by their absorption, which is an all or none process. The early human visual system is therefore capable of acting (under certain conditions) as a perfect photon counter . With a difference measurement, the trade-off with the degree of precision that is possible at low magnitudes is that for large magnitudes the measurement becomes increasingly imprecise. This imprecision stems from the fact that the interval between unique measurements increases logarithmically as a function of magnitude.

A visual system that relies upon difference measurement must therefore code for an inherently logarithmic sensor response (treating them as if they were linear throughout) or it must correct for the non-linearity in the early stages of sensor processing. The latter requires the evaluation of an inverse logarithm (a root between square and cube), which is computationally non-trivial as it must be performed for each individual measurement. As the information is produced by an initial logarithmic measurement, to linearise it would also require a larger code space. A visual system that is capable of a 100:1 luminance contrast range and a *just noticeable difference* of 1% would require code words of 14 bits. Maintaining a non-linear coding therefore allows a broad range of luminance measurement to be coded more compactly. Due to the amount of information that must be processed and stored, a primary requirement for any visual system is that information must be stored as compactly as possible.

Perceptual studies show that the sensor response of the human visual system is best modelled as a power function, and that the *gamma* value introduced by the process of taking accurate measurements is maintained throughout all further stages of visual processing, including perception Poynton [1993].

All of these factors indicate that the design of a visual system should employ logarithmic luminance coding. The inherent complexity in this can be substantially reduced by incorporating it into the earliest stages of sensor processing and by maintaining the information consistently in this form.

8.3   Colour Information

In natural visual systems all colour information originates with spectrally restricted luminance measures. The same principles in respect of taking and coding such measurements therefore apply as with general luminance measurements.

The study of perception and colour reproduction show that the nature of colour information is more complex than luminance information. A luminance model needs little more than a specification of how luminance measurements are to be coded. Colour information, however, even when based on luminance measurements is not luminance information (that is, it is not information about how bright a light stimulus is or alternatively how bright specific spectral regions of a light stimulus are) and therefore the nature of colour information first needs to be specified. A colour model is such a specification, and its primary role is to impose a higher level paradigm onto the underlying nature of the phenomena being measured. Such a paradigm does not necessarily reflect the true nature of the phenomena (electro-magnetic radiation) as it may in part reflect the use that is made of the information. The needs of colour reproduction have led to the development of a variety of colour models, which may be divided between the reflective and transmissive means of colour reproduction. Alternatively, studies of perception and colour matching have led to colour models which attempt to provide a

standardized perceptual colour matching space. The information that is required for colour reproduction are instructions for the production of light stimuli (the *primaries*, produced directly by transmission or indirectly by absorption). How the light stimuli are perceived is not specified and is left implicit to the specific technology used to produce the light stimuli. The colour information produced by these colour models therefore does not predict colour perception and this leads to difficulty in translating between colour models where the underlying technology is incompatible. To resolve this, a colour matching space is used as an intermediary colour model. These colour models therefore define colour information either implicitly or explicitly only in relation to the human visual system, and therefore do solve the fundamental problem of specifying the nature of colour information.

Following the known design elements of natural visual systems leads naturally to the central hypothesis proposed by this thesis that colour information is the attempt to effectively code spectral information. The neural circuitry of natural visual systems show how this may be done. A double opponency design indicates that it is possible to effectively code monochromatic stimuli within the visible spectrum by dividing the spectrum into opponent regions (or poles) that are linked by a linear transition between neighbouring regions. Where the poles reside on the physical spectrum is arbitrary, and may vary with the needs and vagaries of the specific visual system, but the central requirements of such a system are that the poles are equidistantly placed on the area of the spectrum being measured and that there is a uniform linear transition between them. This therefore establishes the paradigm of how colour information codes for monochromatic spectral frequency. Dividing the spectrum into equidistant poles and establishing the transition relationship between them lead to an inherently circular representation of the spectrum. The coding of colour information that is done by the visual system is therefore done with reference to a circular spectral multi-polar colour map (or *colour wheel* – see Appendix D for a brief history).

Any system that aims to determine spectral frequency information using only luminance measurements must first normalize the information with respect to a known point of reference and at the same time remove any element of equiluminance. This information is therefore fundamental to the design of a system that seeks to determine spectral frequency using luminance measurements. This leads to a paradigm of colour as a three dimensional space: with *hue* representing the spectral wavelength, *saturation* representing the degree of equiluminance and *lightness* representing the degree of normalization.

Hue, *Saturation* (sometimes called *chroma*) and *lightness* (sometimes called *value*) are the fundamental dimensions of human colour perception. The double opponency design presented by this thesis shows that these parameters are not arbitrary choices and allows the relationship between luminance sensor measurements and the three parameters of colour to be defined without reference to human perception. Specifying that all light stimuli are assumed to be monochromatic, normalized and with an element of equiluminance, is a colour paradigm (a *colour model*). It is an attempt to simplify

the underlying complexity of the phenomena being measured. Coding monochromatic stimuli using a circular polar colour wheel with the added dimensions of *saturation* and *lightness* is an information paradigm placed upon the underlying spectral power distribution of electromagnetic radiation. Its purpose is to simplify the underlying complexity of the phenomena. The design parameters are information, computational complexity and accuracy. Any colour model would aim to minimize the amount of information and the complexity of producing the information while at the same time maximizing how accurately the information represents the underlying phenomena. There are two aspects to the underlying phenomena. The first is the set of all possible light stimuli within the specified range and the second is the set of light stimuli found within a specific environment along with the probability of encountering the stimuli. Practical systems may choose to restrict themselves to the latter. Even with this restriction there are many possible colour models. The claim made here is that a colour model that is as simple as possible and involves the least amount of computational complexity which is able to represent the most probable light stimuli as accurately as possible is the colour model chosen by natural visual systems. This claim will not be proven in this thesis, but is set as a subject for further study.

The design choices for a visual system are therefore primarily predicated upon the nature of the sensor, with the available computational resources as a secondary factor. The sensors used by natural visual systems are the simplest of possible sensors and therefore retracing the design choices made by natural visual systems leads naturally to an optimal minimal colour model.

8.4   Ubiquity of Colour Information

All colour information is based on spectral contrast measurements rather than spatial contrast measurements. As sensor measurements which provide no useful contrast (that is, which are below the threshold) are discarded, the set of luminance information measurements sent from the sensor array to the visual system may be said to be *sparse*. Furthermore the size of this set of information may be said to be related to the information content of the *image* presented to the sensor array. Colour information on the other hand is ubiquitous; that is, not inherently bounded by information content of an image. One means to achieve such a bound is to simply link colour information to luminance information. Each *colour sensor* may be linked to a *luminance sensor* and if the luminance measurement is below threshold then both colour and luminance measurements are discarded, making colour information at least as sparse as luminance information. This linkage is not necessarily bi-directional; luminance may on the other hand not require an associated colour measurement (indeed this is a requirement for backward compatibility with luminance-only designs or dual function luminance/colour designs). This solution to the ubiquity of colour information, however, has the drawback of failing to produce information in conditions of equilumi-

nant colour contrast. Equiluminance exists only in respect of the degree of precision luminance is measured, and therefore the problem of equiluminance may be alleviated by increasing the precision of the luminance measurements.

Studies of perception as well as physiological studies show that the human visual system does indeed have difficulty discerning visual detail in conditions of equilumance Livingstone and Hubel [1988]. Moreover, studies in image compression have found that the amount of colour information can be reduced by approximately half in respect of the luminance information with no loss in perceptible image quality. This suggests that the human visual system measures luminance with approximately twice the precision that is measures colour. Studies of sensors used by the human visual system for colour vision indicate that they are able to discriminate approximately 128 levels Greenwood [2012]. Studies of human perception have found that the number of perceptible levels (expressed in binary) of luminance is between approximately 128 and 512 Poynton [1993]. Estimates of the number of colours that the human visual system can discriminate vary widely, but typically range between one million [Poynton, 2003a, 229], two million McCamy [1998] and ten million Greenwood [2012]. Considering that all colour information is the result of the measurement of two opponent values, and assuming that luminance is measured with a precision of 256 levels (an 8 bit code) then colour must be measured at a precision between 64 and 128 levels (6-7 bits). This gives a range of between one and four million colours. The development of digital image reproduction has led to colour standards which use typically 24 bits to code colour information. A visual system designed on the basis of conventional technology might therefore employ 10 bit luminance measurements and 7 bit colour measurements, for a total of approximately eight million colours.

Difficulties in discerning visual detail in conditions of equiluminance indicate that the limiting factor on colour measurements by the human visual system is likely to be a linkage to luminance. Colour information is produced only as part of luminance information, and therefore is as sparse as luminance information. Furthermore, while colour is linked to luminance, luminance is not necessarily linked to colour. The disparity in colour discrimination in relation to luminance indicates that colour information is produced only for approximately every fourth luminance contrast sensor pair and therefore is four times as sparse as luminance information. While all the colour information may be represented with two opponent values, as shown above, it must be divided into four separate values. This does not increase the amount of information, as the four values may be coded with reduced precision (resulting in no increase of information). If sensors are placed into groups that produce four luminance measurements for each colour measurement and the colour measurement is linked to the measure of luminance, then the amount of information produced by a visual system that supports colour is approximately twice that of a visual system that supports only luminance.

Chapter 9

Discussion

Opponency allows colour to be defined as a code which is computed in relation to a property of the physical world. If the code accurately reflects the measurement of the physical property then the perception of colour may be said to be a measure of that physical property. How sensor responses are mapped into the code is, however, derived from a relative measurement of difference between sensors that differ in how they respond to wavelengths of light. Colour perception is therefore dependant upon the ability to calibrate with respect to the spectral distribution of the physical sensors employed. Natural visual systems have no a-priori knowledge of the spectral sensitivities of their luminance sensors and therefore they are strictly limited to a relative representation. Colour is therefore the product of natural visual systems attempting with limited sensor means to represent spectral frequency information about light stimuli.

The aim of natural visual systems, however, is not to accurately measure spectral frequency as an end in itself but to code spectral frequency effectively to meet the needs of a general visual system that provides useful information to an organism about the external environment. The needs of such systems require colour calculations to be readily and efficiently performed (particularly with respect to the computational limitations of the system) and the code must reflect this. Opponency coding divides colour into the opponent poles (the primaries) and sets the arithmetic relationship between these poles. This relationship is inherently circular and therefore this circularity is imposed upon the physical spectrum, which is inherently continuous and linear.

The aim of colour vision is therefore not only to determine and code for the spectral frequency of monochromatic stimuli accurately but to impose a useful model that provides information within a coding framework used by the general visual system. Coding colour information from raw sensor values is only the first stage of how visual information is processed. Visual systems typically require information that is constant irrespective of global factors such as the spectral properties of the illuminant, something which cannot be corrected for by local pre-processing within the sensor array. Correcting for such factors requires higher level layers of processing, with each layer

of processing lessening the relation between the raw sensor values produced by the stimuli and the code values produced by the visual system to represent those stimuli. Colour may therefore be said to reflect spectral frequency within an idealized model of the physical world specific to the visual system employed by the organism. Under ideal conditions the mapping from the initial opponent values computed from the raw sensor input to the code values used to produce the final percept is the identity function. However, if the conditions are not ideal the intermediary processing framework intervenes. The complexity of this processing, however, should not distract from the essential nature of colour information; which is to usefully and accurately represent light stimuli.

Human colour vision is based on three spectral frequency-differentiated luminance sensors, and therefore it is often said that it is *impossible* for us to know what the perception of colour for animals which possess a greater or lesser number of such sensors is like Goldsmith [2006]. Under the model of colour vision presented here the human visual sensory arrangement is simply an aberrant variation of the standard dual channel tetrachromacy. It is proposed that visual systems should not be classified by the number of colour sensors but by the number of colour opponent channels. The human visual system, it is proposed, is a dual channel colour opponent system and therefore differs only from other dual channel opponent systems in being limited to three physical sensors rather than the required four sensors. It is predicted that an important element of early visual processing in humans would be to interpolate four sensor values from the available three physical sensors. This introduces additional complexity and reduces the accuracy at the early sensor processing stage but the format of the information produced by the early visual system would be identical to that of any dual channel colour opponent system.

Birds are one group of animal which possess a well developed visual system with a full complement of four sensors. The spectral range of one of these sensors extends into the ultra-violet frequency range and it is on this basis that the claim is sometimes made that birds "see colours that we cannot even envision" Goldsmith [2006]. The set of bird sensor peak spectral sensitivities are typically [370 nm, 445 nm, 508 nm, 565 nm] Goldsmith and Butler [2005], which is very close to an equidistant distribution with a spread of about 60 nm. Bird sensors also have a built-in colour filter (an oil droplet) Hart et al. [2000] which narrows the absorption spectra to almost exactly the requirement of the proposed double opponency (that is, no spectral overlap between non-adjacent sensors). The set of three human sensor values are [424 nm, 530 nm, 560 nm] Dacey [2000], Dartnall et al. [1983], DeMarco et al. [1992], but these are unfiltered and therefore have a much broader response (see Figure 6.2). This makes it possible to interpolate a value for a missing sensor, but requires that the sensor response values themselves must be modified before they can be used as dual opponent values (as the distribution is too broad). If the human visual system employs dual opponent coding then not only would the value for a fourth sensor have to be interpolated but the values from the three physical sensors would need to be recoded to

present a set of *virtual* sensors to the early visual system that meet the requirements for opponency.

The human visual system relies upon three physical sensors and therefore any system of colour reproduction that seeks to reliably reproduce colour perception may limit itself to stimulating the physical sensors. On the other hand, if the early visual system seeks to emulate a four sensor system then it would be expected that the relationship between three primaries set to the wavelengths of the three physical sensors exhibit perceptual non-linearities when mapping the visible spectrum.

It has been observed in the development of colour display technologies calibrated for human colour vision that the use of primaries that match the physical human colour sensors does not produce perceptually accurate colours. The development of display technologies (designed for perceptually reliable image reproduction for the human visual system) has led to the use of a variety of primaries, which have changed as the understanding of the human visual system has increased. The primaries for systems of subtractive colour reproduction historically have been [430 nm, 530 nm, 630 nm] Hunt [1995]. Once additive colour reproduction became possible, the primaries chosen for the colour cathode ray tube were [450 nm, 530 nm, 610 nm] (see Figure C.1). These systems rely on dyes or phosphors where the choice of spectral range is restricted by the availability of suitable chemistry. The more recent development of alternative luminance sources such as high brightness light emitting diodes (LED) have allowed greater choice in the selection of primaries. Figure C.2 shows the spectral response of a multi-colour LED commonly used for colour projection, with primaries at [455 nm, 518 nm, 635 nm]. It may be seen therefore that with the traditional (and the most inaccurate) method of colour reproduction the primaries selected were positioned with the twin aims of matching the primaries with the human sensor sensitivities and positioning the primaries approximately equidistantly on the spectrum. This leads to a compromise where the red primary is shifted from 560 nm to 630 nm, a distance of 70 nm.

As colour reproduction technology evolved (with the aim of reliably reproducing colour more accurately attuned to the human visual system) the selection of primaries has shifted; with the spectral distance between the blue (short wavelength) and the green (medium wavelength) primaries decreasing and the spectral distance between the green and red (long wavelength) primaries increasing. This spectral shift has led to the spectral distance between the blue and green primaries to be approximately 60 nm and the spectral distance between green and red primaries to be approximately 120 nm (see Appendix C for further technical details of primaries used by a variety of display technologies). This distance is consistent with the hypothesis of a fourth *virtual* primary between the green and red primaries.

Colour reproduction has traditionally employed systems that are trichromatic and therefore produce, reflect or absorb light intended to be received by a visual system which employs three colour sensors. This has led to most colour coding systems linked to

these display technologies also being trichromatic. However, in the case of image transmission where bandwidth is restricted and therefore where the design of any optimal coding system necessarily reflects the underlying mechanics of the human visual system (based on the requirement to represent only information that is needed), the use of dual channel opponent coding is common. Compression algorithms such as JPEG, MPEG, and H.264/AVC use a dual channel opponent code to represent colour (or *chroma*) information. The fourth primary inherent to these codes is typically simulated by trichromatic displays with the use of the neighbouring primaries (in equal proportions).

Under reference opponency the primary colours are a reflection of the peak response of the sensors used. Any display technology designed to accurately reproduce colour for a visual system that employs dual channel opponent coding would require primaries that match those of the sensors. Opponency requires that the primaries be equidistantly placed on the visible spectrum. If the human visual system uses dual opponency to code colour information then four primaries are required and they must be located equidistantly on the visible spectrum. If the green primary for the human visual system is approximately 520 nm and the blue primary is shifted from 430 nm to 460 nm, then dual channel opponency would predict the primaries to be [460 nm, 520 nm, 580 nm, 640 nm], with a spread of 60 nm (identical to that of bird vision). Studies of perception have long indicated that the human visual system functions on the basis of four primaries, with the named colour *yellow* as the fourth primary. If the red primary is shifted to 640 nm then the spectral gap between the green and red primaries is precisely 60 nm, and it would therefore be predicted that the yellow primary be located at 580 nm (see Appendix E for a technical review of using a yellow primary for colour reproduction).

In a break with traditional trichromacy, some recently developed large screen display devices that use liquid crystal display (LCD) technology have been designed to produce four primary colours, with the fourth primary being yellow Dipert [2010]. The light emiting diodes that are commonly used to drive liquid crystal displays typically have peak wavelengths at [455 nm, 520 nm, 635 nm] (see Appendix C). Many high brightness light emitting diode modules designed to produce a reference standard *white* light do so with the use of modules that consist of four separate light emitting diodes; one for each primary colour (blue, green, yellow and red) . The yellow primary is typically set to produce a peak wavelength of 590 nm. It may be seen therefore that the evolutionary development of display technology has led to the use of primaries which, although were initially linked to the sensitivies of the three physical sensors of the human visual system, have moved toward and ultimately closely approached the primaries predicted by dual channel opponency. Particularly significant is the shift from the use of three primaries to the use of four.

The evolution of display technology is therefore consistent with the hypothesis that the human visual system acts as if it employed four colour sensors rather than the three physical sensors available to it. If a virtual sensor response based on the study of

human perception is assumed then the spread between the virtual primaries of human sensors is very similar to that of bird sensors, with the human spectral range simply mapped approximately 75 nm towards the long wavelengths. The difference between bird vision and human vision may therefore be expressed as a simple spectral shift. Assuming that bird sensor values are coded directly, we can predict for example that a bird would perceive a monochromatic light of 445 nm as *green*, 505 nm as *yellow*, and 535 nm as *chartreuse*. This would be a reasonable prediction because both birds and humans have a common ancestor that had a visual system which employed dual channel opponency colour with full complement of four colour sensors. The area of the spectrum to which the visual system is tuned (the visible spectrum) may over time change with the specific needs of the animal, but if the principles of the colour system remain unchanged then dual opponency may be used to accurately predict perception for any animal that employs opponent coding. Moreover, as the opponent code is a relative measure, translation between different visual systems is simply a matter of shifting the spectrum and adjusting for the spectral distance between sensors.

Chapter 10

Conclusion

The most important element in the design of any information system is the nature of the information itself. Natural visual systems that developed colour perception began with sensors capable only of broad luminance measurement and no a-priori knowledge of electromagnetic radiation. It is likely therefore that the origin of colour information is closely linked to the sensors that are the source of all colour information. Natural visual systems measure luminance by comparing the response of sensors pairs. If a light stimuli is monochromatic, it is possible under optimal conditions to use luminance sensor pairs whose spectral response has been modified to calculate the dominant spectral frequency of the stimuli. Extending this to cope with non-optimal stimuli leads to a necessary state of initial processing where sensor values must first be normalized with respect to a known reference and any non-monochromatic elements must be removed. Information is lost by these two processing steps, and therefore any system that aims to represent light stimuli as fully as possible must preserve these parameters as colour information.

Natural design choices that follow from the use of modified luminance sensors therefore lead to an information model that consists of spectral frequency (hue), degree of monochromaticity (saturation) and degree of normalization (lightness). Saturation requires a sensor in addition to a sensor pair used to calculate hue. This requirement as well as the need to cover a broad area of the spectrum naturally leads to the use of dual sensor pairs to calculate hue. If hue is calculated from a dual sensor pair design then it is useful to maintain hue as two separate but linked values (opponent channels). The use of two separate but linked values to represent hue allows for a *no colour* state which enables backward compatibility with a luminance-only representation. It also organizes hue into multiple poles (two poles per channel), thereby preserving opponent information which would otherwise be lost.

Given the sensor limitations of natural visual systems, the simplest colour model therefore employs a dual channel four sensor design which produces four *channels* of information. If four sensors are unavailable, it is possible to drive the two opponent sensor

channels with only three sensors using an emulation layer. This reduces accuracy and adds computational complexity.

The dual channel opponent design presents a universal model for more fully representing light stimuli. Specifically, this is a model suited to early sensor specific processing that places sensor information into a standard format within very limited time constraints. This function is generally referred to as *pre-processing*. The human visual system is known to have a thin layer of neural circuitry that lies directly behind the visual sensor array that carries out a pre-processing function on the information produced by the sensors before it is sent via a dedicated communication link to the visual system. It is left as a subject for further study to test the hypothesis that the function of this layer of neural circuitry is functionally equivalent to the proposed pre-processing functions set out in this thesis.

Appendix A

A Worked Example

Assume (for this example) that luminance is determined on the basis of the following equation:

$$Y = 0.05b + 0.45g + 0.45y + 0.05r \tag{A.1}$$

Early visual processing begins with the raw sensor values (whose range is between 0 and 1) produced by the light stimuli:

$[B = 0.2, G = 0.5, \bar{Y} = 0.4, R = 0.2]$

Natural visual systems use luminance contrast rather than absolute luminance. Luminance contrast is equivalent to luminance at maximal constrast. Assuming the neighbouring sensors values of the opposing luminance contrast pairs are $[0, 0, 0, 0]$ then:

$[0.2, 0.5, 0.4, 0.2]$

$Luminance(Y) = 0.425$

The maximum sensor value is 0.5, which is taken as the colour luminance (*lightness*) value. The minimum sensor value is 0.2, which is taken as the equiluminance (*saturation*) value. This value is subtracted from the sensor values.

$[0, 0.3, 0.2, 0]$

$Luminance = 0.425$
$Saturation = 0.2$
$Lightness = 0.5$

The non-zero sensor values are now normalized, the factor for which is $Lightness - Saturation$. The sensor values are multiplied by the inverse of the normalization factor, which is $\frac{1}{0.3}$.

$[b = 0, g = 1.0, y = 0.667, r = 0]$

$Luminance = 0.425$
$Saturation = 0.2$
$Lightness = 0.5$

Once sensor values are in a standard form, hue is calculated. $hue = <r-g\,,\,b-y>$

$Luminance = 0.425$
$Hue = <-1, -0.667>$
$Saturation = 0.2$
$Lightness = 0.5$

These five values are the final product of the early visual system, and are subsequently sent to the visual system for higher level processing.

The reference value for hue indicates that the dominant element of the percept is expected to be the primary colour green, and the measured value indicates the distance from the primary and the direction. It would be expected that the dominant wavelength of the stimuli is located at ⅔ of the spectral distance between green and yellow. If there are no further higher level corrections or amendments then it would be expected that this would also be the perceptual distance from the primary. If the *green* sensor is located at 520 nm and the *yellow* sensor is located at 580 nm then it would be predicted that the dominant wavelength of the stimuli is 560 nm.

The value for hue represents wavelength relative to the sensors used. If the spectral sensitivities of the sensors are known then the spectral frequency of a stimuli may be determined. Conversely, storing the values of lightness and saturation maintains reversability of the transforms. The wavelength of any light stimulus with a monochromatic element maps to a unique hue value, and the lightness and saturation values allow a mapping from hue to the sensor values produced by the stimulus.

Appendix B

Colour Models

Figure B.1 shows the CIE color-matching functions, $\overline{x}(\lambda)$, $\overline{y}(\lambda)$, and $\overline{z}(\lambda)$, which can be thought of as the spectral sensitivity curves of three linear light detectors that produce the CIE XYZ tristimulus values "X", "Y", and "Z". The tabulated numerical values of these functions are known collectively as the CIE standard observer [Poynton, 2003a, 216].

The tristimulus values for a color with a spectral power distribution $I(\lambda)$ are given in terms of the standard observer by:

$$X = \int_0^\infty I(\lambda)\,\overline{x}(\lambda)\,d\lambda \qquad Y = \int_0^\infty I(\lambda)\,\overline{y}(\lambda)\,d\lambda \qquad Z = \int_0^\infty I(\lambda)\,\overline{z}(\lambda)\,d\lambda \quad \text{(B.1)}$$

Where λ is the wavelength of the equivalent monochromatic light (measured in nanometers).

The spectral distribution of Y was deliberately defined to match the human perceptual response of luminance, and therefore this term is often used to refer to luminance (measured linearly).

If X, Y, Z are expressed as proportions then one of the terms becomes redundant:

$$x = \frac{X}{X+Y+Z} \qquad y = \frac{Y}{X+Y+Z} \qquad z = \frac{Z}{X+Y+Z} = 1 - x - y \quad \text{(B.2)}$$

Given a luminance value Y then $X = \left(\frac{x}{y}\right)Y$ and $Z = \left(\frac{1-x-y}{y}\right)Y$. The terms x and y therefore may be said to contain the colour (*chromaticity*) information. Figure B.2 shows the colour information (*chromaticity*) for monochromatic stimuli in relation to the central white point (ideally equiluminant). The diagram is best understood as an

aberrant circle with its center as the white point with the monochromatic colours as the outer bounds of all possible intermediary colours (shown by Figure B.3). The anomalous non-circular boundary of the colour area may be considered an artefact of how the human visual system in practice departs from the ideal circular colour space. The colour space is nominally circular, but the lowest monochromatic wavelength is connected by a non-spectral straight line to the highest. From Figure B.3 it may be seen that by removing the luminance component the dimension of *lightness* is transformed into a dependant variable, with a maximum lightness at the centre of the diagram and a minimum (its lower limit without luminance is the mid-point) at the outer boundary. The center of the coloured space is also the point of minimum saturation, and therefore the corresponding colour would be perceived as white (it is for this reason that it is called the *white point*, although care must be taken not to conflate this with the perception of luminance). The dimension of lightness is not independent but it varies linearly (but only in the positive dimension) as a ratio of the distance from the boundary to the centre. At the boundary therefore saturation is at maximum and lightness is at the mid-point (between maximum lightness and minimum). Expressed in this way, colour information may be referred to as *chromaticity* or *chroma*. All chromatic information for the XYZ colour system may therefore be expressed using the $[x, y]$ chromaticity coordinate space. Any colour may therefore be specified by its chromaticity and the additional dimension of Y. This is known as the CIE xyY colour space.

CIE XYZ values represent a linear measure of luminance and they are not normalized (the $I(\lambda)$ function's output range is traditionally between 0 and 100). The values are always normalized with respect to a reference light source (normally referred to as the *illuminant*, which typically will approximate an ideal *white*). The optimal illuminant is an equal energy illuminant (*CIE Illuminant E*), whose CIE [x,y] coordinates are [⅓,⅓]. The values X_n, Y_n and Z_n are the XYZ luminance measurements produced by the illuminant. Once normalized, the linear luminance values must be transformed into a perceptually linear mapping, and this is best achieved by an exponent of ⅓ (a *gamma* value of 3).

$$X' = \left(\frac{X}{X_n}\right)^{\frac{1}{3}} \qquad Y' = \left(\frac{Y}{Y_n}\right)^{\frac{1}{3}} \qquad Z' = \left(\frac{Z}{Z_n}\right)^{\frac{1}{3}} \tag{B.3}$$

Once normalized (values between 0 and 1) and translated into a logarithmic mapping of luminance, these values may be put into opponent form. This consists of luminance and two opponent colour values.

$$L' = Y' \qquad a' = X' - Y' \qquad b' = Y' - Z' \tag{B.4}$$

Put into opponent form, the colour space would be known as L'a'b' (with the prime denoting non-linear). The CIE modifies this transform slightly to accommodate conditions of low luminance and scales the parameters in a rather non-uniform way. The

CIE refers to this as the CIE L*a*b* colour space (with the * denoting the non-uniform transform).

Once in Cartesian coordinate form, the Lab colour space can also be expressed in polar coordinate (circular) form. The [a,b] opponent pair can be used to determine chroma and hue:

$$C' = \sqrt{a'^2 + b'^2} \qquad h = tan^{-1}\frac{b'}{a'} \tag{B.5}$$

The chroma value C is equivalent to saturation. Saturation may alternatively be expressed as:

$$s = \frac{C'}{L'} \tag{B.6}$$

The motivation in expressing the XYZ colour space in polar coordinate opponent form is that it allows colours to be specified and manipulated by the *natural* parameters of *hue* and *saturation*. The natural parameters of colour are defined by the tradition of artists and technicians who work with colour to produce images that the human visual system perceives as analogous to stimulation by the light reflected in the natural environment. This requires an ability to systematically manipulate colour. The purpose of carrying out the computationally complex transforms on the original XYZ sensor values is therefore to place them into a standard form that make colour calculations possible. Whether the CIE colour system models the colour system used by the human visual system with sufficient accuracy and consistency to enable accurate and colour computation is a matter of ongoing research.

Figure B.4 shows an approximation of an ideal colour space and B.6 shows the expected mapping between hue and monochromatic spectral frequency (using the conventional $atan2$ function). It should be noted that the ideal colour circle is perfectly linear throughout the spectral range. There is therefore in this case a simple transform that maps hue to spectral frequency. Figure B.7 shows how the hue of the CIE Lab colour system maps to spectral frequency. It may be observed from B.7 that while there is an approximately linear relation between hue and spectral frequency for two narrow regions within the range (approximately blue to green and green to yellow), information is lost in broad areas of the periphery of the spectral range where the mapping is approximately flat. Information cannot be recovered if the mapping is flat irrespective of the way the information is coded.

Under the CIE colour model, hue is defined as "the degree to which a stimulus can be described as similar to or different from stimuli that are described as red, green, blue, and yellow (the unique hues)" Fairchild [2004]; Sun et al. [2012]. This defines hue in terms of the subjective perception provided by the human visual system. It does not provide a useful definition of how measurements of a physical phenomena

are codified. A visual system whose purpose is to provide well defined information about the light stimuli being measured might simply define hue as spectral frequency, in which case the mapping from hue to spectral frequency would be expected to be linear throughout the spectral range being coded for. The CIE XYZ-Lab colour space is flat for significant areas of the visible spectrum (see Figure B.7), which makes an accurate mapping to spectral frequency impossible.

The ultimate aim of the CIE XYZ colour space is to enable perceptually accurate colour reproduction. Figure B.1 and Equation B.1 define how a light stimulus is measured. Figure 1.1 shows how a light stimulus may be created that accurately maps the CIE XYZ colour space. Unfortunately this mapping shows large areas of the spectrum where *negative light* is required. While negative light is produced in colour matching by switching one of the primaries to the light stimulus that is to be matched (thereby *subtracting* it) this is not possible when reproducing light stimuli (outside of the colour matching environment). The XYZ primaries are therefore sometimes referred to as *imaginary primaries*. A well established colour space for the practical colour reproduction is sRGB, which uses the primaries defined by ITU-R Recommendation BT.709. The sRGB colour space maps directly to the luminance levels produced by the primaries. Equation B.7 defines the transform from the XYZ colour space into the sRGB colour space. Figure B.5 shows the relative luminance levels of the sRGB primaries in linear form for all monochromatic colours. It may be seen from the colour matching values of the XYZ colour space that all monochromatic colours within the range of approximately 410-680 nm (that is, most of the spectral range) cannot be directly reproduced with sRGB (as sRGB does not support negative luminance values). As the human visual system is based on colour opponency, any set of primaries with an approximate equidistant distribution should be able to match the primaries and most of the intermediary monochromatic colours, as shown by the original CIE colour matching functions (Figure 1.1).

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 3.240479 & -1.537150 & -0.498535 \\ -0.969256 & 1.875992 & 0.041556 \\ 0.055648 & -0.204043 & 1.057311 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} \tag{B.7}$$

Figure B.1: The CIE 1931 XYZ color matching functions. Source: CIE [1932].

Figure B.2: CIE 1931 2° $[x, y]$ chromaticity diagram for monochromatic stimuli, assuming an optimal equiluminant illuminant (CIE*illuminant E*). If the semi-enclosed region is taken as circular, then the white point at $[0.333, 0.333]$ is at the centre of that circle.
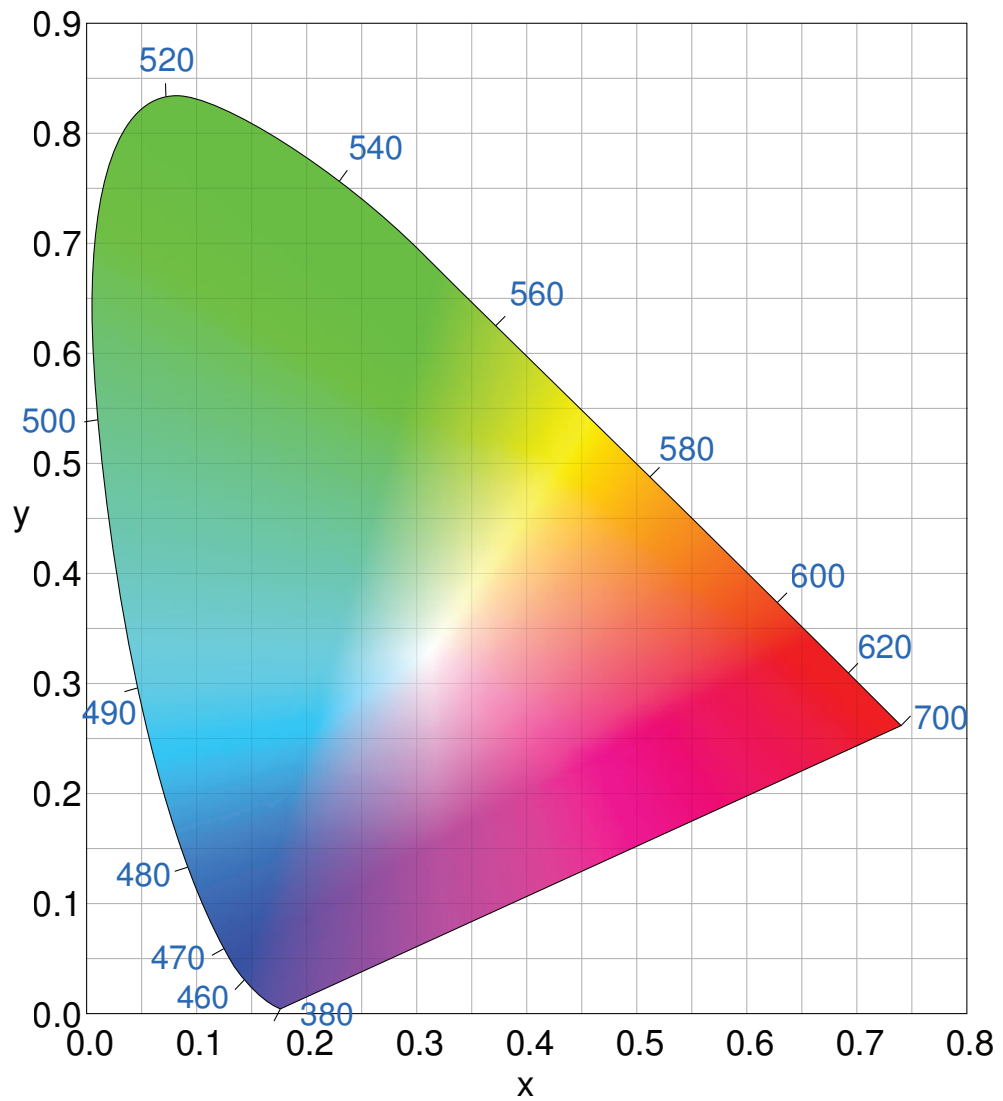
Figure B.3: CIE 1931 2° $[x, y]$ chromaticity diagram. All possible colours lie within the bounded region. Colour produced by monochromatic wavelengths (fully saturated colours) map onto the curved boundary, with common wavelengths indicated in nanometres. Source: CIE [1932] and Wikimedia.

Figure B.4: An approximation of an ideal circular mapping of chroma. The white point is at the centre, with the four primaries of red, yellow, green and blue at the cardinal points of the bounding circle. All intermediary fully saturated colours track the bounding circle.

Figure B.5: XYX to sRGB, with primaries defined by ITU-R Rec. 709.
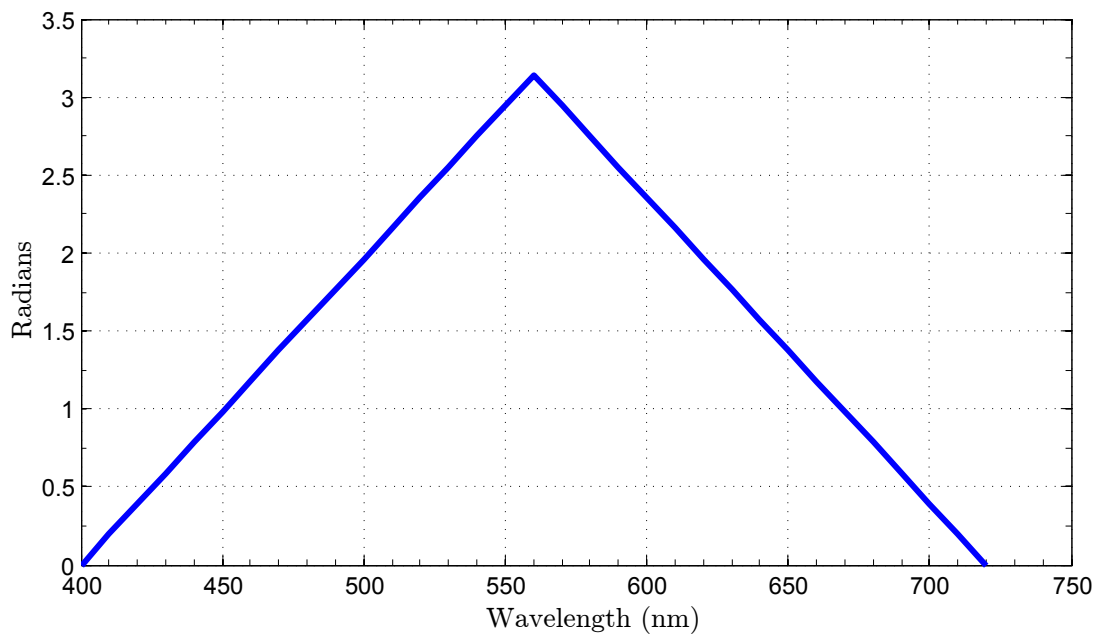
Figure B.6: Mapping of polar angle (in radians) to wavelength for an ideal circular colour space.

Figure B.7: Mapping of polar angle (in radians) to wavelength for CIE L'a'b'.

Appendix C

The Visible Spectrum

Figure C.1 shows the spectral power distribution of the phosphors of a colour display based on the cathode ray tube technology. The *blue* phosphor has a flat peak response of approximately 40 nm, whereas the peak response (production of light) of the green phosphor is approximately half this spectral distance. Both phosphors have an approximate *normal* spectral distribution. The *red* phosphor on the other hand has a very divided response. It produces spectral peaks within three narrow spectral ranges [590 nm, 625 nm and 700 nm]. If a single primary had to stimulate two or more sensors each with a different spectral response then a phosphor design with this type of disparate response would be selected.

Figure C.2 shows the spectral response of a common four colour LED module, with the primaries at [455 nm, 518 nm, 598 nm, 635 nm]. All the colours except green show a narrowed spectral distribution of approximately 50 nm, with green having approximately double the spectral range. The human visual system measures luminance primarily from the *green* sensor (approximately $\frac{2}{3}$g + $\frac{1}{3}$r) with only a nominal contribution from the *blue* sensor (which is entirely absent in the central regions of the fovea). The green primary should therefore have the broadest spectral distribution.

Figure C.3 and C.4 show the spectral power distributions of common artificial light sources which are designed to approximate the appearance produced by the light source that illuminates the human natural environment. The aim of high quality artificial light sources is generally to approximate the appearance of *white* (hence the term *white light*). The very narrow range of the wavelengths produced by a high quality tri-phosphor fluorescent lamp (see Figure C.3) indicates a design which seeks to maximise luminous efficiency by producing only wavelengths which usefully contribute to producing the desired appearance. This may be contrasted with Figure C.4 which shows an older less efficient design. It may be seen that the halophosphate fluorescent lamp produces a broad spectral distribution between 550 nm and 630 nm. The bulk of the wavelengths produced are therefore in the *yellow* range and this results in the typically yellow appearance produced by these lamps.

Figure C.1: Spectral response of the blue, green and red phosphors in a common cathode ray tube. Blue: 445-485 nm, Green: 525-540 nm, Red: 595 nm, (618 nm and 625-630) nm, 705 nm. Source: Wikimedia .

Figure C.2: Spectral response of the Perkin-Elmer four colour LED (part no. E001704). Source: PerkinElmer, Inc. [2009].
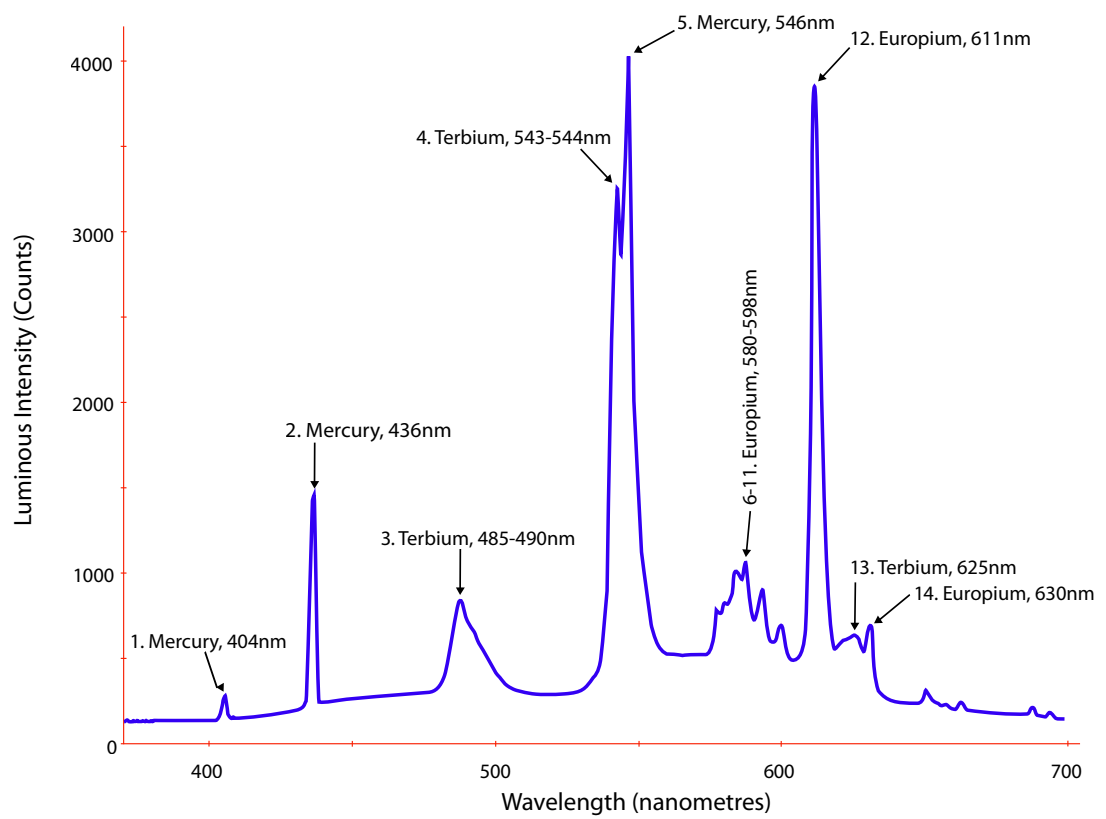
Figure C.3: Spectral power distribution of common tri-phosphor fluorescent lamp with an apparent *colour temperature* approaching that of the diurnal light source of the natural human environment. Source: Wikimedia .
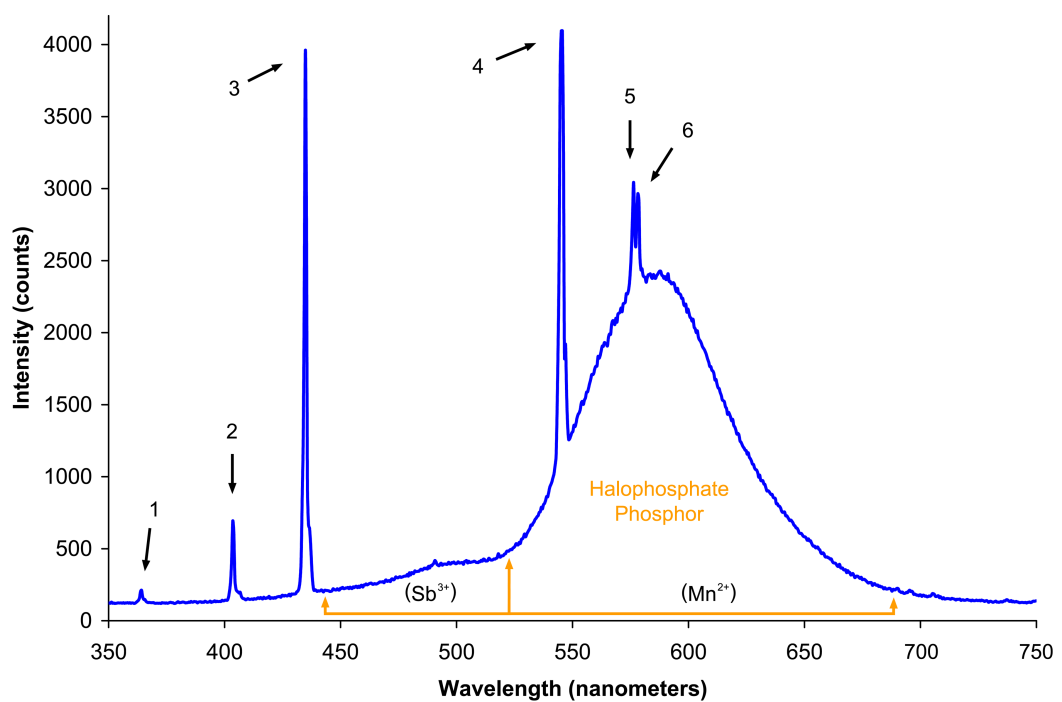
Figure C.4: Spectral power distribution of halophosphate fluorescent lamps, an older technology which produces a lower apparent colour *temperature* with a typically *yellow* appearance. Spectral peaks at: (2) 404 nm, (3) 436 nm, (4) 546 nm and (5-6) 577-579 nm. Source: Wikimedia .

Appendix D

Colour Wheels

The colour circle is said to have originated from Newton (see Figure D.1) Lowengard [2006]. Newton took the linear range of colors created by passing light through a prism and transformed it into a segmented circle, where the size of each segment differs according to its associated spectral frequency and width in the spectrum. Newton was an adherent of allowing only *spectral colours*, but his colour circle led to a gap between red and violet which he filled with purple, which is not, strictly speaking, a spectral colour.

Goethe's symmetric colour wheel (see Figure D.5) Ribe and Steinle [2002] is the first known reference to colour opponency:

> In order at once to see what colour will be evoked by this contrast, the chromatic circle * may be referred to. The colours are here arranged in a general way according to the natural order ... for the colours diametrically opposed to each other in this diagram are those which reciprocally evoke each other in the eye. Thus, yellow demands purple; orange, blue; red, green; and *vice versâ*: thus again all intermediate gradations reciprocally evoke each other; the simpler colour demanding the compound, and *vice versâ* [Goethe, 1840, p. 50].

Goethe also investigated the psychological effects of colour, and as a result included aesthetic qualities in his colour wheel – associating red to the beautiful, orange to the noble, yellow to the good, green to the useful, blue to the mean, and violet to the unnecessary.

Once the idea of arranging colours by systematic transition between a set of primaries was established, a wide variety of colour wheels were proposed. Figure D.2 shows an early 18th century artists colour circle, while Figures D.3 and D.4 show more formal attempts to define the colour circle. One of the more practical colour circles is the RYB colour system which was an early standard for colour printing, shown by Figure

D.6. A colour wheel that still has some contemporary resonance is the Munsell colour system, shown by Figure D.7.

True colour opponency has its origin in the early 20th century with Hering, who showed that the intermediary colour yellow is not perceived by the human visual system as an equal mixture of red and green as predicted under the RGB model. Hering showed that yellow is perceived as a primary colour by his studies of perception which indicate that yellow is inherently linked to the blue primary, and that in the same way the red primary is linked to the green primary. The primary colours are therefore a set of two linked pairs of primary colours. A primary pair forms a linear axis along which the intermediary colours are arranged. Taken together they form a two dimensional space, which expressed in polar form gives a colour circle (see Figure D.8).

Once transmissive colour display technologies became available they rapidly became the dominant force in static colour image reproduction. This led to the RGB colour model, and Figure D.9 shows an RGB circular transform (HSL/HSV). The transmission of large quantities of colour image information, particularly with respect to moving colour images, led to colour models that were more closely linked to the human visual system, allowing only information that is needed by the human visual system to be stored. This requires colour information to be represented by two opponent channels. Figure D.10 shows a circular opponent colour model. Implicit with a two opponent channel model is the requirement for four primaries, with *yellow* as the fourth primary. A circular RGB colour space (as shown in Figure D.9 has three primary colours and treats yellow as a transition hue (like cyan and magenta). By contrast, with an opponent colour space yellow is a primary and the transition hues are chartreuse, orange, cyan and magenta.

Figure D.1: Newton's colour wheel (1704). Clockwise from A, the colours are: Blue, Indigo, Violet, Red, Orange, Yellow and Green. Source: Lowengard [2006].

Figure D.2: An early 18th century colour circle from an artist's manual. The first circle contains seven colours: violet, blue, green, yellow, orange, scarlet and crimson. The second circle adds golden yellow, red, purple, sea green and yellow-green for a total of twelve colours. These color circles are the earliest published examples of a practical application of Newton's color circle. Source: Lowengard [2006].

Figure D.3: Schiffermueller's colour wheel (1772) consists of three primaries (blue, yellow and *carmine*), three secondaries formed by their combination (*sea green*, *orange yellow*, and *violet red*), and six tertiary colours (green, *olive green*, red, *fire red*, *violet blue* and *fire blue*). He named *carmine* as the primary red and placed scarlet (*fire-red*) and true red in secondary and tertiary positions, respectively. Source: Lowengard [2006].

Figure D.4: Harris' theory of colour is based on the principle that all colors are formed from red, blue, and yellow, but his colour circle (1776) is based on a prismatic sequence of six colors rather than on the artists' primaries. He resolves the numerical discrepancy (three primaries, six colours) by designating two types of principal color. Red, yellow, and blue are the grand principals; and orange, green, and purple are compound principals. Source: Lowengard [2006].

Figure D.5: Goethe's colour circle (1810). Inner Ring: red "schön", orange "edel", yellow "gut", green "nützlich", blue "gemein", violet "unnötig". Outer Ring: red-orange "Vernunft", yellow-green "Verstand", green-blue "Sinnlichkeit", violet-red "Phantasie". Source: Goethe [1840].

Figure D.6: The RYB colour star. In the 18[th] century, the RYB colour model became the standard for colour printing. As a result it became the foundation of early theories of color vision, and established the primary colors red, yellow and blue as the fundamental sensory qualities that are blended in the perception of all physical colors. Source: Wikimedia .

Figure D.7: The Munsell colour circle (1905) is divided into five principal hues: Red, Yellow, Green, Blue, and Purple, along with 5 intermediate hues halfway between adjacent principal hues (e.g., YR – Yellow-Red). These 10 hues are each given the number 10. These hues are then broken into a further 10 sub-principals, with each sub-principal given the number 5. This is repeated recursively until the desired hue precision is achieved (typically 40 hues). Adapted from Wikimedia .

Figure D.8: The colour wheel by Hering (1920). The primary colours are red-green and yellow-blue, placed on the colour wheel as opponent poles of two perpendicular axes. Intermediate colours are formed by additive mixing of adjacent primaries. Source: Douma and WebExhibits [2008].

Figure D.9: The HSL/HSV colour wheel (1978) is a circular transform of the RGB colour space. The primary colours are red, green and blue (RGB), with the transition colours as yellow (red-green), magenta(red-blue) and cyan(blue-green).

Figure D.10: The ideal four primary colour model, arranged into two opponent pairs (red-green, blue-yellow). The primary colours are red, yellow green and blue, and the transition colours are orange, magenta, cyan, and chartreuse.

Appendix E

RGB, YUV, and why Yellow can be considered harmful

YUV, YIQ, YC_BC_R and YP_bP_r are a family of opponent colour spaces used to encode RGB information. They are employed when the amount of colour information needs to be minimized. This is achieved by separating luminance (Y) from colour (C) information, and by placing the latter into opponent form. Following on from how the human visual system is known to organize colour information, two opponent colour channels are defined: a red-green chroma channel (C_R) and a blue-yellow chroma channel (C_B). All the colour opponent spaces use the same transform (given below by Equation E.1) but may differ in respect of the luminance coefficients used. These coefficients have evolved over time to more generally reflect the accurate colour reproduction with respect to the human visual system. For convenience, these colour spaces will collectively be referred to as YUV, and it is assumed that the coding is logarithmic rather than linear (that is, for example, Y should be read as Y' – sometimes referred to as *luma* to distinguish it from linearly coded luminance).

The general equation for transforming RGB into colour opponent form is as follows:

$$Y = k_R R + k_G G + k_B B$$

$$C_B = \frac{1}{2} \cdot \frac{B - Y}{1 - k_B} \tag{E.1}$$

$$C_R = \frac{1}{2} \cdot \frac{R - Y}{1 - k_R}$$

The constant k defines the respective luma coefficients used for the RGB primaries. YUV, as defined for standard-definition television (ITU-R Recommendation BT.601),

uses the following constants for determining luminance (k_R, k_G and k_B):

$$k_R = 0.299 \qquad k_G = 1 - (k_R + k_B) \qquad k_B = 0.114$$

This leads to the following transform:

$$\begin{bmatrix} Y \\ C_B \\ C_R \end{bmatrix} = \begin{bmatrix} 0.299 & 0.587 & 0.114 \\ -0.168736 & -0.331264 & 0.5 \\ 0.5 & -0.418688 & -0.081312 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} \tag{E.2}$$

ITU-R Recommendation BT.709 sets out newer standards for the high definition television format and in most respects supercedes Recommendation 601 (although Recommendation 601 remains in use, with for example the commonly used JPEG image compression standard). The luma coefficients set out by Recommendation 709 are:

$$k_R = 0.2126 \qquad k_B = 0.0722$$

YUV is a tranform of RGB, and as such it is used only as an intermediary medium to transmit or process image information. It is translated into RGB form for transmissive display purposes. This means that consistency is more important than the precise luma constants. If the luma constants are set to:

$$k_R = {}^1\!/_3 \qquad k_B = 0$$

then this allows the transform to be simplified ...

$$Y = \tfrac{1}{3}R + \tfrac{2}{3}G \qquad C_B = \tfrac{1}{2}(B - Y) \qquad C_R = \tfrac{3}{4}(R - Y) \tag{E.3}$$

Given Y in simplified form:

$$C_R = \tfrac{3}{4}(R - (\tfrac{1}{3}R + \tfrac{2}{3}G))$$

therefore

$$C_R = \tfrac{1}{2}(R - G) \tag{E.4}$$

The inverse is:

$$R = Y + \tfrac{4}{3}C_R \qquad G = Y - \tfrac{2}{3}C_R \qquad B = Y + 2C_B \tag{E.5}$$

This simplification relies on the coefficient for blue to be 0. Perceptual studies of the human visual system show that blue is a small but significant component of luma; approximately 10%. It is however also known that the human visual system is able to

function without blue sensors in the central area of the fovea, which produces the most detailed visual information. It is therefore reasonable to initially determine luminance and calculate the chroma values without a blue component (Y) and then to add the correct ratio to luma from the derived chroma blue value (Y^*).

$$Y^* = Y + \tfrac{1}{5}\left|C_B\right|$$

The inverse is:

$$Y = Y^* - \tfrac{1}{5}\left|C_B\right|$$

The individual red, green and blue primaries produce the following YUV values (without the blue component added to luma):

	R	G	B	Y	C_B	C_R
Red	1	0	0	$\tfrac{1}{3}$	$-\tfrac{1}{6}$	$\tfrac{1}{2}$
Green	0	1	0	$\tfrac{2}{3}$	$-\tfrac{1}{3}$	$-\tfrac{1}{2}$
Blue	0	0	1	0	$\tfrac{1}{2}$	0
$Yellow$	1	1	0	1	$-\tfrac{1}{2}$	0
cyan	0	1	1	$\tfrac{2}{3}$	$\tfrac{1}{6}$	$-\tfrac{1}{2}$
magenta	1	0	1	$\tfrac{1}{3}$	$\tfrac{1}{3}$	$\tfrac{1}{2}$

(E.6)

Because YUV is defined as an RGB transform, the opponent channels should be *clean*; that is, the red or green primary should have no blue-yellow opponent value and a blue primary should have no red-green opponent value. Two opponent channels imply a fourth primary, however, RGB is limited to three. RGB is able to emulate a yellow primary by treating the red and green primary jointly as a single primary (R=G). The *yellow* primary should also produce no red-green opponent value. However, as can be seen from Table E.6, the red primary produces an aberrant chroma blue value of -1/6 and the green primary produces an aberrant chroma blue value of -1/3. It may also be seen from Table E.6 that the opponent values for the intermediary transition colours of magenta (red-blue) and cyan (blue-green) are incorrect. Intermediary colours are produced by pairs of RGB primaries and should translate into opponent values with both channels at equality of magnitude ($\left|C_B\right| = \left|C_R\right|$). These errors have their origin in the lack of a yellow primary in the RGB colour system, and they reflect a failure to evaluate a conditional inherent to emulating a four primary colour system.

The conditional can be expressed as a yellow correction coefficient (c_{ycc}) which evaluates on the basis of whether $R > G$.

If $R < G$ then

$$c_{ycc} = \left. R \middle/ G \right.$$

else if $R > G$ then

$$c_{ycc} = {}^{G}\!/_{R}$$

and when $R = G$ then $c_{ycc} = 1$.

A simplified luma allows an approximation of the yellow primary with the luma value to the degree that $R = G$. Chroma blue with yellow correction may be defined as:

$$C_B = \tfrac{1}{2}\Big(B - c_{ycc}Y\Big) \tag{E.7}$$

The inverse is:

$$B = c_{ycc}Y + 2C_B \tag{E.8}$$

With R, G and the yellow correction coefficient defined above, table E.9 shows the corrected YUV values for the respective RGB primaries:

	R	G	B	Y^*	C_B	C_R
Red	1	0	0	$\tfrac{1}{3}$	0	$\tfrac{1}{2}$
$Green$	0	1	0	$\tfrac{2}{3}$	0	$-\tfrac{1}{2}$
$Blue$	0	0	1	0.1	$\tfrac{1}{2}$	0
$Yellow$	1	1	0	1.1	$-\tfrac{1}{2}$	0
$cyan$	0	1	1	$\tfrac{2}{3}{+}0.1$	$\tfrac{1}{2}$	$-\tfrac{1}{2}$
$magenta$	1	0	1	$\tfrac{1}{3}{+}0.1$	$\tfrac{1}{2}$	$\tfrac{1}{2}$

$$\tag{E.9}$$

For convenience, luma values for YUV are allowed to go 10% above unity. This may be dealt with by reducing the range by $^{1}\!/_{1.1}$ or practically by simply clipping the luma value, and inversely assuming that the value has been clipped.

The YUV colour system allows intermediary positions for the yellow primary (yellow-red and green-yellow) that cannot be translated back into RGB (that are *out of gamut*). Table E.10 shows the closest RGB approximation to a true yellow-red intermediary colour and a green-yellow intermediary colour.

	R	G	B		Y^*	C_B	C_R
$Yellow$	1	1	0		1.1	$-\tfrac{1}{2}$	0
$yellow\text{-}red$	1	$\tfrac{1}{2}$	0	$\rightarrow$	$\tfrac{2}{3}{+}0.1$	$-\tfrac{1}{6}$	$\tfrac{1}{4}$
$green\text{-}yellow$	$\tfrac{1}{2}$	1	0	$\rightarrow$	$\tfrac{5}{6}{+}0.1$	$-\tfrac{5}{24}$	$-\tfrac{1}{4}$

$$\tag{E.10}$$

Table E.11 shows how true yellow-red and green-yellow in the YUV colour space would

translate into RGB. The values are out of gamut.

$$
\begin{array}{lcccccc}
 & R & G & B & Y^* & C_B & C_R \\
Yellow & 1 & 1 & 0 & 1.1 & -\frac{1}{2} & 0 \\
yellow\text{-}red & 1\frac{2}{3} & \frac{2}{3}+\frac{1}{10} & 1\frac{2}{5} \;\leftarrow & 1.1 & \frac{1}{2} & \frac{1}{2} \\
green\text{-}yellow & \frac{1}{3} & 1\frac{1}{3} & 1\frac{1}{4} \;\leftarrow & 1.1 & \frac{1}{2} & -\frac{1}{2}
\end{array}
\tag{E.11}
$$

Table E.12 shows the translation of the closest RGB approximation to yellow-red and green-yellow into YUV.

$$
\begin{array}{lcccccc}
 & R & G & B & Y^* & C_B & C_R \\
Yellow & 1 & 1 & 0 & 1.1 & -\frac{1}{2} & 0 \\
orange & 1 & \frac{1}{2} & 0 \;\leftrightarrow & 0.7 & -\frac{1}{6} & \frac{1}{4} \\
chartreuse & \frac{1}{2} & 1 & 0 \;\leftrightarrow & \frac{7}{8} & -\frac{5}{24} & -\frac{1}{4}
\end{array}
\tag{E.12}
$$

Index

References

Anderson, K. F. (1995). A conversion of wheatstone bridge to current-loop signal conditioning for strain gages. Technical Memorandum 104309, NASA. 34

Bangert, T. (2008). Trianglevision: A toy visual system. In Kurkov, V., Neruda, R., and Koutnk, J., editors, *Artificial Neural Networks - ICANN 2008*, volume 5163 of *Lecture Notes in Computer Science (LNCS)*, pages 937–950. Springer, Berlin/Heidelberg. 6, 9

CIE, editor (1932). *Commission internationale de l'Eclairage proceedings, 1931*, Cambridge UK. Cambridge University Press. vi, 2, 53, 55

CIE (1998). The cie 1997 interim colour appearance model (simple version). Technical report, CIE, Vienna, Austria.

Conway, B. R. (2003). Colour vision: A clue to hue in v2. *Current Biology*, 13(8):R308–R310. 4

Conway, B. R. (2009). Color vision, cones, and color-coding in the cortex. *The Neuroscientist*, 15(3):274–290. 9, 12

Dacey, D. M. (1996). Circuitry for color coding in the primate retina. *Proceedings of the National Academy of Sciences*, 93(2):582–588.

Dacey, D. M. (2000). Parallel pathways for spectral coding in primate retina. *Annual Review of Neuroscience*, 23(1):743–775. 41

Dacey, D. M. and Packer, O. S. (2003). Colour coding in the primate retina: diverse cell types and cone-specific circuitry. *Current Opinion in Neurobiology*, 13(4):421–427.

Dartnall, H. J. A., Bowmaker, J. K., and Mollon, J. D. (1983). Human visual pigments: Microspectrophotometric results from the eyes of seven persons. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 220(1218):115–130. 29, 41

DeMarco, P., Pokorny, J., and Smith, V. C. (1992). Full-spectrum cone sensitivity functions for x-chromosome-linked anomalous trichromats. *J. Opt. Soc. Am. A*, 9(9):1465–1476. 30, 41

Demarsh, L. E. and Giorgianni, E. J. (1989). Color science for imaging systems. *Physics Today*, 42(9):44–52. 1, 3, 4

DeValois, R. L. and DeValois, K. K. (1993). A multi-stage color model. *Vision Research*, 33(8):1053–1065. 3

Dipert, B. (2010). Display-technology advancements: Change is the only constant. *EDN*, pages 24–33. 43

Douma, M. and WebExhibits (2008). Cause of color / curator, michael douma. Retrieved August 13, 2012. 73

Fairchild, M. D. (2004). Color Appearance Models: CIECAM02 and Beyond. In *The 12th Color Imaging Conference*. IS&T. 51

Feisner, E. A. (2006). *Colour*. Laurence King, London, UK.

Goethe, J. W. v. (1840). *Theory of Colours*. John Murray, London, England. Translated by Charles Lock Eastlake. 64, 70

Goldsmith, T. and 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(10):933–951. 8, 41

Goldsmith, T. H. (2006). What Birds See. *Scientific American*, 295(1):68–75. 8, 10, 20, 22, 30, 41

Greenwood, V. (2012). The Humans With Super Human Vision. Published online, Retrieved August 13, 2012. 1, 39

Hart, N., Partridge, J., Bennett, A., and Cuthill, I. (2000). Visual pigments, cone oil droplets and ocular media in four species of estrildid finch. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, 186(7):681–694. 41

Hunt, R. (1995). *The Reproduction of Colour*. Fountain Press, Kingston-upon-Thames, England. 1, 42

Hurvich, L. M. (1997). Chromatic and achromatic response functions. In Byrne, A. and Hilbert, D. R., editors, *Readings on Color: The Science of Color*, pages 67–91. MIT Press. 34

Hurvich, L. M. and Jameson, D. (1957). An opponent-process theory of color vision. *Psychological Review*, 64(6, Part 1):384–404. 3

ITU (2002). ITU-R BT.709-5: Parameter values for the HDTV standards for production and international programme exchange. Technical report, ITU, Geneva, Switzerland.

ITU (2011). ITU-R BT.601-7: Studio encoding parameters of digital television for standard 4:3 and wide screen 16:9 aspect ratios. Technical report, ITU, Geneva, Switzerland.

Jordan, G. and Mollon, J. (1992). Do tetrachromatic women exist? *Invest. Ophthal. & Vis. Sci.*, 33:754. 1

Lennie, P. (2000). Color vision: Putting it together. *Current Biology*, 10(16):R589–R591. 1, 3, 8

Lennie, P. and D'Zmura, M. (1988). Mechanisms of color vision. *Critical Reviews in Neurobiology*, 3(4):333–400. 17

Livingstone, M. and Hubel, D. (1988). Segregation of form, color, movement, and depth: anatomy, physiology, and perception. *Science*, 240(4853):740–749. 39

Lotto, R. B. and Purves, D. (2000). An empirical explanation of color contrast. *Proceedings of the National Academy of Sciences*, 97(23):12834–12839. 12

Lowengard, S. (2006). *The Creation of Color in Eighteenth-Century Europe*. Gutenberg⟨e⟩, Columbia University Press. 64, 66, 67, 68, 69

Mancuso, K., Mauck, M. C., Kuchenbecker, J. A., Neitz, M., and Neitz, J. (2010). A multi-stage color model revisited: Implications for a gene therapy cure for red-green colorblindness. In Anderson, R. E., Hollyfield, J. G., LaVail, M. M., and Mandal, M. N. A., editors, *Retinal Degenerative Diseases*, volume 664 of *Advances in Experimental Medicine and Biology*, pages 631–638. Springer, New York. 3, 30

McCamy, C. S. (1998). On the number of discernible colors. *Color Research & Application*, 23(5):337–337. 39

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. 9

Moroney, N., Fairchild, M. D., Hunt, R. W. G., Li, C., Luo, M. R., and Newman, T. (2002). The CIECAM02 color appearance model. In *The 10th Color Imaging Conference: Color Science and Engineering Systems, Technologies, Applications*, volume 10, pages 23–27, Scottsdale, Arizona. IS&T.

Nathans, J. (1989). The Genes for Color Vision. *Scientific American*, 260(1):42–49. 30

Neitz, J., Carroll, J., and Neitz, M. (2001). Color vision: Almost reason enough for having eyes. *Opt. Photon. News*, 12(1):26–33. 30

Neitz, J., Geist, T., and Jacobs, G. H. (1989). Color vision in the dog. *Visual Neuroscience*, 3(2):119–125. 10, 30

Neitz, J. and Neitz, M. (2008). Colour vision: The wonder of hue. *Current Biology*, 18(16):R700–R702. 4, 7, 26

PerkinElmer, Inc. (2009). ACULED VHL Product Note. Retrieved August 13, 2012. 61

Poynton, C. (2003a). *Digital Video and HDTV Algorithms and Interfaces*. Morgan Kaufmann, San Francisco, CA. 39, 49

Poynton, C. A. (1993). "Gamma"and Its Disguises: The Nonlinear Mappings of Intensity in Perception, CRTs, Film, and Video. *SMPTE Journal*, 102(12):1099–1108. 35, 36, 39

Poynton, C. A. (2003b). YUV and luminance considered harmful. In *Digital Video and HDTV Algorithms and Interfaces*, pages 595–600. Morgan Kaufmann, San Francisco, CA. 35

Ribe, N. and Steinle, F. (2002). Exploratory experimentation: Goethe, Land, and Color Theory. *Physics Today*, 55(7):43–49. 64

Rieke, F. and Baylor, D. A. (1998). Single-photon detection by rod cells of the retina. *Rev. Mod. Phys.*, 70:1027–1036. 35

Shepard, R. N. (1997). The perceptual organization of colors: An adaptation to regularities of the terrestrial world? In Byrne, A. and Hilbert, D. R., editors, *Readings on Color: The Science of Color*, pages 311–356. MIT Press.

Shepard, R. N. and Cooper, L. A. (1992). Representation of colors in the blind, color-blind, and normally sighted. *Psychological Science*, 3(2):97–104. 19

Solomon, S. G. and Lennie, P. (2007). The machinery of colour vision. *Nat Rev Neurosci*, 8(4):276–286. 4, 8, 9, 30

Sun, X., Chang, Y., Ye, Y., Ma, Z., Liang, Y., Li, T., Jiang, N., Xing, W., and Luo, L. (2012). The effect of dietary pigments on the coloration of japanese ornamental carp (koi, cyprinus carpio l.). *Aquaculture*, 342343(0):62–68. 51

Thornton, W. A. (1999). Spectral sensitivities of the normal human visual system, color-matching functions and their principles, and how and why the two sets should coincide. *Color Research & Application*, 24(2):139–156. 3

Valberg, A. (2001). Unique hues: an old problem for a new generation. *Vision Research*, 41(13):1645–1657.

Vorobyev, M. (2006). Evolution of colour vision: The story of lost visual pigments. *Perception*, 35:168. 8, 30