## 8

# THE SCIENCE OF COLOUR AND COLOUR VISION 

Alex Byrne and David R. Hilbert

## 1 Colour vision and colour science

Colour science concerns the process of colour vision and those features of the environment that affect the colours that we see and how we see them. Colour vision has been studied systematically from a variety of points of view since the nineteenth century. The science we discuss below draws on optics, psychology, neuroscience, neurology, ophthalmology, and biology. And, although the relevant basic facts of optics and physiology and their contribution to colour vision have been known for at least a century and half, there are still many aspects of colour vision-including some quite fundamental ones-that are poorly understood. In what follows we will provide an overview of what is known and indicate matters of current controversy. We will concentrate on giving the background necessary to understand those parts of colour science that are potentially relevant to philosophical work on colour. Our account is necessarily quite sketchy and we won't be able to do more than provide a starting point for those interested in the topic. ${ }^{1}$

## 2 The optical process

The process of vision typically begins with a source of light that illuminates the objects in a scene. The light is reflected from the surfaces of objects and some of it enters the eye where the cornea and lens combine to focus the light and produce an image of the scene on the retina. This textbook scenario oversimplifies in numerous respects but it captures two central truths about vision. First, it is light as modified by the surfaces and objects in the environment that enables vision, including colour vision. Differences between objects that don't affect light (or aren't correlated with differences that affect light) are not visible. Accordingly, it is the variable effects of different parts of the scene on the light falling on them that enables us to see objects and their colours. Second, a crucial part of the visual process is the formation of an image in the back of the eye. The retinal image is important because it separates the light coming from the different parts of the scene, enabling spatial vision and with it the ability to visually attribute different colours to different locations within the scene. Although we will temporarily set aside the fact that the immediate stimulus for vision is an image of the scene before the eyes, it will re-emerge when we turn to the important topic of the interaction between spatial vision and colour vision.

### 2.1 Light

Light is a form of electromagnetic radiation, and so can be described in both wave and particle terms. The particles of light, photons, are usefully characterized in terms of their energy (the usual unit is the electron-volt ( eV ), $1.6 \times 10^{-19}$ joules) while the waves associated with the photon are usefully characterized by their wavelength (the usual unit is the nanometer ( nm ), $10^{-9}$ metres). These are not independent characterizations: specifically, the energy of a photon is inversely proportional to its wavelength. The intensity or power of a light is the amount of energy it delivers per unit time. Most light sources emit light at a variety of wavelengths so a complete characterization of a light in these terms requires describing how its power is distributed across wavelengths. The spectral power distribution (SPD) of a light specifies the proportion of the total power of that light that is carried by the photons at each wavelength. For many purposes in colour science, overall intensity is held fixed and it is the varying SPD that is the explanatory variable.

Only a very small segment of the total electromagnetic spectrum is relevant to most questions in colour science because the receptors in the eye only respond directly to a narrow range of wavelengths. The precise boundaries are somewhat arbitrary but the visible spectrum runs roughly from $400 \mathrm{~nm}(3.1 \mathrm{eV})$ at the violet end to $700 \mathrm{~nm}(1.8 \mathrm{eV})$ at the red end. The range of intensities that are relevant is much larger-the ratio of the intensities of the illumination provided by direct summer sunlight to that available on a moonless night is about 10 billion to one. Normal indoor lighting typically lies somewhere near the centre of this range.

The light sources that initiate the process of vision can be described in terms of two kinds of characteristics: spatial and spectral. First, light sources can be divided into those that are of significant spatial extent, like the sky on an overcast day or a bank of fluorescent tubes behind a diffusing panel, and those that approximate point sources, like the sun or a street lamp. An extended source can provide much more uniform illumination across the scene while a point source illuminates objects in a way that depends much more strongly on their position and orientation with respect to the source and the other objects in the scene. Second, as just noted, the light emitted by a source can be characterized in terms of its overall intensity and spectral power distribution. The SPD of light sources is critical to understanding how the process of colour vision works.

## 3 Colour in the environment

### 3.1 Objects

When light falls on an object some proportion of the light at each wavelength is reflected, some proportion is absorbed and-for transparent and translucent objects-some proportion is transmitted. Reflection can be quite complicated but for many purposes it is useful to separate the reflected light into two components. First, a diffuse component, in which the intensity of the reflected light displays relatively little dependence on the angle between the eye, the object's surface, and the light source. Second, a specular component in which the reflection is mirrorlike and highly directional. Typically, the diffuse component is much more influenced by characteristics of the object, while the specularly reflected light often approximates the SPD of the light source. A number of the characteristics of an object affect the way in which it modifies the light it reflects, most notably its chemical composition and the roughness of its surface. Since many objects are heterogeneous in their composition the reflecting characteristics of an object are typically variable and the variation often is found at several different spatial scales, giving rise to both visible patterns and visible texture.

The reflectance of an object (or surface) at a given wavelength is the ratio of the light (number of photons) it reflects at that wavelength to the incident light at that wavelength. The surface spectral reflectance (SSR) of an object is the reflectance of the object at each wavelength (in practice narrow bands of wavelengths) in the visible spectrum. Displaying an object's SSR graphically results in its spectral reflectance curve. In order to achieve a widespread system of colour measurement the illuminants need to be standardized. The most important of these is CIE illuminant $C$-an approximation to average daylight that has the virtue of being reproducible in the laboratory using a standard light source and filter.

The visible light reaching the eye from an (opaque, non-luminous) object is the joint product of its SSR and the SPD of the incident light. Ignoring the effects of scene composition, these exhaust the physical characteristics of objects and light relevant to predicting colour appearance. ${ }^{2}$ What is missing, however, from this physical description is any way of relating this information to perceived colour. First, not all differences in the SSR of the object or the SPD of the illuminant are perceptually detectable. Second, and more importantly, a pair of spectral reflectance curves is little help by itself as to whether or not the corresponding two objects will appear to match in colour when viewed in a given illuminant. Unsurprisingly, the physics of light and its interaction with objects is not enough to explain how we perceive colour.

## 4 Basic physiology of colour vision

Perceived colour is, in complicated ways, dependent on the spectral power distribution of the light reaching the eye from the objects in the scene. This entails that there are mechanisms in the eye/brain that respond differentially to light of different wavelengths. A large amount of research in colour science, going back to the early nineteenth century, concerns the properties of those visual mechanisms that generate the differential response to wavelength.

The process of vision is initiated by the absorption of light by specialized cells in the retina called photoreceptors. A given photoreceptor will respond strongly to light at some wavelengths and much less strongly at other wavelengths, keeping intensity constant. The specification of how strongly a photoreceptor responds to light in the visible spectrum is known as the spectral sensitivity of the photoreceptor. Displaying a photoreceptor's spectral sensitivity graphically results in its spectral sensitivity curve-very roughly a bell shape, with the peak centred over wavelengths to which the photoreceptor is maximally sensitive, and tails of diminishing sensitivity on either side. In spite of responding differently to light of different wavelengths, the behaviour of a single photoreceptor does not by itself contain any information about the SPD of the light to which it is responding. Photoreceptors provide the same response to an absorbed photon, no matter what its wavelength. Although photons of different wavelengths have different probabilities of being absorbed, the response of a single photoreceptor is the same to a dim light at a wavelength to which it is highly sensitive and a brighter light at a wavelength to which it is less sensitive. Since colour vision requires the ability to distinguish between lights with different wavelengths, that means that colour vision requires contributions from at least two types of photoreceptors that differ in their spectral sensitivity. In fact, as we will discuss in the next section, human colour perception is primarily driven by three distinct photoreceptor types.

### 4.1 Rods and cones

The human retina contains two morphologically and physiologically distinct classes of photoreceptors. The rods, so-called because of their characteristic shape, are active mainly at low light levels and play little role in colour vision. ${ }^{3}$ The photoreceptors that play the major role are the
cones (similarly so-called), active mainly at high light levels. The cones are subdivided into three types on the basis of their differences in spectral sensitivity. One type has a peak sensitivity in the short-wavelength end of the visible spectrum and the other two types have closely spaced peaks near the middle of the spectrum. The three cone-types are morphologically indistinguishable, and although their existence was inferred in the nineteenth century in order to explain the observed characteristics of human colour vision, it was only in the late twentieth century that direct measurements of their spectral sensitivities were made, and the light absorbing photopigments they contain were isolated (see Merbs and Nathans 1992).

Since the ability to discriminate between spectrally different stimuli depends entirely on the differences in spectral sensitivity among the three cone-types it is possible to compare the spectral sensitivities required to explain discrimination performance to the measured characteristics of the cones and their photopigments. The agreement is in general very good and simple colour discrimination tasks are an unusual case in which human behaviour (of a very specialized kind) can be predicted on the basis of knowledge of basic neurophysiology. This is possible because the later stages of visual processing preserve the information present in the cone responses and the behavioural response (under carefully controlled conditions) makes use of all the information available.

Although the cone spectral sensitivities largely determine the ability to discriminate among coloured stimuli, their relation to colour appearances is much more complicated. Since the visible spectrum, under ordinary viewing conditions, has a characteristic colour appearance, it is tempting to apply colour labels to the individual cones based on the appearance of the region of the spectrum to which they are most sensitive. The usual labels are "blue" for the short wavelength receptors (S-cones), "green" for the middle wavelength receptors ( $M$-cones), and "red" for the long wavelength receptors ( $L$-cones). This labelling can suggest the theory-sometimes found in popular discussions-that the perceived colour of a light is the result of mixing blue, green, and red, in proportion to the excitation of the corresponding cone-type. However, the usual labelling is misleading and the theory is incorrect. One reason why the labelling is misleading is that the wavelength of peak sensitivity for the L-cones is actually in the yellow-green part of the spectrum. And even if the "red" cones were well-named, the idea that all colours are mixtures of blue, green, and red doesn't fit the phenomenological facts. Admittedly, purple is, in some intuitive sense, a mixture of red and blue, but what about yellow? That seems to be just as basic as red, green, and blue. In any event, yellow doesn't appear to be a mixture of these colours in the way that purple appears to be a mixture of red and blue. Further, how does the mixing theory explain the appearance of a green light that is neither yellowish or bluish? Presumably this is because the light excites only the "green" cones-but because of the overlap in the spectral sensitivities of the three cone-types, there is no such light. As we will see, the problem of explaining colour appearance is a difficult one that does not yet have a fully satisfactory solution.

One important fact about photoreceptors, and neurons in general, helps explain one of the difficulties in predicting colour appearance given just the characterization of the stimulus. Although the relative sensitivity of the photoreceptors to light of different wavelengths is fixed, the absolute sensitivity of the photoreceptors dynamically adjusts to the light level. This adaptation allows the cones to provide usable signals at the very wide range of light intensities that we encounter as we move about the environment. One consequence of this is that the cone outputs provide relatively little information about the absolute intensity of the light stimulating them. The darkest areas of a scene lit by direct daylight are comparable in absolute intensity to the brightest areas of a scene viewed under a typical reading light, even after correcting for the change in pupil size. Another consequence is that the same stimulus can produce very different
cone outputs depending on the recent history of stimulation of the cones. After adaptation to short-wavelength light the S -cones will have decreased sensitivity and a given stimulus will tend to look less blue than it would if the adapting stimulus had consisted of long-wavelength light. Adaptation of various kinds is not unique to the cones but plays a role throughout visual processing.

### 4.2 Chromatic processing in the retina

The processing of visual information begins within the retina itself and its output neurons, the ganglion cells, have very different response properties, both spatial and spectral, from the photoreceptors themselves. ${ }^{4}$ A ganglion cell receives inputs (via other cells) from multiple photoreceptors arranged in a patch on the back of the retina-the cell's receptive field. Ganglion cells have centre-surround receptive fields, meaning that they are excited/inhibited by light in the centre of the receptive field and inhibited/excited by light in the periphery or surround. Importantly for understanding colour vision, the centre and surround can also differ in their sensitivity to light of different wavelengths. In foveal or central vision, where both spatial and spectral discrimination are best, in many cases the centre response is driven by a single photoreceptor while the surround draws on inputs from neighbouring photoreceptors. Consequently, ganglion cells respond best to spectral and spatial contrast. For example, a $+\mathrm{L}-\mathrm{M}$ cell-one whose centre is excited by L-cone input and whose surround is inhibited by M-cone input-will respond well to a small red or white spot on a dark or blue background, less well to uniform red light (which will stimulate the M -cones to some degree) and poorly to uniform white light. Cells with this kind of opponent structure transform the original three cone channels into new channels based on contrast.

Retinal processing also begins a tendency towards specialization that continues through later stages of the visual system. The most important is the subdivision of retinal ganglion cells into two separate processing streams known as the parvocellular $(\mathrm{P})$ and magnocellular $(\mathrm{M})$ streams. The P -stream carries chromatic information ${ }^{5}$ and information about sustained, high spatial resolution aspects of the retinal image. The M -stream is responsive to rapidly changing stimuli, has lower spatial resolution, and is relatively insensitive to chromatic information. These two pathways are driven by the M - and L-cone outputs; the S-cone signal is carried by a separate pathway whose properties are less well understood.

It is important to note that there is no purely chromatic channel originating in the retina. Not only are the outputs of the three cone-types subject to an opponent transformation almost immediately, but the cells in the P -stream combine spectral, intensity, and spatial information. It is only by comparing the responses of multiple cell-types to the same stimulus that it is possible to separate the chromatic information from the spatial and intensity information. It is not until the cortex that cells are encountered whose responses disambiguate the spatial and spectral information that jointly determines the activity of cells earlier in the visual pathway.

## 5 The psychophysics of colour

So far we have looked at colour vision from the point of view of physiology. Alternatively, we could look at how people (and other animals) behave in response to coloured stimuli. This kind of approach, in which very constrained responses to carefully constructed and varied stimuli are measured and analysed, has been central to colour science. As we saw in discussing the cone sensitivities, the physiology is intimately connected with measures of psychophysical performance, like spectral discrimination. Colour science has been traditionally characterized by an
unusually integrated approach to its subject matter with studies of animal behaviour motivating and justifying physiological theorizing and vice versa. To give just two examples, the most widely used values for the cone spectral sensitivities derive from behavioural data, and what is known about the colour discrimination behaviour of many non-human animals is largely based on properties of the photopigments found in their eyes.

### 5.1 Trichromacy, primaries, and colour spaces

Any colour can be matched with an appropriate mixture of only three primaries. As might be suspected, this is a consequence of trichromacy, that exactly three types of photoreceptors contribute to human colour vision.

The claim about matching needs to be qualified, in large part because of the many complicated effects of the viewing context on perceived colour. These effects can be largely discounted if we create a very simple perceptual situation, e.g. a bisected circle on a uniform neutral background. The two halves of the circle will appear identical in colour if and only if the light reaching the eye from each half produces the same output from each of the three cone-types. In this (somewhat artificial) situation, we can choose three lights such that, for any light projected on to the left half of the circle, an appropriate weighted mixture of the three lights projected on to the right half will result in uniform cone output across the circle's retinal image.

All this only applies to additive mixtures, like mixtures of lights in which each element of the mixture simply adds to the light reaching the eye. In subtractive mixtures, like pigment mixtures, the contributions of the components of the mixture to the visual stimulus are much more complicated and it may take more than three elements to match an arbitrary stimulus. Another qualification is that some matches will require the addition of one of the primaries to the light to be matched rather than to the other two primaries-in effect, negative amounts of one of the primaries. A final point to note is that there are numerous sets of primaries. In fact, any three lights, no two of which can be mixed to match the others, will serve as primaries. The traditional red, green, and blue additive primaries used in television and computer screens have the virtue of matching a very large set of lights without using any negative amounts, but this is only of technological significance.

These facts about matching and primaries lead to an obvious method for a systematic representation of colour stimuli: represent the colour of each stimulus by the amounts of a certain set of primaries required to match it. In such a system, stimuli with the same coordinates will appear the same colour (at least in highly constrained viewing conditions). And given the coordinates of a stimulus in such a system, it will be possible to produce a new stimulus that will be an exact match by adding together the specified amounts of the three primaries. Since coordinates in one system can be transformed into corresponding coordinates in any other, the new stimulus need not even be constructed using the original primaries to guarantee a match. Many of the standard colour spaces used in science and industry employ this basic method. For example, the widely used CIE XYZ space is just a set of functions that take the spectral power distribution of a light into the amounts of three specially chosen primaries that match that light. These functions are based on colour-matching data collected on a relatively modest number of individuals in the early twentieth century. Many other more recent standards have a similar structure. RGB coordinates use an idealized set of monitor primaries to represent colour and although the primaries are very different the basic principle is the same.

Such systems for representing colour based on three primaries are very useful for many purposes in research and industry, but they have two significant drawbacks. First, they do a relatively poor job of representing perceived colour similarity, especially for stimuli that are distant
from each other in the space. ${ }^{6}$ Second, a system based solely on matching will fail to capture perceived colour since two stimuli may change their colour appearance substantially while still remaining matched. The fundamental problem is that the simple colour matching experiment that motivates these systems idealizes away from many factors that profoundly affect perceived colour.

### 5.2 Colour appearance and opponent-process theory

Neither the physics of light, nor the cone outputs, nor the primaries used in matching provide an adequate basis for understanding colour appearance. One very influential attempt to provide the outlines of a theory of colour appearance involves combining psychophysical experimentation with speculative physiology. As we saw earlier (section 4.1), attempting to account for colour appearance in terms of the three cone-types leaves us with one too few basic colours. Red, yellow, blue, and green all have a plausible claim to being basic colours, unlike purple, orange, turquoise, and olive which appear to be mixtures (in some intuitive sense) of the basic colours. In addition, these four basic colours are naturally sorted into two "opponent" pairs: red and green on the one hand and blue and yellow on the other. Red and green are opposed in the sense that there are no reddish greens or greenish reds, and similarly for yellow and blue. Red and green are so famously opposed that there is a significant philosophical literature devoted to explaining the nature of the opposition. ${ }^{7}$ Opponent-process theory is a physiological hypothesis put forward to explain these observations, together with many others.

The core of opponent-process theory is that information about the spectral characteristics of a stimulus is carried by two opponent channels (plus a non-opponent channel for intensity). In the simplest model, one channel is generated by subtracting the M-cone signal from the L-cone signal (L-M) while the other channel results from subtracting the sum of the L- and M-signals from the S-cone signal (S-(L+M)). The L-M (or red-green) channel results in the perception of reddishness when positive and greenishness when negative, while the $\mathrm{S}-(\mathrm{L}+\mathrm{M})$ (or yellow-blue) channel results in the perception of bluishness when positive and yellowishness when negative. Thus a stimulus that looks bluish-red will produce a high positive value for the L-M channel and a (less high) positive value for the $\mathrm{S}-(\mathrm{L}+\mathrm{M})$ channel. Since no channel can produce a signal that is both negative and positive, the hue incompatibilities mentioned above are explained. This framework, motivated by phenomenological observations about basic colours and opponency, proved to be a powerful unifying tool that allowed a simple and intuitive understanding of diverse set of colour phenomena. When the chromatic opponency of cells in the peripheral visual pathway was first discovered in the 1960s it seemed as if direct experimental support for the hypothetical opponent processes had been found. Unfortunately, in the subsequent decades the status of opponent-process theory has become less clear. Although chromatic information is encoded in the visual pathways using opponent coding, the response properties of these cells don't match the characteristics of the psychophysically characterized opponent-processes. Unlike the good fit between the measured cone spectral sensitivities and the hypothesized sensitivities required to explain the psycho-physical discrimination data, the hoped-for match between physiology and psychophysically characterized opponent processes has failed to materialize. Although this is an area of current controversy it seems safe to say that the simple opponentprocess model that seemed so promising in the late twentieth century is at best a very rough approximation.

The uncertain status of opponent-process theory leaves the field with no unified physiological account of the elementary facts about colour appearance that helped motivate it. Although there have been claims to find some basis for the special status of the unique hues in the response
properties of some cortical neurons, the claims are controversial and anyway don't provide the kind of unifying framework that earlier looked to be on the cards (see Stoughton and Conway 2008; Conway and Tsao 2009; Mollon 2009; Wool et al. 2015). Further, the phenomenological foundations have themselves been disputed, with some claiming that there are more than four basic colours, or even that the notion of a basic colour is suspect (Saunders and van Brakel 1997).

These controversies aside, there is still a need for colour order systems that capture central facts about colour appearance and that provide a more natural representation of colour similarity than the primary based models that were discussed in section 5.1. There are a number of such systems and they all share one significant feature: the colours are represented in terms of three dimensions. It is tempting to assume that this is because the three-dimensionality that originates with the cones is maintained in the ultimate cortical representation of colour, but if so this is a peculiarity of colour vision, not an instance of general truth about perception. The human auditory system samples the frequency spectrum much more densely, but the representation of pitch is essentially one-dimensional. Moreover, there are reasons to doubt that three-dimensions are, in fact, capable of fully capturing all of the variation in colour appearance (Fairchild 1998). Nevertheless, three dimensions do an efficient job for most practical purposes.

One way to construct an ordering system that reflects colour appearances starts with the phenomenological claims that underpinned opponent-process theory. The Natural Colour System (NCS) is an example of a system with this structure (Hård et al. 1996a, 1996b; Kuehni 2003: 301-9). The NCS represents colour using two opponent axes (red-green and blue-yellow) and a non-opponent lightness axis. No physiological interpretation is associated with this system, and it is not directly tied to any system of primary-based matching. To classify colours with the NCS, samples are matched to standards generated in accordance with the underlying opponent model. As the name suggests, the system is intended to be a better fit to our perceptual representation of colour than other alternatives. In this form, the representation of colour embodied in opponent-process theory can be maintained independently of its success or failure as a physiological theory.

A widely used alternative is to represent colour in terms of three dimensions of hue, brightness, and saturation (HBS). These representations give rise to the familiar colour solid with hue being represented by a circle around the origin, brightness by the vertical axis, and saturation by horizontal distance from the origin. The popular Munsell system is a variant of the HBS system with its three dimension of hue, value (brightness), and chroma (a relative of saturation). One reason for the popularity of the Munsell system is that brightness and saturation are very difficult to estimate visually and the Munsell system has a physical realization that allows colours to be placed in the system by comparison to samples. Although the system was constructed to do a good job of capturing perceived similarity, the visual inaccessibility of the brightness and saturation dimensions suggest that it is not a good match for the way colour is represented by the visual system.

### 5.3 Contrast, adaptation, and other psychophysical effects

As we saw in the discussion of basic physiology above, the cones do not provide a fixed response to a fixed stimulus, and the channel carrying chromatic information from the retina to the brain combines spatial and spectral information. These and other physiological features have measurable (and sometimes very large) effects on how we perceive colour.

To start with a simple example, we are all familiar with the large changes in perceived lightness and colour when going inside on a bright day. Many parts of the visual system (pupil, cones,
retinal ganglion cells, etc.) have adapted to the bright light and, at varying speeds, will then adapt to the much dimmer (and spectrally different) illumination indoors. The initial perception of dark and desaturated colours gradually moves back towards the brighter and more saturated colours perceived outside and there may be shifts in hue as well. One way to understand the overall effect of adaptation at the various levels of visual processing is that the visual system changes to maximize the amount of information it can extract from the visual stimulus. For example, as noted in section 4.1, the range of responses that the cones can produce is orders of magnitude smaller than the variation in the intensity of the stimulation they receive. If the cones did not adapt to changes in light intensity then they would provide useful information about only a very narrow range of stimuli. By becoming less sensitive as the stimulus intensity increases and more sensitive as it decreases, the cones preserve their ability to signal differences in stimulation across a much broader range of stimuli. One consequence of the various forms of adaptation is that large changes in the stimulus (resulting from changes in the illumination) typically produce much smaller changes in perceived colour once adaptation has run its course. Adaptation contributes to the relative stability of perceived colour across changes in illumination known as colour constancy (discussed in more detail in section 5.4 below).

As we have seen, the chromatic and spatial characteristics of stimuli interact in early colour processing. One illustration of this fact can be found in the familiar phenomena of colour contrast. If a neutral grey square is viewed surrounded by a larger coloured background it will appear tinted with a hue contrary to that of the background: reddish backgrounds thus induce greenish tints and greenish backgrounds induce reddish tints. Not all spatial effects involve the induction of a contrasting hue and, in assimilation, the colour of thin, but clearly visible, lines spreads to neighbouring areas. It needn't be only the directly adjacent regions of a scene that influence perceived colour. In the watercolour illusion, the colour of an appropriately chosen border spreads to large areas of the white space it encloses (Pinna et al. 2001). Even simple patterns, like a disc surrounded by concentric rings, can produce greater effects on perceived hue than a uniform background (Monnier and Shevell 2003). The causes of these kinds of effects are understood to varying degrees but in general they fall into two overlapping classes. First, there is averaging over stimulus areas at different spatial scales resulting from the underlying physiology. For example, assimilation is due, in part, to the fact that the visual system has higher resolution for achromatic contrast than for chromatic contrast. The dark lines in a typical stimulus that produces assimilation are visible to the luminance channel but not resolvable by the chromatic channels which then averages their lower lightness in with surrounding areas. Similar effects can occur with hue alone since there are many fewer S-cones than there are L- and M-cones so the averaging occurs over larger areas for the $S$-cone input than for the other two cone-types. ${ }^{8}$ A different way of looking at these kinds of effects is that they are consequences of the visual system's attempt to use all of the information available to it in arriving at a representation of the spatial layout of the perceived scene and to assign visual features to different regions of it. Chromatic information is useful in extracting the spatial features of the scene from the stimulus and the spatial layout is useful in generating stable and useful colour assignments to the different areas of the scene. We will return to some of these issues later in the discussion of colour constancy.

The variety and quantity of informative and sometimes surprising interactions known to exist between perceived colour and various features of the stimulus other than the SPD of the light coming from an object is much too large to catalogue here. There are two important points worth keeping in mind with respect to the large literature on the psychophysics of colour vision. First, the psychophysics is often very informative as to the underlying physiological mechanisms, and much of the empirical literature in colour psychophysics is aimed at illuminating the
underlying physiology using behavioural data collected in response to carefully controlled stimuli. Knowledge of the existence and response characteristics of the three human cone-types was almost entirely based on psychophysical data. For these purposes, the choice of stimuli need not reflect important features of the kind of stimuli encountered outside the laboratory. Second, the fact that many factors other than the character of the light reaching the eye from an object can influence its perceived colour should not be surprising. The point of vision is not to accurately characterize the proximal stimulus but rather to guide action. For the purpose of guiding action the properties of a distal object are important, and so to ignore factors other than the light an object sends to the eye would be to throw away valuable information about it.

### 5.4 Colour constancy

We have already mentioned simultaneous contrast, in which the perceived colour of an object is influenced by the colour of its surround. This phenomenon illustrates the important point that the relation between stimuli and perceived colour cannot be fully understood by taking each point in the scene before the eyes in isolation. Holding the subject's perceptual apparatus constant, the perceived colour of an object is determined by the character of the light produced by the entire scene before the eyes.

Colour constancy, the stability of perceived colour across alterations in the character of the illuminant, is another manifestation of these non-local influences. ${ }^{9}$ Recall that the light reaching the eye from an area of a surface is the joint product of the SPD of the illuminant and the SSR of the surface. As the illuminant varies, so does the SPD of the light reaching the eye. In spite of this variation in the local visual stimulus, under many conditions the perceived colour of an object will not appreciably change. However, it is an important (although entirely unsurprising) fact that colour vision (in humans and other animals) is only approximately colour constant. (Similarly, shape constancy is only approximate.) It is easy to devise scenes and viewing conditions for which constancy effects are minimal or non-existent and, as it happens, these kind of viewing conditions are favoured for colorimetric and many experimental tasks. An interesting but virtually intractable question is how much colour constancy human colour vision displays under natural conditions. The difficulty is partly conceptual: is it constancy in colour phenomenology or colour judgement that we are attempting to measure? It is also partly technical: how can we construct a representative sample of natural viewing conditions and scenes in order to make laboratory measurements? In spite of these problems there has been a great deal of both experimental and theoretical work done on the nature of the constancy mechanisms.

One important but controversial approach to colour constancy treats it as the result of the visual system's attempt to estimate object reflectances from the light reaching the eye. The perceived colour of objects is approximately constant under many conditions because under those conditions the reflectance estimate generated by the visual system is reasonably accurate. In this framework, the most common strategy is to first generate an estimate of the SPD of the illumination in a scene and use that estimate to compute the reflectance of an object from the light reaching the eye from that object. A simple example of a theory of this kind involves the assumption that the environment, on average, is grey. That is, if the reflectances of the objects in a scene are averaged together the resulting curve will be flat across the visible spectrum and approximately $1 / 2$. Given this assumption, averaging the light reaching the eye across the entire scene and dividing at each wavelength by $1 / 2$ gives an estimate of the illuminant on the scene. Unfortunately, the grey world assumption is false for many scenes in which humans have reasonably good constancy, so this cannot be the entire explanation. More sophisticated theories of this kind have been developed and this is still an area of active research.

A wide variety of other factors have been invoked to explain constancy effects in various circumstances. Comparing the ratios of cone outputs across a scene contains important information about whether changes in the retinal image are due to changes in the illumination or to changes in the surface (although not about the absolute reflectance) (see Foster 2003). There are many different types of contrast, spatial and spectral, that seem to have some relationship with colour constancy. There can also be a powerful influence of perceived scene geometry on how the visual system disentangles illumination and surface properties. Although human colour vision displays some degree of several different kinds of constancy there is no current consensus on the best explanation of the various constancy phenomena or even of the best way to characterize those phenomena. ${ }^{10}$

## 6 Colour in the cortex

### 6.1 The role of chromatic information in the cortex

Chromatic discrimination is extraordinarily precise in some ways and extraordinarily coarse in others. Extremely small differences in SPD are discriminable in the right circumstances and, by some measures, the visual system is better at detecting this type of chromatic contrast than achromatic contrast (differences in the overall intensity of illumination). On the other hand, the spatial and temporal resolution for chromatic contrast is much worse than for achromatic contrast, and consequently chromatic contrast makes hardly any contribution to high-resolution spatial or temporal vision. This and other factors lead to a picture of cortical colour processing in which chromatic and achromatic information are combined in the eye and mid-brain areas but separated in the cortex, and contribute very differently to visual processing. In particular, chromatic contrast does not contribute to spatial vision in ordinary contexts and only plays a role in perceiving colour (and via that in tasks like object identification). Contrariwise, information derived from the achromatic signal plays only a minimal role in perception of colour. The one thing that's safe to say about colour in the cortex is that this picture has been rejected, at least in anything like its original form. In cortical area V1, the first cortical visual area, there are very few cells that are responsive only to chromatic signals and even that small minority are also orientation sensitive, so their behaviour reflects both chromatic and spatial information. The overwhelming majority of cells in V1 are sensitive to both chromatic and achromatic inputs. S-cone input, which does not contribute to the achromatic pathway, is found throughout the visual areas that receive inputs from V1 including areas that have nothing to do with perceiving colour, like area MT which is thought to play a role in motion perception. Similarly, chromatic contrast plays a role in spatial vision and vice versa, as can be shown using psychophysical methods. Although the precise details are still a matter of controversy, it's clear that colour is a cue used by the brain to perform a variety of tasks and that the information about the SPD of the stimulus delivered by the cones is utilized for many purposes other than that of discriminating and recognizing colour. ${ }^{11}$

### 6.2 The organization of cortical colour processing

There are two central issues involved in accommodating this new understanding of the role of chromatic information in the cortex. The first is partly conceptual. The results of the previous section are often described as showing that colour vision contributes to spatial vision. Although all that is intended is that chromatic information contributes to spatial vision, it can be read as implying that colour as perceived contributes to spatial vision-and this is a much more
controversial claim. It's important to keep separate the role of chromatic information in, for example, the perception of shape and the perception of hue. It is unlikely that perceived hue is an input to the perception of shape even though both draw on the chromatic signal originating in the cones. This is supported by studies of achromatopsia (colour blindness resulting from cortical damage). Some achromatopsics can continue to perceive shapes that are defined solely by chromatic information even though they cannot discriminate, sort, or recognize hues at all. They have lost the ability to see colour but not the ability to utilize chromatic information for other visual functions. ${ }^{12}$

The second issue is primarily empirical. Although the explanation of why human beings experience colour as they do is presumably to be found in the cortex, the identified cortical cells and cortical areas do not seem well-suited to explaining the details of how we visually represent colour. Related to this is the extended controversy over whether there is a cortical area specifically dedicated to colour and, if so, where it is. Much of this controversy has centred on Zeki's controversial identification of the human analogue of macaque V4 as the brain area responsible for the perception of colour (Lueck et al. 1989). What does seem clear is that there are neurons responsive specifically to chromatic information in V1 and there are clusters of such neurons in areas outside of V1 as well. Our ability to discriminate and identify colour presumably relies on these neurons but going beyond that is highly speculative. Cortical processing of colour, beyond the clarification of the role of chromatic information in spatial vision, remains a confused (and confusing) topic. (For recent overviews see Conway 2014 and Johnson and Mullen 2016.)

## 7 Defects of colour vision and naming

Colour vision, like any other biological characteristic, varies from individual to individual. A familiar and extreme example of such variation is that a non-negligible proportion of human beings are colour "blind", most of them being specifically insensitive to the difference between red and green. In light of the salience of colour and, in particular, the striking difference between red and green for those of us with normal colour vision, it is a surprising fact that colour blindness was first clearly characterized around 1800. Thus colour blindness does not appear to be a functionally significant problem in most practical contexts.

Colour blindness is of great theoretical interest. Study of such defects has proven very illuminating in understanding normal colour vision and also raises some interesting questions about the contribution of the photoreceptors to the character of colour experience. Most colour blind individuals are not, in fact, colour blind in any strict sense of the phrase. Rather their colour vision differs from that of colour normal individuals in several well-defined respects, none of which amount to a complete loss of colour vision. The most common form of colour blindness is dichromacy. Dichromats require only two primaries in matching experiments, and lack the ability to discriminate some stimuli that are readily discriminable by normal (trichromatic) subjects. For example, all dichromats will accept a match between some monochromatic lights and a white light. Dichromacy results from a loss of function of one of the three cone photoreceptor types, and comes in three corresponding forms. ${ }^{13}$

What is commonly called red-green colour blindness actually consists of two different defects depending on whether it is the long or middle wavelength receptor whose function has been lost. Protanopes have no functioning long wavelength receptor and deuteranopes have no functioning middle wavelength receptor. They can be differentiated by, among other methods, the loss of long wavelength sensitivity relative to normals that is found in protanopia but not in deuteranopia. Although both protanopes and deuteranopes are unable to distinguish spectral lights in the middle to long wavelengths that appear green to red to normal observers, thus the
name red-green colour blind, their ability to discriminate non-spectral lights is substantially different. Subjects having any of the three forms of dichromacy will accept all matches made by a normal observer, although not vice versa. Protanopia and deuteranopia are the overwhelmingly most common forms of dichromacy, and most cases are the result of recessive inherited abnormalities in genes on the X chromosome which code for the photopigments contained in the long and middle wavelength photoreceptors. Consequently, red-green colour blindness is much more common among males than among females. The third form of dichromacy, tritanopia, is much less common and is due to the loss of function of the short wavelength receptor. Monochromacy is much rarer than dichromacy and is most often due to the loss of all cone function. Monochromatic individuals are only able to make light-dark distinctions and are strictly speaking colour blind.

The genes coding for the three cone photopigments have now been isolated and sequenced. This achievement has provided new methods for understanding the early stages of colour vision and also for investigations of the evolution of colour vision. It is now known precisely what genetic abnormalities are responsible for the two varieties of red-green dichromacy and how these abnormalities affect the spectral sensitivity of the photopigments in colour blind individuals. The genetics has also helped in the discovery of the detailed structure of the photopigment proteins themselves which in turn has led to a more detailed understanding of normal variation in human colour vision (Neitz and Neitz 2011). In addition, it is now possible, using the methods of molecular genetics, to trace evolutionary relationships among the photopigments found in different species.

With the precise characterization of the different forms of colour blindness in the nineteenth century arose a puzzle as to what the visual experience of colour blind individuals is like. Protanopes and deuteranopes, for example, perceive only a single hue in the regions of the spectrum between 550 and 700 nm , but it is difficult to get empirical evidence for which hue it is. Opponent process theory suggests that, as protanopes and deuteranopes have no functioning red-green opponent channel, they should see only yellow, blue, black, and white. But colour blind subjects talk about colour just like the rest of us, only making mistakes normal observers would never make. They know that grass is green and tomatoes are red and although deuteranopes may have trouble telling the difference between ripe and unripe tomatoes, they will not say they are yellow or blue. Some very unusual individuals have normal vision in one eye and a colour deficiency in the other. These subjects might seem ideal, since they are familiar with the full range of colour experience due to their normal eye, and so can report on what they see through their colour-deficient eye. Unfortunately, the small but much discussed literature on such subjects has produced more controversy than consensus. (For a brief review see Boynton 1979: 380-2).

Most defects of colour vision are due to receptoral abnormalities. These cases are in most respects well understood, partly because there are many examples to study and partly because the role the photoreceptors play in colour vision is well understood. But receptoral abnormalities are not the only cause of defects of colour vision: as mentioned in section 6.2, damage to areas of the visual cortex is another cause. These achromatopsic disorders are, in general, less well characterized and understood than the much more common disorders discussed above. In addition, there is very little understanding of what contribution the damaged areas make to normal colour vision.

In some well studied cases of achromatopsia it has been established that all three cone-types are present and contributing to visual functioning. Even more striking is that serious impairments of colour vision can be accompanied by essentially normal perception of luminance resulting in subjects who appear to perceive the world in shades of white, grey, and black. Not
all cases of achromatopsia are total and there is a great deal of variation in the severity of the impairment. There can be some remaining degree of colour vision and the defect may even be limited to some areas of the visual field. However, the specific characteristics of the colour abnormality in at least some cases of achromatopsia are very different from the forms of dichromacy.

Cortical damage can cause other kinds of colour-related deficits where the pattern of which abilities are spared and which are preserved is complicated. Colour agnosia is an inability to recognize the colours of seen objects with other aspects of colour vision remaining apparently intact. One colour agnosic performed normally on many non-verbal tests of colour perception, had a normal colour vocabulary and was able correctly to remember common colour associations, for example that grass is green and blood is red. When presented with an object and asked for its colour he would reply with a colour term, but his performance was no better than chance. He performed well on tasks that involve arranging colour samples in terms of similarity but poorly on sorting them into categories on the basis of similarity (van Zandvoort et al. 2007).

## 8 Animal colour vision

Some degree of colour vision is widely distributed throughout the animal kingdom, and appears to have evolved independently in several groups. Almost all vertebrates that have been studied possess some form of colour vision, although many only have a rudimentary ability which may not play a significant role in guiding behaviour. Although comparatively few have been tested, many invertebrates also possess colour vision, which in some (e.g. bees) is highly developed. The number of photoreceptor types and the spectral characteristics of the photoreceptors varies from species to species. Among mammals only (some) primates are known to have trichromatic colour vision. All other species of mammal that have been studied are dichromats with possibly a few, such as rats, lacking colour vision altogether. Some birds and fish are tetrachromats. ${ }^{14}$ Further, the spectral range over which their vision extends is broader, particularly into the ultraviolet. Colour vision in these groups is phylogenetically older, and some respects more highly developed, than it is among mammals.

An organism is said to have colour vision if and only if it is able to discriminate between some spectrally different stimuli that are equated for brightness (or luminance). ${ }^{15}$ There are two basic methods for determining the presence or absence of colour vision in non-human organisms. The first is behavioural: the organism's ability to discriminate equiluminant stimuli is tested directly. A complication arises because stimuli that are equiluminant for a human observer will not, in general, be equiluminant for a non-human observer. The luminance of stimuli for an organism can be equated if its spectral sensitivity function (the function from stimulus wavelength to stimulus brightness) can be determined. Alternatively, the relative luminance of the stimuli can be randomly varied over a wide range, assuming that consistently successful discrimination can only be based on colour differences. (For a review of these techniques see Jacobs 1981: 5-11.) Both techniques are somewhat tedious, and consequently have only been used to investigate a relatively small number of animals. The second method is physiological: the visual capacities of an organism are inferred from information about the physiological characteristics of its visual system. For example, it is possible to measure the absorption spectra of individual photoreceptor cells using a technique known as microspectrophotometry. Establishing the existence of two cone photoreceptor types in this way provides reasonably good evidence that the organism in question is a dichromat. These measurements, and other physiological techniques, although not easy to perform, are often less time-consuming than behavioural methods.

What selective advantages does colour vision confer? This is a large and complex question, but it is broadly accepted that the selective advantages for colour vision systems like ours include object recognition, detection of targets against variegated backgrounds, and perceptual segregation of figure from ground by similarity in colour. An even larger and more complex question is how variation in colour vision across species is connected with variation in the visual environment and more generally the species' ecological niche. As we have also seen the sensitivity to the SPD of stimuli that is crucial to colour vision also plays a role in spatial vision. Any explanation of the evolution of colour vision will have to consider the full range of visual tasks in which chromatic information is involved.

## Notes

1 Two notable omissions from this chapter include the science behind technologies like paint systems and colour reproduction systems, and colour language, which displays interesting patterns across the world's languages, and whose connection to perception is disputed. The second is of more relevance to philosophy: for recent reviews and further references see Kay and Regier 2006; Regier and Kay 2009.
2 Also ignored is the fact that some surfaces have direction-dependent reflectances, and the phenomenon of fluorescence - the absorption of light at one wavelength and its re-emission at a longer wavelength.
3 Although it is standard to ignore rod input into colour vision, rods do influence perceived colour at intermediate and low light levels. We will follow the usual practice and ignore rod intrusion. See also footnote 5 below.
4 Very recently it has been discovered that some ganglion cells are intrinsically photosensitive and although the primary function of this sensitivity is synchronizing circadian rhythms with the day-night cycle they may, under some conditions, influence perceived colour (Horiguchi et al. 2013).
5 That is, information about wavelength. Chromatic or wavelength information may be used for detecting features other than colour, for instance edges.
6 There are modifications of the CIE standard, like the CIE $L \star A B$ space, that attempt to correct for this problem but they are only partially successful and almost entirely ad hoc.
7 See Hardin 1993: x-xii. As Hardin points out, the perception of reddish-green (and bluish-yellow) can be induced under special conditions. See Crane and Piantanida 1983; Billock and Tsou 2010.
8 The sparse distribution of S-cones has other consequences for vision. There are no S-cones at all in the very centre of the fovea rendering all normal human observers yellow-blue colour blind for small, centrally presented stimuli. This is presumably an adaptation to support high-resolution spatial vision which is driven by the L-and M-cone inputs. One benefit of the low-resolution of the S-cone channel is that it substantially mitigates the very high chromatic aberration of the optics of the eye.
9 Strictly speaking this is illumination-independent colour constancy, not full colour constancy. Full colour constancy requires constancy as both the illuminant and the arrangement of objects that make up the scene are varied (Brainard and Maloney 2011: 4). (For completeness one can also add constancy as the viewing medium is varied: see Brown 2003: 253-4.)
10 For a survey of the illuminant estimation approach see Brainard and Maloney 2011. For a recent theoretical challenge to that approach see Logvinenko et al. 2015. For more general survey of what is and isn't known about colour constancy see Foster 2011; Olkkonen and Ekroll 2016.
11 A useful recent overview is Johnson and Mullen 2016. For the S-cone picture see Conway 2014.
12 For more discussion of these issues see Akins and Hahn 2014.
13 More common than dichromacy is anomalous trichromacy. Although anomalous trichromats have three functioning receptor types, one of the receptors has its spectral sensitivity shifted from the normal position. Typically, this results in poorer than normal colour discrimination performance as well as other abnormalities. Corresponding to each form of dichromacy described below there is a form of anomalous trichromacy.
14 The receptoral story for birds and non-mammalian vertebrates generally is particularly complex. Their cones contain oil droplets through which light is filtered before interacting with the photopigments, and different types of droplet may be found in combination with the same type of photopigment. (See, for example, Bowmaker 1977.)
15 There has been some discussion of the adequacy of this criterion as a sufficient condition for possession of colour vision. See Hilbert 1992; Thompson et al. 1992; Thompson 1995: 141-214.

## References

Akins, K. A., and M. Hahn. 2014. More than mere colouring: the role of spectral information in human vision. British Journal for the Philosophy of Science 65: 125-71.
Billock, V. A., and B. H. Tsou. 2010. Seeing forbidden colors. Scientific American 302: 72-7.
Bowmaker, J. K. 1977. The visual pigments, oil droplets and spectral sensitivity of the pigeon. Vision Research 17: 1129-38.
Boynton, R. M. 1979. Human Color Vision. New York: Holt, Rinehart and Winston.
Brainard, D. H., and L. T. Maloney. 2011. Surface color perception and equivalent illumination models. Journal of Vision 11: 1-18.
Brown, R. 2003. Background and illuminants: the yin and yang of colour constancy. Colour Perception: Mind and the Physical World, ed. R. Mausfeld and D. Heyer. Oxford: Oxford University Press.
Conway, B. R. 2014. Color signals through dorsal and ventral visual pathways. Visual Neuroscience 31: 197-209.
Conway, B. R., and D. Y. Tsao. 2009. Color-tuned neurons are spatially clustered according to color preference within alert macaque posterior inferior temporal cortex. Proceedings of the National Academy of Sciences 106: 18034-9.
Crane, H. D., and T. P. Piantanida. 1983. On seeing reddish green and yellowish blue. Science 221: 1078-80.
Fairchild, M. D. 1998. Color Appearance Models. Reading, MA: Addison-Wesley.
Foster, D. 2003. Does colour constancy exist? Trends in Cognitive Sciences 7: 439-43.
Foster, D. H. 2011. Color constancy. Vision Research 51: 674-700.
Hård, A., L. Sivik, and G. Tonnquist. 1996a. NCS, natural color system-from concept to research and applications. Part I. Color Research \& Application 21: 180-205.
Hård, A., L. Sivik, and G. Tonnquist. 1996b. NCS, natural color system-from concept to research and applications. Part II. Color Research \& Application 21: 206-20.
Hardin, C. L. 1993. Color for Philosophers (expanded edition). Indianapolis: Hackett.
Hilbert, D. R. 1992. What is color vision? Philosophical Studies 68: 351-70.
Horiguchi, H., J. Winawer, R. F. Dougherty, and B. A. Wandell. 2013. Human trichromacy revisited. Proceedings of the National Academy of Sciences 110: E260-9.
Jacobs, G. H. 1981. Comparative Color Vision. New York: Academic Press.
Johnson, E. N., and K. T. Mullen. 2016. Color in the cortex. Human Color Vision, ed. J. Kremers, R. C. Baraas, and N. J. Marshall. Cham, Switzerland: Springer.
Kay, P., and T. Regier. 2006. Language, thought and color: recent developments. Trends in Cognitive Sciences 10: 51-4.
Kuehni, R. G. 2003. Color Space and Its Divisions: Color Order from Antiquity to the Present. Hoboken, NJ: Wiley-Interscience.
Logvinenko, A. D., B. Funt, H. Mirzaei, and R. Tokunaga. 2015. Rethinking colour constancy. PLoS ONE 10: e0135029.
Lueck, C. J., S. Zeki, and K. J. Friston. 1989. The colour centre in the cerebral cortex of man. Nature 340: 386-9.
Merbs, S. L., and J. Nathans. 1992. Absorption spectra of human cone pigments. Nature 356: 433-5.
Mollon, J. D. 2009. A neural basis for unique hues? Current Biology 19: R441-2.
Monnier, P., and S. K. Shevell. 2003. Large shifts in color appearance from patterned chromatic backgrounds. Nature Neuroscience 6: 801-2.
Neitz, J., and M. Neitz. 2011. The genetics of normal and defective color vision. Vision Research 51: 633-51.
Olkkonen, M., and V. Ekroll. 2016. Color constancy and contextual effects on color appearance. Human Color Vision, ed. J. Kremers, R. C. Baraas, and N. J. Marshall. Cham, Switzerland: Springer.
Pinna, B., G. Brelstaff, and L. Spillmann. 2001. Surface color from boundaries: a new 'watercolor' illusion. Vision Research 41: 2669-76.
Regier, T., and P. Kay. 2009. Language, thought, and color: Whorf was half right. Trends in Cognitive Sciences 13: 439-46.
Saunders, B. A., and J. van Brakel. 1997. Are there nontrivial constraints on colour categorization? Behavioral and Brain Sciences 20: 167-79.
Stoughton, C. M., and B. R. Conway. 2008. Neural basis for unique hues. Current Biology 18: R698-9.
Thompson, E. 1995. Colour Vision. London: Routledge.

Thompson, E., A. G. Palacios, and F. J. Varela. 1992. Ways of coloring: comparative color vision as a case study for cognitive science. Behavioral and Brain Sciences 15: 1-74.
van Zandvoort, M. J. E., T. C. W. Nijboer, and E. de Haan. 2007. Developmental colour agnosia. Cortex 43: 750-7.
Wool, L. E., S. J. Komban, J. Kremkow, M. Jansen, X. Li, J.-M. Alonso, and Q. Zaidi. 2015. Salience of unique hues and implications for color theory. Journal of Vision 15: 1-11.

