2

Colour description/specification systems

B R I G G, University of Leeds, UK

2.1 Introduction

Colour has always been important, but especially so in the modern world. We only have to look around us to see the variety of colours produced on textiles, painted surfaces, paper and plastics. In most cases the colour is an important factor in the production of the material and is often vital to the commercial success of a product. Textiles often sell as much on colour and design as on quality of materials. In our homes, colours of carpets, curtains and paints are carefully chosen to produce a pleasing effect. The same colours may be represented as colour prints, for example in mail order catalogues, or on computer screens for design work or in e-commerce. In some cases, such as foods, we use the colour to judge the quality of the material, while in other cases, such as packaging, the colour is important in attracting customers. In almost all applications it is difficult to produce exactly the colour required and to ensure that successive batches are the same colour as each other. Raw textile materials such as cotton and wool vary from batch to batch as do dyes and pigments, despite strenuous efforts by producers to try to ensure good quality control. There are normally two main considerations for the producer of coloured products (or computer images). Firstly, whether the colour is reasonably close to that desired. Secondly, whether repeat samples or batches match each other. For example, successive batches of paint, or arrays of clothing items displayed together should match very closely. In this case, it is the differences in colour between samples or within a sample that matter. Two alternative methods are available to assess the correctness of the colours. For many years, only visual assessment was available. This worked fairly well, but there were problems. It is difficult to remember colours accurately and difficult to describe differences in colour. Experienced colourists have a good idea as to the size of difference which might be tolerated, but there is no objective way to settle differences of opinion between say a dyer who considers that a match is close enough, and a customer of the opposite opinion.

2.1.1 The nature of measurement

In all branches of science and engineering, measurement plays an important part. In commerce materials are usually bought and sold by weight or volume. Without standardised systems for measuring mass, length and time, modern life would be very difficult. It is obvious that a standard system for measuring and specifying colour is equally desirable. However, there are important differences between colour and, for example, length. The colour of an object depends on many factors, such as lighting, size of sample, and background and surrounding colours. Much more importantly, colour is a subjective phenomenon and depends on the observer. The measurement of subjective phenomena, such as colour, taste and smell, is obviously more difficult than that of objective phenomena such as mass, length and time. Provided that good instruments are used carefully, we can be confident that, for objective measurements, the results will be correct. With colour the situation is completely different. If a colour does not look to be correct, it is not correct, no matter what any instrument indicates. The final arbiter is the human eye.

One problem with regard to colour measurement is that the human eye is readily available and particularly sensitive to colour; millions of different colours can be distinguished. Any measurement that is less reliable than the unaided eye will be of limited value. Another consideration is that, with colour, we are often concerned with differences in colour rather than with colour itself. For many purposes the exact colour is not as important as uniformity of colour. When buying a blue shirt, the exact shade of blue is unlikely to matter, but any appreciable difference between the collar, sleeves and other parts would not be acceptable. Small differences between different shirts on display at the same time would give an impression of carelessness on the part of the manufacturer. Similarly, when buying paint, a customer will often accept any colour within reasonable range of the desired colour, but would be much less tolerant of even a small difference between the paint from two different cans. With colour, we should never forget that the final objective is to produce something that is pleasing or satisfactory to the eye. If the colour *looks* wrong, it *is* wrong.

When assessing the usefulness of the colour specification system to be described in this chapter, the reader should carefully consider how far the system enables us to deal with real problems with respect to colour and how far the system fails to deal with the subjective nature of colour. For the present we will simply examine how a system of colour specification can be set up. We will return to the question of how far such a system fulfils our needs at the end of the chapter. In considering the appearance of an object, factors such as texture and gloss are important, as well as colour. However, in this chapter these factors will be ignored.

Almost all modern colour measurement is based on the CIE system of colour specification. The initials come from the French title of the international committee (Commission Internationale de l′Eclairage) who set up the system in 1931. Although additions have been made since, the basic structure and principles are unchanged and the system is widely used. The reader should bear this in mind when determining how useful such a system is likely to be, while we consider the development of the CIE system in the next few sections. It should also be emphasised at this point that the system is empirical, i.e. it is based on experimental observations rather than on theories of colour vision.

When discussing colour in general, we could be considering coloured lights, coloured solutions or coloured surfaces such as painted surfaces, plastics and textiles. In most practical situations we are concerned with coloured surfaces, although, as we shall see, the properties of coloured lights are used in the specification of the colour of surfaces. It is important to realise that the colour of an object depends on the light source used to illuminate its surface, the particular observer who views it, as well as the properties of the surface itself. The nature of the surface is the most important factor. A piece of white paper will look white under all normal light sources when viewed by any observer with normal colour vision. (This statement is not completely true, for example small areas of 'white' paper viewed as part of a pattern formed by a variety of colours could look to be a different colour. In trying to understand the principles involved in colour measurement, it is probably best to ignore such exceptions for the moment and adopt a simplified view of colour. In practical applications, however, such factors must be considered whenever appropriate.) Because our white paper does remain white, there is a tendency to think of colour as a property of the surface. In considering the measurement of colour, the light source and the observer cannot be ignored; we should think about surface, light source and observer.

2.2 Basic facts

It would be very difficult to design a system of colour measurement that attempted to describe the colours that we see. We simply have to think how we might describe a colour. Colour names such as red, yellow, green and blue are reasonably well understood, but names such rose, salmon and cerise are less well standardised. What might be called rose by one person could be called pink by another. Even for red there is a considerable range of colours which any one person would accept as red. This is not necessarily the same range as for a second person. Strictly speaking, this discussion concerns the names which we give to colours rather than what people actually see. The latter is much more difficult to pin down, but there is no doubt

that different people do see colours differently. Even for one person, the appearance of a particular coloured surface can change as the circumstances (e.g. surrounding colour) change. The CIE system basically attempts to tell us how a colour might be reproduced (by a mixture of three primary light sources) rather than described. The amounts of the three primaries required to match a particular colour provide a numerical specification of that colour. A different colour would require different amounts of the primaries and hence the specification would be different. It turns out that, in many applications, this is all that is required. As we will see later, some idea of the colour seen can be deduced from CIE colour specifications; furthermore, we would never attempt to reproduce a colour by actually mixing the CIE primaries.

It is well known that colour is three-dimensional. This is apparent in various ways. Colour atlases such as the Munsell atlas arrange colours using three scales (hue, value, and chroma in the case of the Munsell system). We can match a wide range of colours using a mixture of three dyes, and the concentrations of three specific dyes required on a particular substrate to match a colour could be used to specify that colour. We would need to specify the light source under which the colour was seen, but the three concentrations would give a numerical specification of the colour; the specification would be different for different colours and it would be possible for someone else to reproduce the colour. Unfortunately, dyes or pigments tend to be impure and the precise colour depends on the method of manufacture. Even repeat batches from the same plant will not match exactly, and properties of mixtures are not completely predictable.

Coloured lights are much easier to define and reproduce. Imagine a red light obtained by isolating the wavelength 700 nm from the spectrum. All laboratories in the world capable of measuring wavelength accurately (an objective physical measurement) could produce the same red colour. A green colour corresponding to 546.1 nm could be produced even more easily. A mercury lamp emits light at only four wavelengths in the visible region (404.7 nm, 435.8 nm, 546.1 nm and 577.8 nm). By filtering out the other three, the required green wavelength could be obtained. The wavelengths 404.7 nm and 435.8 nm could be obtained in a similar manner. Small variations in the operating conditions have no significant effect on the wavelengths emitted by a mercury lamp. Hence, three primary light sources could be defined simply as appropriate wavelengths and be easily reproduced.

Mixtures of three coloured lights can be produced in various ways, but the simplest is to imagine three spotlights shining on the same area on a white screen. The colour produced would be a mixture of the three colours being mixed and it is possible to produce a wide range of colours by varying the amounts of the three primaries. The colours to be matched could include surface colours illuminated by a particular light source.

2.3 Additive and subtractive mixing

Most of us are familiar with the colours produced by mixing paints or coloured solutions. The results depend on the exact colours mixed, but usually

red + blue = purple

red + yellow = orange

yellow + blue = green

while

red + yellow + blue in the correct proportions = grey or black.

To people accustomed to mixing dyes or pigments, the colours produced by mixing lights might be surprising in some cases. For example, a blue light mixed with a yellow light might well give white, while red and green lights could be mixed to give a yellow. Quite obviously there is something fundamentally different about mixing dyes (and pigments) and mixing coloured lights. The latter is an example of additive mixing, while dye or paint mixture is an example of a subtractive mixing. Since the CIE system is based on additive mixtures of lights, these must be considered further.

Additive mixtures occur when two or more coloured lights are shone at the same time so that we see the two lights together. Consider red and green lights shone onto a white screen. The screen will reflect almost all the light incident upon it, and the mixture of red plus green in the same appropriate proportions will reach our eye. If the colour seen (yellow) is surprising, this is simply because we are not used to mixing colours in this way. Note that there is no way in which the two colours interact with each other. If the red and green are single wavelengths, both wavelengths reach our eye and are not interfered with in any way by each other. We see the red wavelength *plus* the green wavelength and hence call the mixture an *additive* mixture.

Subtractive mixtures occur much more commonly but usually involve much more complicated processes. The simplest case occurs when we shine light through two coloured glass filters placed one behind the other. The fraction of light transmitted by the combination of filters is simply the product of the fractions transmitted by the filters independently.

The most important examples of subtractive mixtures occur when we mix paints, or when we dye with a mixture of dyes. The results are often predictable on the basis of everyday experience, but the details of the process are very complicated (see Chapter 8). When we see the result of a subtractive mixture, we generally see light reflected at many different wavelengths. For example, a mixture of yellow and blue dyes will absorb some light at all

wavelengths, but mainly the blue and red wavelengths are absorbed and we see the remaining light, probably giving a green appearance. A yellow surface absorbs the short (blue and violet) wavelengths strongly and reflects most of the green, yellow and red wavelengths. The yellow section of the visible spectrum is very short, and what we see is basically an *additive* mixture of green and red, resulting in the yellow appearance.

2.3.1 Properties of additive mixtures of light

As indicated earlier, coloured lights are easy to define and hence are suitable for use as primaries in a system of colour specification. The properties of additive mixtures of coloured light have been studied for many years[1–5] and those that are particularly relevant to their use in systems of colour specification are considered here.

We can match a wide range of colours using an additive mixture of, say, red, green and blue primaries. Suppose our primaries are single wavelengths and we use them to match white light consisting of a mixture of all the wavelengths in the visible region. Although the mixture is physically quite different from the white light, by careful adjustment of the amounts of the primaries, the white light can be matched, i.e. we can produce, with the mixture, a white that looks identical to the white light. If we change the colour or lightness of the surround, our white colours change in appearance. However, over a wide range of conditions, the match holds, i.e. if the colours change, they both do so by the same amount. Hence we can deal with colours without considering their spectral composition, at least in many applications. Furthermore, if colour 1 matches colour 2, and colour 3 matches colour 4, then an additive mixture of colour 1 plus colour 3 will match colour 2 plus colour 4. This is vital when we consider that normal colours are additive mixtures of all the wavelengths in the visible spectrum. We need to consider the effect of the additive mixture of all the wavelengths.

It cannot be stressed too strongly that modern colorimetry is based on the properties of additive mixtures of coloured lights, and that these properties have been determined by experiment. The main properties were established well over a century ago. Subsequent work has confirmed that the simple properties described above are indeed true, but has defined much more closely the range of experimental conditions under which the simple laws hold.[6–8] Theories of colour vision must attempt to explain such laws, and the exceptions.

2.4 The CIE system of colour specification

If we were to select and define three particular primaries (R), (B) and (G), the amounts of these required to match any colour could be used to specify

the colour and be called tristimulus values. Each different colour would require different amounts, the tristimulus values would therefore define the colour and with practice we could deduce the appearance of the colour from the tristimulus values. However, different results would be obtained by anyone using a different set of primaries. It can be shown that sets of tristimulus values obtained using one set of primaries can be converted to the corresponding results that would have been obtained using a second set, provided that the amounts of one set of primaries required to match each primary of the second set of primaries in turn are known. Hence we could either insist that one set of primaries is always used, or allow the use of different sets, but insist that the results are converted to those which would have been obtained using a standard set.

Even using very pure colours for our set of primaries, there would still be some very pure colours that could not be matched. For example, a very pure cyan (blue–green) might be more saturated than the colours obtained by mixing the blue and green primaries. Adding the third primary (red) would produce even less saturated mixtures. A possible solution in this case would be to add some of the red primary to the pure cyan colour, and then match the resultant colour using the blue and green primaries.

Experimentally, it has been found that, if this procedure is followed, all colours can be matched using one set of primaries, the only restriction in the choice of primaries being that it must not be possible to match any one of the primaries using a mixture of the other two. Negative values are undesirable. It would be easy to omit the minus sign or fail to notice it. By careful choice of primaries it is possible to reduce the incidence of negative tristimulus values. The best primaries are red, green and blue spectrum colours. Mixtures of these give the widest possible range of colours. However, there is no set of real primary colours that can be used to match all colours using positive amounts of the primaries, i.e. there is no set of real primaries that will eliminate negative tristimulus values entirely. Since it is possible to calculate tristimulus values for one set of primaries from those obtained using a second set, there is no need to restrict ourselves to a set of real primary colours. Purely imaginary primaries can be used; it is only necessary that these have been defined in terms of the three real primaries being used to actually produce a match. This is not just a hypothetical possibility. Negative tristimulus values would be a nuisance in practice and, in the CIE system, imaginary primaries are indeed used so as to avoid negative values.

If we produced a visual match, the amounts of the primaries required could be noted, and the results converted to the equivalent values for a standard set of primaries. Such a procedure is perfectly possible, the main problem lying in the precision and accuracy achievable. The results would vary from one observer to another because of differences between eyes.

Even for one observer, repeat measurements would not be very satisfactory. Under the controlled conditions (usually one eye, small field of view and low level of illumination) necessary in such an instrumental arrangement, it is impossible to achieve the precision of unaided eyes under normal conditions, e.g. when judging whether an appreciable difference exists between two adjacent panels on a car body under good daylight. Some of the problems could be overcome by using more than three primaries[9], but such an arrangement has found little use in industrial applications.

The matches in the instrumental arrangement being considered are physically very different, and visual judgements are particularly unreliable for this type of match. The CIE realised even before 1931 that this was likely to be the case and decided that the system adopted should allow calculation of tristimulus values from measured reflectance values. The system was set up in 1931 because the required information (basically infomation defining a standard observer) became available at that time. It must be stressed that the CIE system incorporates the features described earlier, i.e. a colour is to be specified by the amounts of the (X), (Y), and (Z) primaries required, in an additive mixture, to match it. The 'standard observer' data is added to the framework already described.

It is possible to calculate tristimulus values (i.e. the amounts of three primaries which, if additively mixed, would match a colour) of a sample specified. The CIE had to define standard primaries, standard light sources and a standard observer, together with standard observing and viewing conditions.

2.4.1 Standard primaries

As discussed earlier, the CIE had considerable latitude in selecting three primaries. For the moment, the main property to note is that all real colours can be matched using positive amounts of the chosen primaries (X), (Y) and (Z). These were defined by the following equations:

$$C_{\lambda 1} \equiv 0.73467(X) + 0.26533(Y) + 0.00000(Z) \quad [2.1]$$

$$C_{\lambda 2} \equiv 0.27376(X) + 0.71741(Y) + 0.00883(Z) \quad [2.2]$$

$$C_{\lambda 3} \equiv 0.16658(X) + 0.00886(Y) + 0.82456(Z) \quad [2.3]$$

$$S_E \equiv 0.33333(X) + 0.33333(Y) + 0.33333(Z) \quad [2.4]$$

where $\lambda_1 = 700\,\text{nm}$, $\lambda_2 = 546.1\,\text{nm}$, $\lambda_3 = 435.8\,\text{nm}$ and S_E = equal-energy stimulus, i.e. a stimulus having equal amounts of energy at all wavelengths through the visible spectrum.

The relative sensitivity of the eye to light of different wavelength had been determined previously, and a particular curve (denoted by V_λ) adopted

as standard by the CIE in 1924. The CIE 1931 standard colorimetric system was made consistent with the 1924 V_λ curve by choosing primaries such that the $\bar{y}_\lambda$ curve (see Section 2.4.3) was identical to the V_λ curve.

2.4.2 Standard light sources and standard illuminants

The percentage (or fraction) of light at different wavelengths reflected from a coloured surface is independent of the amount of light incident on the surface, but the amount of light reflected is the amount incident multiplied by the fraction reflected. The amount of light reflected, and hence the appearance, depends on the light source. In practice, we use many different light sources, particularly various phases of daylight, and various types of fluorescent tube and tungsten light. However, if we wish to check that a paint sample is the correct colour, it would probably be satisfactory to check using one form of daylight, tungsten light and a fluorescent tube. Any sample that was satisfactory under all three conditions would almost certainly be satisfactory if seen under any other light source. Hence it was only necessary for the CIE to specify a small number of light sources rather than all possible sources. For standardisation purposes the minimum number is desirable. In 1931 fluorescent tubes were unimportant and the CIE specified three standard illuminants as follows. (The CIE distinguishes between sources and illuminants. A source refers to a physical emitter of light such as the sun or a lamp, while an illuminant refers to a specified spectral energy distribution. Thus an illuminant can readily be specified, but may not be realisable in practice. In calculating tristimulus values from reflectance values, the tabulated energy distribution is used, but may be different from the actual distribution of the light source in the spectrophotometer.)

CIE standard illuminant A

The spectral energy distribution of Illuminant A (plot of E vs. wavelength) represents a black body radiator at an absolute temperature of 2856 K. Source A can be realised by a gas-filled coiled tungsten filament lamp operating at a correlated colour temperature of 2856 K. The energy distributions of source A and illuminant A can be very close if a calibrated lamp is used.

CIE standard illuminants B and C

Illuminants B and C correspond to different phases of daylight; the former is intended to represent direct sunlight with a correlated colour temperature of 4874 K and the latter to represent average daylight with a correlated

colour temperature of 6774 K. The CIE gave details of how sources B and C may be obtained in the laboratory.[12] However, neither corresponds very closely to real daylight, particularly in the near UV region. The differences between the illuminants, e.g. A and C should be compared with the differences between the expected reflectance curves for different colours. Since the amount of light of any one wavelength reaching the eye is proportional to the energy of the source multiplied by the reflectance factor $E_\lambda R_\lambda$, it is clear that the amounts of any given wavelength of light reaching the eye from a single surface illuminated by two different sources can be quite different. Thus, considering all wavelengths of the visible spectrum, the tristimulus values for a surface under two different illuminants may well be very different, even though, after allowing time for adaptation, the colours seen under the two sources may look very similar.

2.4.3 Standard observer-colour matching functions

As explained above, we can calculate, for any light source, how much light is reflected at each wavelength throughout the visible region. The tristimulus values of any one wavelength are the amounts of the three chosen primaries required to match the light of the particular wavelength. If we know these tristimulus values we can calculate the tristimulus values for the sample. The amounts required depend on the observer and results for an average (or 'standard observer') are required. The values were actually determined experimentally as follows.

Wright[10] and Guild[11] used visual tristimulus colourimeters in which one-half of the field of view consisted of a mixture of (R), (G) and (B) primaries, while the colour in the other half was light of a single wavelength. To produce a match experimentally, it was necessary to add some of (R), (G) or (B) to the wavelength to be matched. This was quite possible experimentally, as discussed earlier, and the resultant tristimulus values for the wavelength included at least one negative value. Each used a somewhat different technique, and in particular different primaries were used. Both considered each wavelength throughout the visible spectrum and averaged results from a number of observers. The results differed from one observer to another (as expected), but when the average results from the two experiments were converted to a common set of primaries, the agreement was considered to be satisfactory. The results were expressed as the tristimulus values for an equal energy spectrum, i.e. using primaries (R), (G) and (B) the results were expressed as the amounts $\bar{r}$, $\bar{g}$ and $\bar{b}$ required to match one unit of energy of each wavelength throughout the visible region. Since (R), (G) and (B) were real primaries, some of the values were negative. The CIE adopted three unreal primaries (X), (Y) and (Z) and the colour matching functions in terms of these primaries are denoted by $\bar{x}$, $\bar{y}$ and $\bar{z}$

and are always positive. This ensures that tristimulus values for all real colours are always positive. The values are called the CIE colour-matching functions and they define the colour-matching properties of the CIE 1931 Standard Colorimetric Observer. Note that all values are positive. This results from the choice of primaries, as does the fact that $\bar{z}$ is zero for almost half of the spectrum.

2.4.4 Standard illumination and viewing conditions

The CIE specified that opaque samples should be illuminated at 45° from the normal to the specimen surface and viewed at an angle close to the normal; alternatively, the specimen should be illuminated at an angle close to the normal and viewed at an angle 45° to the normal.

2.4.5 Calculation of tristimulus values

The values of $\bar{x}_\lambda$, $\bar{y}_\lambda$ and $\bar{z}_\lambda$ are the amounts of the CIE primaries (X), (Y) and (Z), respectively, needed to match one unit of energy of the wavelength λ. For any sample, the amount of energy reflected from the surface, i.e. E_λ multiplied by R_λ. Hence the amounts of each of the CIE primaries needed to match this is $E_\lambda R_\lambda$ multiplied by $\bar{x}_\lambda$, $\bar{y}_\lambda$, or $\bar{z}_\lambda$, respectively. Adding the products for all the wavelengths gives us the amounts required to match the colour. These amounts are called tristimulus values and denoted by X, Y, and Z. Thus, for example, X is calculated from $X = \sum E_\lambda \bar{x}_\lambda R_\lambda$ where the greek letter sigma means 'sum' as in many other mathematical equations.

So far, the question of units has been avoided. Using the CIE system as described we could calculate the tristimulus values for a sample. However, the units used for E_λ and $\bar{x}_\lambda$, $\bar{y}_\lambda$, and $\bar{z}_\lambda$ are arbitrary. The values of E_λ for one wavelength relative to another are correct, but the absolute values have not been specified. Similarly, the $\bar{x}_\lambda$ values for one wavelength are correct relative to other $\bar{x}_\lambda$ values at other wavelengths, and correct relative to the $\bar{y}$ and $\bar{z}_\lambda$ values at the same wavelength; however, the absolute size of the values is arbitrary. For opaque samples (object colours), the usual practice is to normalise the tristimulus values by calculating using (2.5):

$$\begin{aligned} X &= \sum E_\lambda \bar{x}_\lambda R_\lambda \Big/ \sum E_\lambda \bar{y}_\lambda \\ Y &= \sum E_\lambda \bar{y}_\lambda R_\lambda \Big/ \sum E_\lambda \bar{y}_\lambda \\ Z &= \sum E_\lambda \bar{z}_\lambda R_\lambda \Big/ \sum E_\lambda \bar{y}_\lambda \end{aligned} \qquad [2.5]$$

Thus if R_λ is expressed as a percentage, Y runs from zero (for a sample which reflects no light) to 100 (for a sample which diffusely reflects all the

Table 2.1 Tristimulus values for a sample reflecting 100% of the incident light based on the CIE 1964 supplementary standard observer and weighting factors at 10 nm intervals

Illuminant	X	Y	Z
A	111.14	100.00	35.20
C	97.29	100.00	116.15
D_{65}	94.81	100.00	107.30

light incident on it) and is independent of any units used. The ranges for X and Z depend on the illuminant. For a sample reflecting all the light incident upon it, i.e. $R_\lambda = 100\%$ at all wavelengths, the X, Y and Z values for the sample under illuminants A, B, C and D_{65} are given in Table 2.1. It can be seen that the Z tristimulus values, in particular, vary greatly with the illuminant.

The fact that the tristimulus values of a sample take no account of the intensity of light incident on the sample causes no problems in normal practice. If, for example, we produce a sample with the same tristimulus values as a target (for a particular illuminant), the two will match for all normal levels of illumination by the chosen illuminant. The actual appearance of object colours is almost independent of the level of illumination. A piece of white paper looks white in weak sunlight. If the intensity of the sunlight increases, we recognise this, but the paper does not look any lighter. A medium grey is recognised as such whether viewed in daylight on a very dark day or on a very bright day. This is true even though the amount of light reflected by the grey sample on a bright day would probably be much more than that reflected by the white paper on a dark day. To a very large extent, the appearance of object colours seems to be judged relative to the light source. We separate the properties of the sample from those of the source. This applies to changes in the distribution of light throughout the spectrum as well as to changes in the intensity of the source. Our white paper will still look white under tungsten light even though the distribution reaching the eye is vastly different from that reaching the eye when the paper is seen in daylight.

In summary, the CIE defined a standard observer, standard illuminants, standard illuminating and viewing geometries and a particular set of primaries (X), (Y) and (Z).

The tristimulus values of a particular colour are the amounts of (X), (Y) and (Z) required to match the colour under standard conditions and are referred to as X, Y and Z. In practice, X, Y and Z for coloured materials

are calculated from measured reflectance values together with $\bar{x}$, $\bar{y}$, and $\bar{z}$ and E_λ for the chosen standard observers and illuminants. For the standard observer, an additive mixture of X units of the (X) primary together with Y units of the (Y) primary and Z units of the (Z) primary would look the same as the sample illuminated by the appropriate standard illuminant.

2.4.6 Additions to the CIE system since 1931

Since the CIE system of colour specification was adopted in 1931, the basic system has remained unchanged, but experience has led to additions being made.

D illuminants

Illuminants *B* and *C* were intended to represent different phases of daylight. Later measurements[13] have shown that neither represents any common phase of daylight at all closely, particularly in the near UV region. This latter point is particularly important when considering samples treated with fluorescent brightening agents.

Illuminant D_{65} is based on measurements of the total daylight (i.e. sun plus sky) in a number of countries. Except for times near sunrise and sunset, the relative spectral energy distribution generally corresponds to correlated colour temperatures between 6000 and 7000 K. If we consider illumination by only part of the sky (e.g. a portion of blue sky from a north-facing window or direct evening sunlight from a west-facing window), the correlated colour temperature and energy distribution can be quite different. Judd *et al.*[13] showed that the different relative energy distributions could be represented quite closely by a series of curves dependent only on the correlated colour temperature of the particular form of daylight. Based on Judd's work, the CIE has defined a series of *D* illuminants ranging from correlated colour temperatures of 4000 to 25 000 K. In the interests of standardisation, the CIE recommends that D_{65} should be used whenever possible.

There seems to be no doubt that the *D* illuminants represent a substantial improvement over illuminants *B* and *C*. Illuminant *B* has rarely been used and in recent years illuminant *C* has been almost completely replaced by the *D* illuminants, D_{65} (for textiles) and D_{50} (for graphic arts), in particular. However, there is a major problem with the *D* illuminants in that there is no satisfactory way of obtaining, say, D_{65} in the laboratory. Problems occur with metameric pairs. A pair that are a close match according to tristimulus values calculated for D_{65} might well be seen to be a poor match when viewed under a so-called daylight source in a viewing cabinet. However, the match might well look better if inspected under real daylight.

1964 supplementary standard observer (10°)

The original 1931 CIE standard observer was based on experiments using a 2° field of view. This is a much narrower field of view than that normally used for critical colour appraisal. In addition, a few problems were encountered using the 1931 observer.[14] New colour matching experiments were therefore carried out by Stiles and Burch[15] and by Speranskaya.[16] The experiments were similar to those by Wright and Guild for the 1931 standard observer in that each wavelength in turn was visually matched using an additive mixture of the primary light sources. The main difference was the much wider field of view (10°). As indicated earlier, the use of wide fields with metameric matches causes problems. These were largely overcome by ignoring the centre 2° of the field of view. The two sets of experimental results were combined and used to define the CIE 1964 Supplementary Standard Colorimetric Observer, often referred to as the 10° Observer. The subscript 10λ is used to distinguish the 10° data from the original 1931 2° standard observer data. It is recommended that the 1964 observer is used whenever a more accurate correlation with visual colour matching of fields greater than 4° is required. However, there is probably little to choose between the two standard observers for some applications.

Standard illuminating and viewing conditions

The original CIE recommendation was that the sample should be illuminated at 45° to the surface and the light viewed normally, i.e. at right angles to the surface. This mode can be represented 45/0. It was assumed that the opposite mode (0/45) would give the same result, but this is not the case if the incident light is polarized[17]. Four possible sets of conditions are now recommended. These are 45/0, 0/45, *d*/0 and 0/*d*. In the third case the sample is illuminated by diffuse light while in the last case the light reflected at all angles is collected (using an integrating sphere, as in many spectrophotometers).

2.4.7 Standard of reflectance factor

The CIE recommends that reflectance measurements should be made relative to a perfect diffuser, i.e. a sample that diffusely reflects all the light incident upon it. No such surface exists, but working standards of known spectral reflectance factors are normally used, allowing the correct results to be obtained. (For example, if the working standard reflects 98% of the light of a particular wavelength, all values measured relative to the working standard need to be multiplied by 100/98.) In practice, instrument manufacturers supply calibrated white tiles with their instruments. Using these, corrected *R* values are obtained automatically.

2.5 Calculation of tristimulus values from R_λ values measured at 10 or 20 nm intervals

The equations given above for calculating tristimulus values from reflectance values may look awesome to a beginner, but involve only simple additions and multiplications, and in any case these days are invariably carried out using computers generally attached to a spectrophotometer. However, there are problems. The CIE standards[20] tabulate $E_\lambda, \bar{x}_\lambda$, etc. at 5 nm intervals over the full visible region. Many spectrophotometers measure R_λ at 10 or 20 nm intervals, often over a restricted wavelength range, e.g. 400 nm, 420 nm, . . . , 700 nm. At one time it was considered sufficient simply to use the approptiate CIE values at the measured wavelengths, and to add the values at, e.g. 720 nm, 740 nm etc. to the values at 700 nm. It was gradually realised[18,19] that large and completely unnecessary errors could be introduced in this way. Substantial work has been carried out in attempts to minimise these errors.

Stearns has described a method that allows appropriate adjustments to the E_λ, $\bar{x}_\lambda$, etc. tables to be calculated for any wavelength interval. The results from using these adjusted tables are consistent with those obtained using 5 nm interval tables and interpolated R values where necessary, and give considerably improved results with no extra effort on the part of the user. However, even very small differences can cause problems, for example when checking computer programs, or when comparing results for the same sample measured by a supplier and his customer, and it is highly desirable that everyone uses the same set of values. In an attempt to promote uniformity of practice, the American Society for Testing and Materials (ASTM) has produced 10 nm and 20 nm tables for commonly used illuminants and for both standard observers. ASTM also give a method for calculating appropriate values for other illuminants. Some problems still remain.

For most users the values to be used are built in to the software supplied by the instrument supplier. However, manufacturers are understandably reluctant to change software. Improved sets of values or improved methods would lead to different results even for identical sets of reflectance values. Users can check whether appropriate sets are being used by inputting R values of 100 at all wavelengths and checking the results.

2.6 Relationships between tristimulus values and colour appearance

It is difficult to relate, in a simple way, the tristimulus values of a sample to the colour appearance. One reason is that the colour depends not only on the stimulus, but also on surrounding colours and the state of adaptation of the eye. Even ignoring such factors, the three-dimensional nature of colour

makes it difficult to determine relationships and it is usual to simplify any relationship involving colour by considering only one or two dimensions at a time. It was indicated earlier that, in choosing primaries for the CIE system, the wide range of possibilities had been used to produce certain desirable features in the final system. A consequence of ensuring that the $\bar{y}_\lambda$ curve corresponded to the V_λ curve was that the Y tristimulus value should roughly represent the lightness of a sample, i.e. the higher the Y value, the lighter the sample appears. The scale is far from uniform and caution should be exercised when comparing the lightness of quite different colours such as reds and greys. Nonetheless, in general it will be found that, the higher the Y value, the lighter a sample will look. Thus if $Y = 80$, we can be sure that the sample will appear light, while if $Y = 3$, the sample will look dark.

2.6.1 Chromaticity diagrams

To represent the other two dimensions of colour, it was usual to first define 'chromaticity coordinates' (x, y and z) and then plot y against x:

$$x = X/(X + Y + Z)$$
$$y = Y/(X + Y + Z)$$

and

$$z = Z/(X + Y + Z) \qquad [2.6]$$

From (2.6) it follows that $x + y + z = 1$ for all colours; it is therefore only necessary to quote two of the chromaticity coordinates and these can be plotted on a normal two-dimensional graph. It can also be shown that X and Z can easily be calculated from x, y and Y; hence the latter set is an acceptable form of specification, and consideration of Y values and plots of y against x should cover all possible colours. A plot of y against x is called a chromaticity diagram and, at one time, was widely used. It is described in detail in many textbooks on colour.

It was stated earlier that the Y scale is far from uniform. The same applies to the x y diagram; equal distances in the diagram do not correspond to equal visual differences. For a fixed difference in x and y, the difference seen would be much smaller for a pair of green samples than for pairs of blue or grey samples. This problem will be considered further in Chapter 4.

A further disadvantage of the x y diagram is that colours measured under different standard illuminants, e.g. Illuminant D65 and Illuminant A are represented by completely different points in the diagram. It has been

emphasised that colour is three-dimensional. Thus any two-dimensional plot cannot represent colour completely. In the case of the chromaticity diagram, it is simplest to regard the missing factor as the *Y* tristimulus value. In general, any one point on any chromaticity diagram corresponds to a range of colours differing in lightness, and this should always be kept in mind when trying to visualise the colours corresponding to particular chromaticity coordinates.

2.7 Usefulness and limitations of the CIE system

In many ways the CIE system of colour specification has been remarkably successful. Almost all important applications of colour measurement use the CIE system. The basic system has survived unchanged for over 70 years. The additions made since 1931 have led to improvements in some respects, but have not changed the basic principles of the system in any way. It is unlikely that any major changes will be made in the foreseeable future. On the other hand, we must consider the limitations of the system. These stem basically from the limited objectives of the system rather than from a failure to meet the objectives.

The CIE tristimulus values for a sample are related to the colour of the sample, but ignore other important features such as surface texture, gloss, sheen, etc. Thus a gloss paint sample and a matt paint sample might have the same tristimulus values, but obviously will not look the same. Whether the colours of the two samples will look the same depends critically on the geometrical arrangements for illuminating and viewing the samples. Only if the instrumental geometry of illumination and viewing conditions are close to that used visually will the colours seen be close. An instrument will always average out the light reflected from the area being measured (typically a 2 cm diameter circle). In judging a colour visually some sort of averaging takes place, but the observer is always conscious of any non-uniformity over the area viewed. Thus a matt paint surface, a woven textile surface and a pile fabric will always look different from one another, but their measured tristimulus values could be the same.

Ignoring all features other than colour, the tristimulus values for a sample give only a limited amount of information. Basically, the tristimulus values tell us the amounts of three imaginary primaries, which if additively mixed will give the same colour as a surface illuminated by a standard source and viewed by a standard observer using one of the standard geometries. It follows that the mixture of the CIE primaries would be unlikely to match the surface if the latter was illuminated by a different source or if the 'match' was viewed by an individual observer or if a different illuminating or viewing geometry was used. Control can be exercised (to some extent) over the source and the geometry. If these are important, we must try to

ensure that the instrumental conditions correspond as closely as possible to those to be used when viewing the object visually. The only choice as far as the standard observer is concerned is whether to use the 1931 (2°) observer or the 1964 (10°) observer. Neither is likely to correspond closely to any individual observer. However, the standard observers may well correspond reasonably closely with the average judgement of real observers, bearing in mind that, in many applications, the product is mass produced and will be seen by many different observers.

A full specification of a colour requires X, Y and Z values (or equivalent sets such as x, y and Y or L^*, a^*, and b^*) for several different illuminants. The results are still only valid for the standard observer and could be unsatisfactory for a real observer. This should not be a problem in practice since we usually require colours to be acceptable to a large number of potential customers; the standard observer is probably a better guide to the population in general than any one observer. However, situations often do arise where one particular individual (e.g. a head dyer or a buyer for a chain store) inspects material and problems will occur when that individual's colour vision is appreciably different from that of the standard observer. Problems will be much more severe for highly metameric pairs of samples.

Strictly speaking, the tristimulus values tell us nothing about the colour of a sample although, as discussed above, with experience we can make a reasonable estimate of the colour from either X, Y and Z or x, y and Y values. In such cases it is essential that the illuminant used for the measurements is known. Chromaticity coordinates of $x = 0.314$ and $y = 0.331$ correspond to a neutral colour if derived from measurements under illuminant D_{65}, but if derived from measurements under illuminant A will correspond to a blue colour.

In many applications the aim is to match a particular target, which might be defined by a set of tristimulus values. If we produce a sample and wish to test whether this matches the target, it is essential that sample and target measurements correspond to exactly the same conditions (illuminant, standard observer, illuminating and viewing geometries, and, in practice, the same instrument). If, for example, the sample is a really good match to the target, but the tristimulus values are to be measured using different standard observers for the sample and target, the resulting tristimulus values and chromaticity coordinates would be appreciably different. Again, the sample and target might have different surface structures, such as those of matt paints, gloss paints or pile carpets. In these cases the tristimulus values could be identical but the surfaces would look appreciably different. Whenever possible, sample and target should have the same surface structure.

It is usually important that the colours within one batch and between repeat batches of material match closely. In these cases the samples will be

of the same material, and the same dyes or pigments will have been used. It would be natural to measure all the samples using the same instrument and (if a spectrophotometer was used) to calculate the tristimulus values for the same standard observer and standard illuminant. Under these conditions, if the tristimulus values for a sample are very close to those for the standard, the sample will be a close visual match to the standard for any normal observer viewing under a light source roughly equivalent to the standard illuminant used to calculate the tristimulus values. (Exceptions can occur, for example, with metallic paints, where the appearance depends very much on the illuminating and viewing geometries.) If the variation of appearance with, say, viewing angle is different for the sample and for the standard, they may not match visually, although according to the instrumental results (obtained with a different viewing angle) they should be a good match.

In many ways the chief limitation of the CIE system is its non-uniformity. Equal changes in x, y or Y do not correspond to the same perceived difference. Many attempts have been made to provide a more uniform system. In each case the basic approach has been to start with the tristimulus values or chromaticity coordinates from the CIE system, and to transform these in some way to give a more uniform system. The end result is a colour-difference formula which, for a pair of samples, gives a number that is intended to be proportional to the difference seen (see Chapter 4).

2.8 Colour order systems

It is often useful to have real physical samples to represent colours. A designer can use such a sample to see the effect of the colour alongside other colours and, when satisfied, to pass the sample to a dyer to represent the colour required. The sample may be a piece of coloured paper or dyed fabric, etc. The main advantage of this is that designer and dyer can both see the colour. One limitation is that such samples are not stable – the colour changes with time, particularly if the sample is handled often and becomes soiled. If both designer and dyer have sets of the same collection, there is no guarantee that the two sets are identical, even when first purchased. A further problem is that, even with collections of more than a thousand samples, it is often found that the required colour is not present in the collection. This problem can be minimised by using a colour order system, in which the samples are arranged in an ordered fashion, enabling interpolation to be made between adjacent samples. Such systems have other advantages: they can be used for designers and students to understand colours and the relationships between them. The most common such system is the Munsell system.

In the Munsell system the samples are arranged along three scales, hue, value and chroma. Hue is the attribute by which we distinguish between red, yellow, green, etc. In the Munsell book, the samples are arranged so that all the samples on one page are the same hue as judged visually. Thus on a page of orange colours none should look to be redder or yellower than the rest. There are 40 such pages, 2.5*YR*, 5*YR*, 7.5*YR*, 10*YR*, 2.5*Y*, and so on, round to 2.5*YR* again. The Munsell value is roughly what is commonly called lightness. On one page of the *Munsell Atlas*, all the samples on one row have the same value, i.e. none appears lighter or darker than the rest. Rows higher up the page are lighter. The value scale runs from 0 (for black) to 10 (for white). The remaining scale, chroma, represents colourfulness or distance from neutral. The samples on one row vary in Munsell Chroma and range from neutral (grey) to the most colourful (or saturated) colour that can be produced without changing the hue or value. The chroma scale runs from 0 (for neutral) to a maximum of up to 14 depending on the samples that can physically be produced. One particular feature of this system is that, for each scale, the differences between neighbouring samples are visually all the same, as far as possible. Since the book was first produced, alterations have been made to improve the spacing. Unusually, the samples are also represented by their CIE specifications. Great care is taken to try to ensure that the samples are stable, and that each set is the same as all the others. Sets are available in matt and gloss paint.

Other colour order systems all involve three dimensions, but the scales are somewhat different. All have the advantage that the colours can be seen, but the accuracy of colour representation is somewhat limited.

The Natural Colour System (NCS) is based on the concept that there are four unique hues: red, yellow, green and blue and that these together with white and black make six basic colours. Colours are represented by the relative amounts of these basic colours perceived to be present. Thus a pure orange might be perceived as 50% each of red and yellow, while a dark brown might be 50% black and 25% red and yellow. Again, the samples on one page of the *NCS Atlas* are the same hue, but the other samples on the page are represented in a equilateral triangular arrangement, black to white again forming the vertical axis, with the purest possible representation of the hue at the other apex of the triangle.

2.9 Colour specifiers

Various sets of specifiers are available commercially, often with a good selection of colours, but are generally not arranged according to any particular pattern. It is particularly useful if the samples are on the same type

of material as the samples to be produced. It is easier to match a wool sample to another wool sample rather than to a gloss paint sample.

2.10 Future trends

The CIE system is well established and it is unlikely that any major changes will be made in the foreseeable future. Some additions are possible. The D illuminants are good representations of daylight, but are almost impossible to reproduce accurately in colour measuring instruments and colour matching cabinets. A major problem lies in the relatively large amount of UV radiation. This leads to problems with materials treated with brightening agents (almost all textiles and papers). Many of these are generally seen indoors, when daylight has passed through glass, hence reducing the UV content. There has been a suggestion that a standard illuminant representing indoor daylight would be more applicable in many cases and easier to reproduce in practice. Details of the procedures for calculating tristimulus values from reflectance values will be tidied up.

With respect to colour order systems, an enormous amount of effort has been put into systems such as Munsell and NCS and, while alternative arrangements are possible, it is unlikely that the effort required will be undertaken. Alternative colour specifier systems will quite possibly be produced for particular applications.

2.11 References

1. Grassman, H. (1853). Zur theorie der Farbenmischung. *Poggendorf's Annalen Physik Chemie* **89**, 69–84.
2. Grassman, H. (1853). Zur theorie der Farbenmischung. *Phil. Mag.* **7** (4), 254.
3. Blottiau, F. (1947). Les défauts d'additivité de la colorimétric thrichromatique. *Rev. D'Opt.* **26**, 193.
4. Trezona, P.W. (1953). Additivity of colour equations. *Proc. Phys. Soc. (London)*, **B66**, 548.
5. Trezona, P.W. (1954). Additivity of colour equations. *Proc. Phys. Soc. (London)*, **B67**, 513.
6. Brindley, G.S. (1953). The effects on colour vision of adaptation to very bright lights. *J. Physiol.*, **122**(28), 332–350.
7. Trezona, P.W. (1972). Colour metrics, 36, Soesterberg: AIC/Holland.
8. Stiles, W.S. and Wyszecki, G. (1973). Rod intrusion in large-field color matching. *Acta Chromatica*, **3**, 155–163.
9. Donaldson, R. (1947). A colorimeter with six matching stimuli. *Proc. Phys. Soc. (London)*, **59**, 554–560.
10. Wright, W.D. (1928–1929). A re-determination of the trichometric co-efficients of the spectral colours. *Trans. Opt. Soc. (London)*, **30**, 141–164.
11. Guild, J. (1931). The colorimetric properties of the spectrum. *Phil. Trans. Roy. Soc. (London)*, **A 230**, 149–187.

12. Davis, R. and Gibson, K.S. (1931). Miscellaneous Publication 114. Washington, DC: National Bureau of Standards.
13. Judd, D.B., MacAdam, D.L. and Wyszecki, G. (1964). Spectral distribution of typical daylight as a function of correlated color temperature. *J. Opt. Soc. Am.* **54**, 1031–1040.
14. Judd, D.B. (1949). A comparison of direct colorimetry of titanium pigments with their indirect colorimetry based on spectrophotometry and a standard observer. *J. Opt. Soc. Am.* **39**, 945.
15. Stiles, W.S. and Burch, J.M. (1958). NPL colour-matching investigation: final report. *Optica Acta*, **6**, 1–26.
16. Speronskaya, N.I. (1959). Determination of spectrum color co-ordinates for twenty-seven normal observers. *Optics and Spectroscopy* **7**, 424–428 (USSR), English translation.
17. Johnson, S. *et al*. (1963). *JSDC*, **79**, 731.
18. Foster, W.H. Jr *et al*. (1970). Weights for calculation of tristimulus values from 16 reflectance values. *Colour Eng.* **8**, 35.
19. Stearns, E.I. (1975). Weights for calculation of tristimulus values. *Clemson Rev. Ind. Managem. Tex. Sci.* **14**, 79–113.
20. Supplement No. 2 to CIE Publication No. 15 *Colorimetry* (E1.3.1) (TC1.3) Paris: Bureau Central de la CIE, 1978.
21. Newhall, S.M., Nickerson, D. and Judd, D.B. (1943). Final report of the O. S. A. Subcommittee on Spacing of the Munsell Colors. *J. Opt. Soc. Amer.* **33**, 385.