4

Specifying, generating, and measuring colours

J. D. MOLLON

4.1 Introduction

In setting up an experiment on colour, the experimenter faces a fundamental choice: should the coloured stimuli be specified in physical terms or only in terms of the colour perception of a standard human observer? A specification of the former kind is called 'spectroradiometric', whereas a specification in terms of the human eye is called 'colorimetric'.

A *spectroradiometric* specification of a stimulus gives a 'spectral power distribution', that is a description of how radiance varies with wavelength over the range of the visible spectrum. However, normal human colour vision depends on the relative rates of quantum catch in only three classes of retinal cone, and so a normal observer can make a colour match by adjusting just three primary lights. Thus any spectral stimulus can be specified by three variables. Typically this *colorimetric* specification is given in terms of an internationally agreed luminance value and two *chromaticity coordinates* (see below). For many purposes (e.g. an experiment on visual search where colour is one of several attributes that distinguish the stimuli), a colorimetric specification is adequate; and it is clearly less clumsy than a full spectral power distribution.

However, if the subjects are unlikely to have the spectral sensitivity of the standard observer, if, for example, they are animals or are human observers with anomalous colour vision, then a conventional colorimetric specification is inappropriate. For the experiment to be replicable, the experimenter must in such cases use monochromatic stimuli of specified radiances and wavelengths or, if the stimulus is broad-band, must report its full spectral power distribution. One new factor that favours a spectroradiometric specification is the polymorphism of normal human colour vision: the colour matches of individual subjects may differ considerably from those of the standard colorimetric observer, and it is now known that much of the variance in such matches derives from inherited differences in the nucleotide sequences of the genes that code for the long- and middle-wave photopigments of the retina (Winderickx *et al.*, 1992). A radiometric specification may also recommend itself when mesopic conditions are being studied, that is when the subject's responses depend on the rods as well as the three types of cone: for two lights that have the same chromaticity and photopic luminance may differ greatly in their effects on the rods, and so a colorimetric specification would be inadequate.

In this chapter we discuss how to specify, generate, and measure coloured stimuli. Where appropriate we consider separately the spectroradiometric and colorimetric approaches. In experimental practice, of course, the two approaches are often mixed and many authors, although making their calibrations spectroradiometrically, report their results in terms of visual mechanisms, e.g. the contrast visible to an individual class of cones (Stromeyer *et al.*, 1985).

4.2 How to specify coloured stimuli

4.2.1 Spectroradiometric specification

If the stimuli are monochromatic (of bandwidth less than, say, 10 nm), a radiometric characterization should give the wavelength, bandwidth, and radiance of the stimuli. If the stimuli are broad-band, then the spectral power distribution should be reported, that is the variation of radiance with wavelength.

4.2.1.1 Spectral limits

In the case of man, the visible spectrum can conventionally be taken to extend from 400 to 700 nm.[1] However, in the case of insects, birds, fish, and even some mammals (Jacobs *et al.*, 1991), sensitivity may extend much further into the ultraviolet or infrared, and the human spectral limits would then be inappropriate. Photopigments with peak sensitivities as low as 360 nm are proving to be quite common.

4.2.1.2 Wavelength vs. wavenumber

For a time it was fashionable amongst colour scientists, especially those influenced by W. S. Stiles, to specify not the wavelength of visible radiation but its frequency or wavenumber. One valid ground for giving actual frequency is that frequency, unlike wavelength (or wavenumber), is independent of the optical medium; but people were probably most influenced by the belief that all photopigments had similar absorption spectra when the spectra were plotted on a frequency abscissa. This is now known to be quite wrong at the level of the photoreceptors (Barlow, 1982) (Mansfield, 1985). In fact, *log* frequency gives the most constant shape for pigment spectra (Baylor *et al.*, 1987), but wavelength is nowadays again the conventional abscissa except in specialist discussions of photopigments.

4.2.1.3 Quanta vs. energy

The energy of a light quantum varies with its frequency. The variable actually thought to determine the response of a cone cell is the total number of quanta absorbed: although the frequency of a given quantum will determine the *probability* that it will be absorbed by the photopigment, no information about the energy or frequency of the quantum is thought to be preserved in the physiological signal. For this reason many colour scientists prefer to specify the radiance of monochromatic lights in quanta rather than in energy units. The energy per quantum is given by the relationship $\varepsilon = hc/\lambda$, where h is Planck's constant (6.626×10^{-34} J s), c is the speed of light (3×10^{8} m s^{-1}), and λ is wavelength expressed in metres. So in laboratory practice the relationship that is most often needed is:

$$N = \lambda \times 5.03 \times 10^{15};$$

where N is the number of quanta per joule, and wavelength (λ) is now expressed in nanometres. Since a watt is a joule per second, one watt of monochromatic radiation at λ corresponds to N quanta per second.

4.2.2 Colorimetric specification

It would seem natural to have a way of expressing colours in terms of the relative excitations of the three cones of the normal retina, and the MacLeod–Boynton diagram, discussed below in Section 4.2.2.2, offers such a representation; but in 1931, when the Commission Internationale de l'Eclairage (CIE), in response to commercial and scientific needs, adopted a system of specifying colours in terms of three variables, the spectral sensitivities of the actual photoreceptors were not securely known. What were available were colour-matching functions, derived from the careful experiments of Guild and of Wright: these allowed any stimulus to be specified in terms of the proportions of three (arbitrary) primary lights needed to match it and they became the basis of the CIE 1931 chromaticity diagram (Wyszecki and Stiles, 1967; 1982).

4.2.2.1 The CIE system

Already in 1972 William Rushton was regretting that the CIE system was too often used 'to instruct the young and bewilder the old'; but the system is well understood by the trades that need to communicate about colours, it has useful features, and it is not really all that difficult to understand. Any light can be described in terms of three 'tristimulus values', X, Y, Z, which could be thought of as the photon catches of three imaginary photodetectors.[2] Their 'distribution coefficients', or relative sensitivities for an equal energy spectrum, are shown in Fig. 4.1. Although these distribution coefficients (designated $\bar{x}$, $\bar{y}$, $\bar{z}$) are not the human cone sensitivities, they are approximately linear transformations of the cone sensitivities of an average observer ($\bar{z}$ is in fact very close to the spectral sensitivity of the short-wave cones). Y was chosen to have the same spectral sensitivity as V_λ, which had already been adopted as the standard function for the luminous sensitivity of the eye (see Chapter 1); so $\bar{y}$ has a maximum value of 1.0 at 555 nm.

To specify the colour of a stimulus in the CIE system, one first finds the values of X, Y, Z by multiplying the power spectrum $P(\lambda)$ of the stimulus with $\bar{x}$, $\bar{y}$, and $\bar{z}$ in turn:

$$X = K \int P(\lambda) \times \bar{x}\ (\lambda)\ \mathrm{d}\lambda;$$

$$Y = K \int P(\lambda) \times \bar{y}\ (\lambda)\ \mathrm{d}\lambda;$$

$$Z = K \int P(\lambda) \times \bar{z}\ (\lambda)\ \mathrm{d}\lambda;$$

where K is a scaling constant. One then gives the corresponding 'chromaticity coordinates' x, y, where:

$$x = X/(X + Y + Z), \text{ and}$$

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

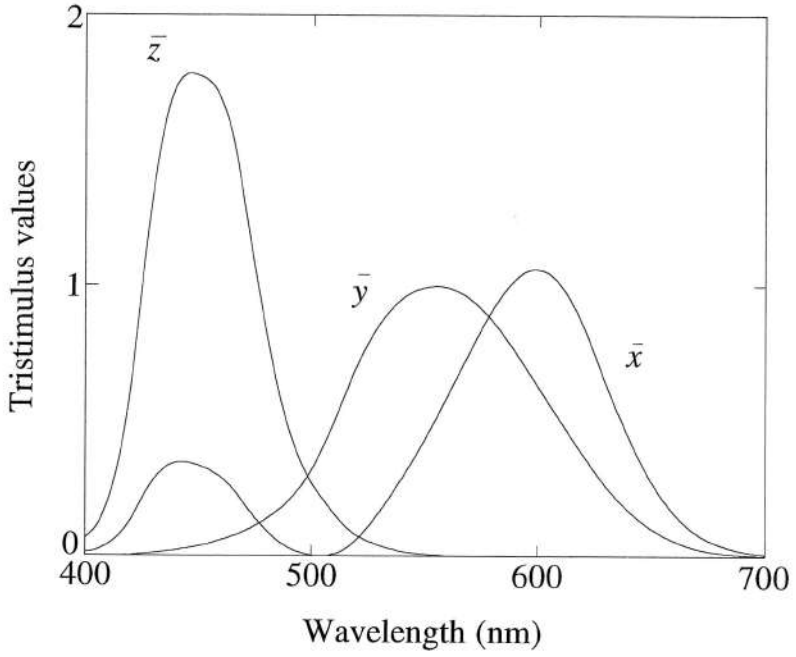

Fig. 4.1 The variation with wavelength of $\bar{x}$, $\bar{y}$, $\bar{z}$ in the system introduced in 1931 by the Commission International de l'Éclairage (CIE). Formally, the curves show the amount of each of three imaginary primary lights required to match the colour of one unit of radiant power of the indicated wavelength; but informally, it may be helpful to think of them as the sensitivity curves of three photodetectors. They are not the sensitivities of the human cones but are linearly related to those of the cones—or at least, the cones of some average observer.

Figure 4.2 shows the CIE chromaticity diagram, which has x and y as its axes.[3] The locus of monochromatic lights is shown as a solid line and the chromaticities of some common stimuli are indicated. To think physiologically about the CIE diagram, it is useful to know that lines radiating from $x = 1$, $y = 0$ correspond to loci of equal excitation of the short-wave cones when luminance is held constant (Fig. 4.3). It is also useful to know that lines radiating from approximately $x = 0.171$, $y = 0$ correspond to sets of chromaticities that are are confused by a tritanope (someone lacking the short-wave cones) and thus are chromaticities for which the ratio of long- and middle-wave cone signals is constant.

4.2.2.2 MacLeod–Boynton diagram

The CIE diagram is now so entrenched in commercial practice that it will not be quickly displaced. However, many visual scientists are adopting a physiologically realistic alternative, the chromaticity diagram of MacLeod and Boynton (1979), shown in Fig. 4.4. The vertical ordinate of the diagram corresponds to the relative excitation of the short-wave cones, expressed as a proportion of its maximum value near 400 nm; the abscissa represents the proportion of the total luminance contributed by the long-wave cones. In this diagram the results of colour mixing can be simply predicted by applying a centre-of-gravity rule using the luminances of the lights as weightings—in contrast to the CIE diagram where the weightings must be in terms of the sum of X, Y, and Z.

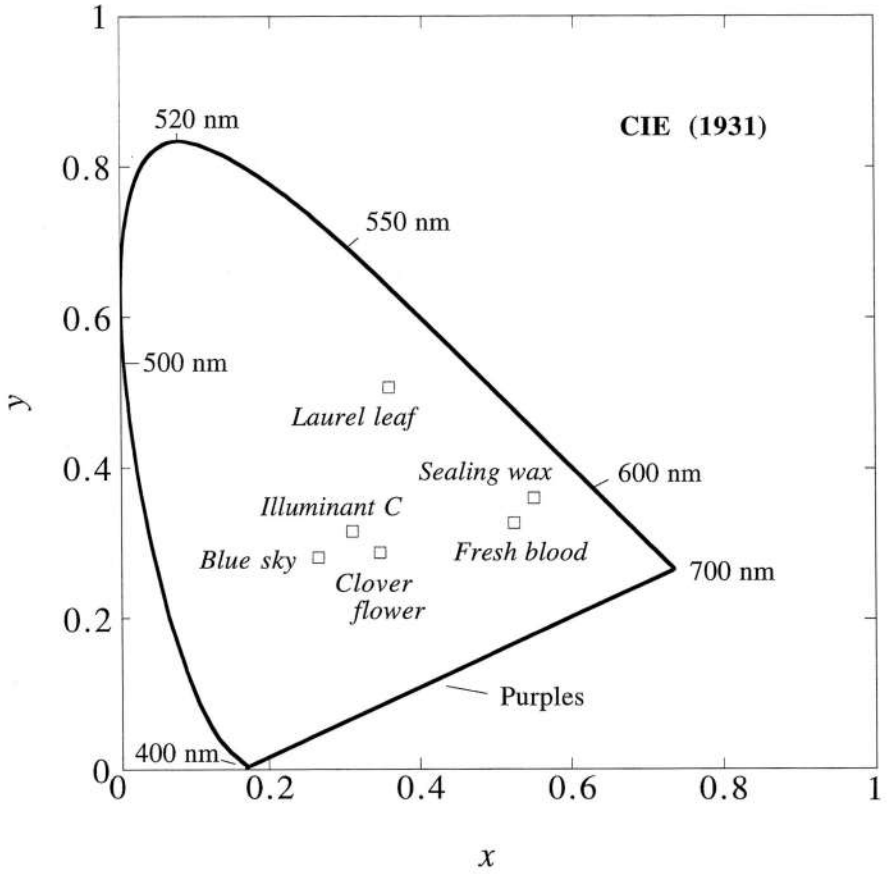

Fig. 4.2 The CIE (1931) chromaticity diagram. The solid line represents the spectrum of monochromatic lights from 400 to 700 nm plus the purple lights that consist of mixtures of long- and short-wave light. The chromaticities of all other physically realizable stimuli lie within this boundary. Coordinates are plotted for some illustrative stimuli. In the case of the reflective materials, the values shown are ones measured in Illuminant C, an approximation to northern daylight (see Section 4.3.8). Notice that quite strongly coloured stimuli may lie well within the interior of the diagram.

The MacLeod–Boynton diagram recommends itself because its axes correspond to the two chromatic channels that have been identified physiologically in the early visual system (Derrington *et al.*, 1984). The vertical axis corresponds to the phylogenetically older subsystem, which compares the signal of the short-wave cones with some combination of the signals of the long- and middle-wave cones. The abscissa corresponds to the phylogenetically recent subsystem, which compares the quantum catches of the long- and middle-wave cones (Mollon and Jordan, 1988). Horizontal lines in the diagram represent lights that give equal excitation in the short-wave cones. Vertical lines are tritan confusion lines, i.e. sets of chromaticities that represent a constant ratio of long-wave to middle-wave cone excitation. So the two physiologically significant sets of converging loci in the CIE diagram (identified in Section 4.2.2.1 and Fig. 4.3) become two sets of parallel loci in the MacLeod–Boynton diagram.

The MacLeod–Boynton diagram is derived from the cone sensitivities given by Smith and Pokorny (1975). The latter are not exact linear transformations of the CIE $\bar{x}$, $\bar{y}$, and $\bar{z}$ functions but rather are derived from the slightly different 1951 $\bar{x}_j$, $\bar{y}_j$, $\bar{z}_j$ functions of Judd (tabulated by Vos, 1978). The functions of Judd are preferred for work in visual science in that they incorporate a more accurate estimate of luminosity at short wavelengths. Vos (1978) gives a formula for transforming between chromat-

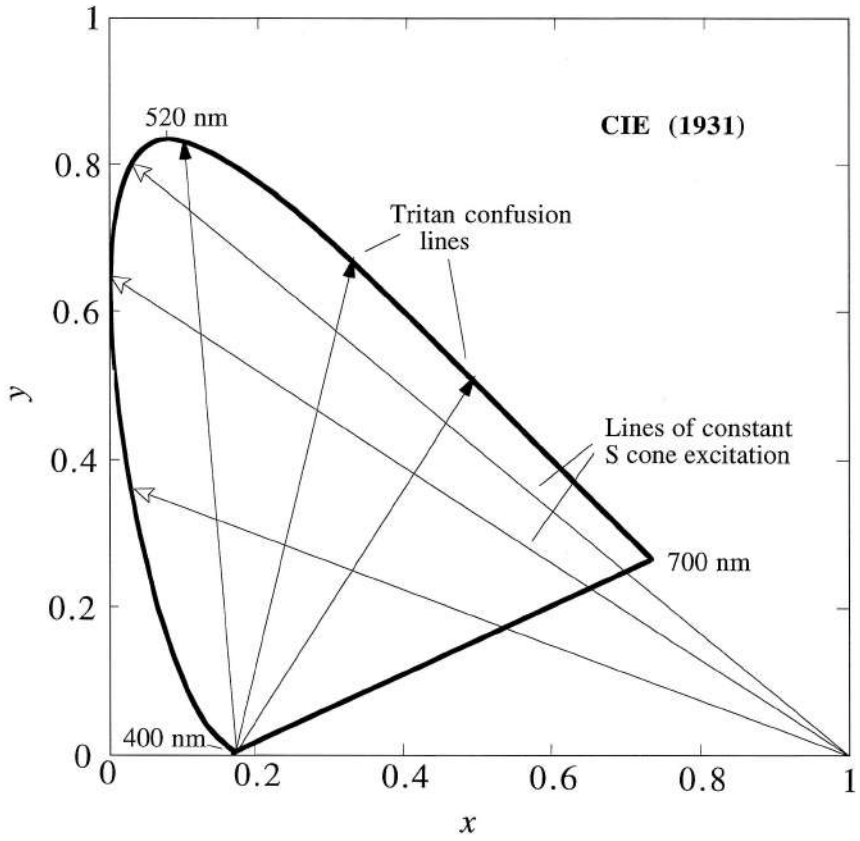

Fig. 4.3 The CIE (1931) chromaticity diagram showing some physiologically significant features. Any line radiating from $x = 1.0$, $y = 0$ (examples are shown with open arrow heads) represents a set of stimuli for which the excitation of the short-wave cones is constant for lights of constant luminance. A special case of such a line is the diagonal running from $x = 1.0$, $y = 0$ to $x = 0$, $y = 1.0$, which corresponds to zero excitation of the short-wave cones; between 700 nm and approximately 550 nm, this line effectively coincides with the spectrum locus, indicating that the standard observer is dichromatic in this part of the spectrum. Any line radiating from $x = 0.171$, $y = 0$ (examples are shown with solid arrow heads) is a tritan confusion line: along such a line the ratio of excitation of long- and middle-wave cones is constant and such stimuli would be confounded by a tritanope (i.e. a subject lacking short-wave cones). These relationships are shown here only to illustrate properties of the CIE chromaticity diagram: for research purposes it will often be desirable to establish them experimentally for a particular subject.

icity coordinates in the CIE (1931) and Judd (1951) systems, but this transformation is valid only for monochromatic lights. There is no unique transformation between the two systems for non-monochromatic lights. Thus, to use the MacLeod–Boynton diagram you strictly must know the spectroradiometric properties of your stimulus; you cannot simply measure the light with a colorimetric instrument that gives CIE (1931) chromaticity and luminance.

4.2.2.3 The limitations of chromaticity diagrams

To know the chromaticity coordinates of a stimulus is equivalent to knowing the relative signals it will produce in the three cone types of a (standard) human observer, but the chromaticity coordinates do not tell us how the stimulus will *look*: by manipulating simultaneous and successive contrast, we can generate almost any hue and any brightness from a given luminance and chromaticity.

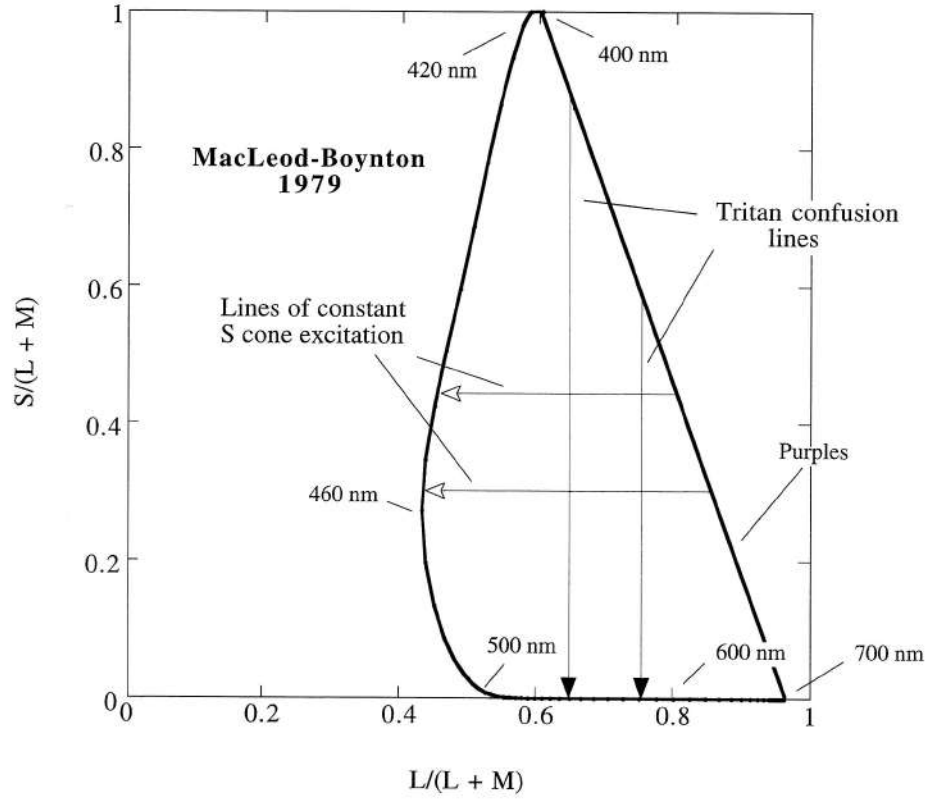

Fig. 4.4 The chromaticity diagram introduced by MacLeod and Boynton (1979). The abscissa represents the relative excitation of the long- (L) and middle-wave (M) cones: it is scaled to indicate the proportion of the total luminance that is accounted for by photons absorbed in the L cones. The ordinate represents the excitation of the short-wave (S) cones, arbitrarily scaled to have a value of 1.0 at the wavelength at which it is maximum at constant luminance. As in Figs 4.2 and 4.3, the solid line represents the coordinates of monochromatic spectral lights, plus the purples that consist of mixtures of short- and long-wave light. Also shown, using the same conventions as in Fig. 4.3, are examples of tritan confusion lines and lines of constant short-wave cone excitation. In the MacLeod–Boynton diagram, these two sets of physiologically significant lines become orthogonal.

Nor does distance in a chromaticity diagram tells us how salient will be the difference in colour *between* two stimuli. The philosophers' stone of colour science has been a *uniform colour space*, that is a representation of all chromaticities such that stimuli equally separated in the space are equally distinct. It is clear today that a space of this kind can be valid only for the stimulus conditions under which it was generated: its structure will vary with the spatial parameters of a display (Regan and Mollon, 1997), with the size of the colour differences (Mollon and Cavonius, 1986), and with the range of colours in the array (Webster and Mollon, 1995). Nevertheless, the CIE (1976) u', v' uniform chromaticity diagram, a linear transformation of the x,y diagram, has an honest use in choosing, say, chromaticities that are roughly equally separated in an experiment on visual search or in choosing a starting set of constants in a discrimination task intended for patients with varying degrees of colour loss. The relationship between the two diagrams is given by the following equations:

$$u' = 4x/(-2x + 12y + 3) \qquad v' = 9y/(-2x + 12y + 3)$$

$$x = 9u'(6u' - 16v' + 12) \qquad y = 4v'/(6u' - 16v' + 12).$$

A rough working rule, for use only in the privacy of the laboratory, is to reduce the factor 9 in the formula for v' if the stimuli are brief (less than 200 ms) or small (less than 0.6 deg). These are the conditions under which the eye tends to tritanopia, and compressing the vertical ordinate of the u', v' diagram is a rough way of allowing for this.

4.2.2.4 Materials

Tables of chromaticity coordinates for the CIE and MacLeod–Boynton diagrams are available from a web site: *http://www-cvrl.ucsd.edu/*. Blocks of graph paper overprinted with the CIE chromaticity diagram are available from Beuth Verlag (see the Appendix at the end of this chapter for this and other addresses).

4.3 Generating coloured stimuli

4.3.1 Monochromators

The monochromator remains the workhorse of the colour laboratory. In almost all modern monochromators a diffraction grating is used to form a spectrum, and the desired wavelength is selected by rotating the grating until the required wavelength coincides with the exit slit of the instrument (Fig. 4.5). The surface of a diffraction grating consists of evenly spaced, parallel grooves. When white light falls obliquely on the grating, each groove can be thought of as a thin individual source. Consider light that is diffracted in a given direction from one groove. A wavelength will exist such that light diffracted in the same direction from the next groove has to travel a path that is longer by that wavelength, or by an integral multiple of that wavelength. The light of this wavelength will thus be in phase and will constructively interfere. Other wavelengths, out of phase, will destructively interfere. The favoured wavelength is different for different directions of diffracted light, and thus the spectrum is formed.

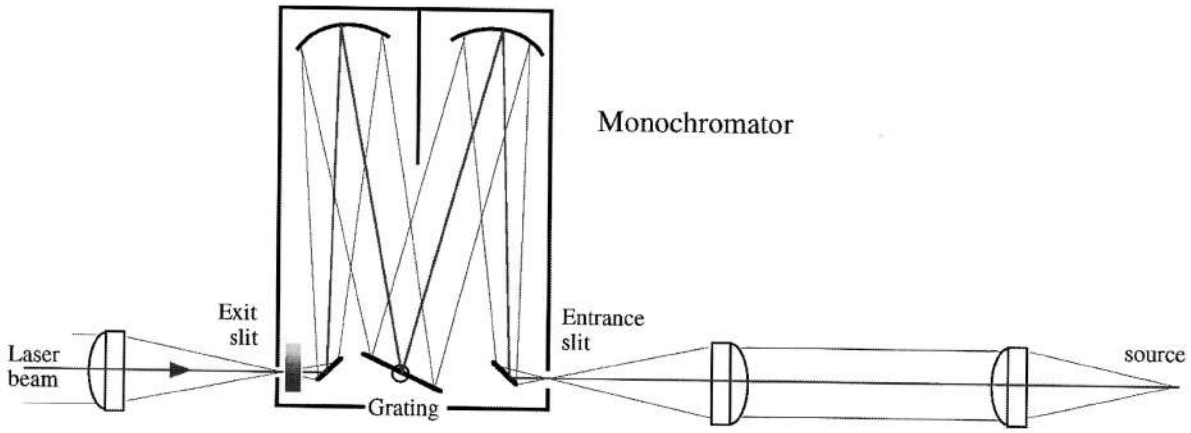

Fig. 4.5 A schematic diagram of a monochromator built into an optical system. White light is drawn from a source to the right and illuminates the diffraction grating, which is mounted on a rotable shaft. As the grating is rotated, different wavelengths are selected from the spectrum formed at the exit slit. For the purposes of lining-up, it is convenient to pass a helium–neon laser beam through the system from the position of the subject's pupil to the source (see Section 4.3.1, and Chapter 1).

The grating surface is exceedingly delicate: felt gloves are recommended when a grating is to be exchanged; there is absolutely no question of cleaning gratings with tissue or cloth; and compressed-air cleaners should be avoided since droplets of liquid may be deposited on the grating surface. Any dust on the grating will scatter white light indiscriminately and impair the purity of the spectral band selected by the exit slit: the best approach is to prevent dust entering the monochromator in the first place.

The bandwidth of the monochromator—the range of wavelengths transmitted—is controlled by the width of both the entrance and the exit slits. If the entrance slit is set to near its minimum and the exit slit is increased, then the width of the selected spectral band will increase linearly. If the exit slit is fixed at some small value, and the entrance slit is increased, a range of overlapping spectra are effectively superposed at the output, again increasing the bandwidth. Most monochromators are designed to be used with the entrance and exit slits set to the same width. Necessarily, the total radiant output will be lower the narrower the slits, and one almost always finds oneself compromising between bandwidth and throughput. One way to maximize the latter without increasing the former is to orient the monochromator (or possibly the source, if its operation will not be impaired) so that the longer dimension of the filament or the arc has the same orientation as the entrance slit of the monochromator.

Monochromators are available with integral stepping motors, allowing the selected wavelength to be changed under computer control. Such instruments make it possible to measure wavelength discrimination by forced-choice psychophysics, although such studies are still surprisingly rare. A suitable design is to present a bipartite field in two successive intervals: in one interval the two half-fields are identical in wavelength and in the other a wavelength change of $\Delta\lambda$ is introduced in one of the half-fields; the subject is required to say in which interval the half-fields differed and $\Delta\lambda$ is titrated according to a staircase routine (Mollon and Cavonius, 1987). In a forced-choice experiment of this kind, the threshold may be as little as 0.3 nm; so an instrument with a step of < 0.1 nm is required if the experimenter wishes to study the limits of human wavelength discrimination. However, this does not require that the *bandwidth* of the stimulus should be equally small: you would need a very powerful source to study photopic discrimination using a bandwidth of 0.1 nm, and the important (and remarkable) fact to remember is that the discrimination of centre wavelengths 0.3 nm apart is being achieved with retinal filters that have bandwidths of 100 nm. So a stimulus bandwidth of 1–2 nm is quite acceptable. A major problem to watch for is play in the coupling between the stepping motor and the grating: this may be equivalent to several nominal steps. The problem can usually be solved by programming your monochromator always to make its final approach to the desired wavelength from a single direction. Even with manual controls, this is good practice.

Manoeuvering a heavy monochromator in the way described in the box (opposite) is much easier if it is mounted on a fully adjustable table. Curiously, suitable tables do not seem to be available commercially: if you ask your workshop to make one, the table should ideally have three legs (each with a small adjustment in height), a lateral adjustment, and a small rotatory adjustment that allows the instrument to be rotated around the centre of the entrance slit. But when the monochromator is aligned in this way, what do you do if you find, as well you may, that the beam from the source does not emerge concentrically with the entering laser beam? If the

Building a monochromator into a Maxwellian-view system can often be a devil of a task. For aligning a monochromator that has entrance and exit slits in-line (Fig. 4.5), the following procedure has been found to be helpful in the writer's laboratory:

1. Pass a laser beam from the position of the subject's pupil to the intended position of the source, using a centre spike to ensure that the beam runs horizontally above the centre of the optical bench.
2. Introduce the source and ensure that the laser beam falls on the centre of the filament (or, in the case of an arc source, close to the tip of the appropriate electrode).
3. Introduce those lenses that are to lie between the source and the monochromator, making sure that the laser beam continues to fall on the centre of the source (see Chapter 1 for detailed guidance on the centring of lenses).
4. Turn on the source, check that its beam is everywhere concentric with the laser beam, and introduce the monochromator so that the source is focused on the centre of the entrance slit.
5. Now gently rotate the monochromator about the latter point until the laser beam passes into the exit slit at its centre.

optics of the monochromator are sealed, then give it back to the manufacturer. If internal adjustments are provided, then you will most probably need to adjust the mounting of the grating to ensure that the grating is perfectly vertical relative to the incoming beam. Another problem to watch for is a wavelength-dependence in the direction of the output beam of the monochromator: any small effect of this kind will mean that the stimulus beam will traverse the subject's pupil as wavelength is changed during the experiment, and the Stiles–Crawford effects will then produce changes in the effective radiance and chromaticity of the beam. If you find this fault, challenge the manufacturer with it.

An unattractive feature of many monochromators, even rather expensive ones, is that a perfectly circular entry beam emerges as elliptical. This distortion arises because spherical mirrors are used to collimate an obliquely incident beam before it is presented to the grating (Fig. 4.5). In principle, you could correct the problem by reflecting the exit beam off an external spherical mirror and adjusting the angle of the mirror to restore the circularity of the beam. The problem is avoided in 'imaging monochromators' by using toroidal mirrors, and in several Jobin–Yvon monochromators by using the same corrected holographic grating to form the spectrum and refocus the beam.

When using a monochromator, it is always necessary to consider the effects of stray light, especially if the selected wavelength is near one extreme of the visible spectrum. Suppose one wishes to use a violet stimulus of 430 nm and suppose the nominal level of stray light is 0.1% throughout the spectrum. When the latter is integrated across the middle of the spectrum, where the eye may be 100 times more sensitive than at 430 nm, it may well be the stray light that the subject detects in a threshold experiment. The problem can be easily eliminated by placing relatively broad-band, gelatin blocking filters in series with the monochromator.

Remember that the output of the monochromator will be partially polarized as a result of reflection at its internal surfaces and that this polarization may interact with later components in the system.

How should the wavelength calibration of the monochromator be checked? I have suggested using a laser to line up the monochromator in a Maxwellian-view system, and so you will necessarily be able to check that the scale reading is 632.8 nm when the helium–neon laser light is maximally transmitted. But the careful experimenter will wish to check at least one other wavelength. Special lamps, emitting one or more spectral lines, are sold for the purpose, and the manufacturers usually imply that you should place these at the position of your usual source. This is the very last thing you would want to do if you had just built the monochromator into a Maxwellian-view system. Instead, you could place the spectral source at the position of the subject's pupil and rely on the principle of reversibility of optical paths. Another, quite cheap, solution is to use a didymium filter. Such filters appear almost neutral in colour but actually absorb in certain, relatively narrow, spectral bands: if you scan your monochromator through the spectrum you can check (with a photodiode or other detector) whether these bands are in their proper places.

4.3.2 Narrow-band interference filters

Interference filters make it possible to select spectral bands only a few nanometres wide. Such filters typically consist of superposed layers of optical materials of different refractive index. Light is multiply reflected to and fro between the internal interfaces and the separations are chosen such that normally incident light of the desired wavelength constructively interferes with itself and is transmitted, whereas other wavelengths are attenuated by destructive interference.

If an interference filter is tilted relative to the incident beam, the favoured wavelength becomes shorter. This property is occasionally useful, but it does mean that interference filters are best used in a well-collimated beam. If the beam is convergent at the filter and rays pass through it at different angles, the selected wavelength in the centre of the beam will be higher than that at its circumference. For precise work, temperature should be controlled, since expansion of the filter may change the separation of the critical surfaces and thus the wavelength selected. Many modern interference filters pass little stray light outside the chosen band, but the warning given above for monochromators should be heeded when working at the extremes of the spectrum.

It is possible to buy 'interference wedges', interference filters that vary linearly in wavelength across their surface. Sometimes they are made in the form of a disc (with wavelength varying circumferentially) so that they can be conveniently mounted on a stepping motor. Such devices may seem seductively cheap by comparison with a monochromator or a set of fixed interference filters, but remember that they are intended to be placed at the focus of a beam and so the considerations of the preceding paragraph mean that they are suitable for use only when a small central area of the beam is to be used as the final stimulus. On the other hand, Palmer and Whitlock (1978) have ingeniously shown how narrow-band stimuli can be produced by projecting on to a single linear interference wedge an inverted image of itself.

Electronically tuneable interference filters of relatively large aperture have recently become available. A device of this kind has been evaluated in the writer's laboratory, but transmission in different regions of the aperture was not uniform enough for the intended purpose (spectral imaging with a digital camera.)

4.3.3 Broad-band filters

A large variety of broad-band colour filters are available in either gelatin or glass. Two standard series are the Kodak Wratten and the Ilford sets, and a booklet of transmission curves can be purchased for each. They are available from photographic dealers. A subset of the Ilford filters (nos 602–626) cover the spectrum in relatively narrow pass-bands. Very cheap filters are available in large sheets from suppliers of stage lighting, e.g. Stage Electrics of Exeter. Such suppliers often give free swatches of sample filters that are themselves big enough to be useful in the visual science laboratory.

Many gelatin filters (especially the yellow, orange, and red ones) are high-pass filters, in that they pass wavelengths above a certain value. The opposite—a low-pass filter that blocks transmission above a certain wavelength—would often be very handy but is not available in gelatin. Interference filters with this property can be bought (e.g. from Melles Griot) but are much more expensive than gelatin filters and usually have a small residual ripple in the low-pass part of their transmission spectrum. 'Dichroic' mirrors are available that transmit part of the visible spectrum and reflect the other half: a mirror of this kind is useful when you are combining two different wavelengths and want to lose as little light as possible when mixing them.

To remove ultraviolet radiation from the beam a Wratten 1A or an Ilford 025 filter is suitable. Infrared radiation can be cheaply removed with HA3 glass (but when you receive it, examine it between crossed polaroids to check for dichroism: otherwise it may interact with polarization in your system to yield striations across the stimulus field). Also available are 'hot' and 'cold' mirrors that reflect or transmit infrared, respectively. If you are working in Maxwellian view and making your calibrations with an unfiltered silicon photodiode, it is very important indeed to remove any infrared radiation that might contaminate your measurements.

4.3.4 Electronic template colorimeters

Potentially the most versatile instrument for producing spectral stimuli is a template colorimeter, in which a spectrum is formed (by means of a prism or a diffraction grating) and an opaque template is used to select the amount of each spectral region to be transmitted. If the light is then recombined (by an integrating sphere, for example), any arbitary spectral power distribution can be offered to the eye. Newton was perhaps the first to use such a device when he filtered his spectrum with a moveable comb, and template colorimeters have often been used in research on colour vision (Ives, 1915; Stiles, 1955; Holtsmark and Valberg, 1969). But what gives them new potential is the availability of electronic masks that can be rapidly varied under computer control. Bonnardel *et al.* (1996) placed a black-and-white liquid-crystal screen (such as sold for use with overhead projectors) in the plane of a spectrum. The device was used to measure the eye's contrast sensitivity for comb-modulated

spectra with varying frequency and phase of comb-modulation, but could in principle be used to generate any spectral stimuli. Until now, two disadvantages of liquid-crystal displays have been that the 'ON'–'OFF' contrast ratio has been relatively poor (as little as 15:1) and that the contrast ratio is wavelength-dependent, being lowest for short-wavelengths. But such devices are improving in contrast, in speed, and in price. Moreover, it is now possible to buy the raw array without the packaging but with the driving electronics (e.g. from CRL Ltd).

The digital micromirror device (DMD), such as recently introduced by Texas Instruments, has great promise as the basis for a template colorimeter. The DMD consists of an array of tiny, hinged mirrors: each individual mirror is about $16\,\mu$m across and can be rotated rapidly between two states (+ 10 degrees and – 10 degrees, corresponding to 'ON' and 'OFF'). Arrays of 800×600 micromirrors are currently available, and each mirror can be thought of as an independent pixel. The effective luminance of the pixel is controlled by the proportion of time the corresponding mirror is in the 'ON' state. A spectrum could be formed by means of a prism or diffraction grating and a DMD placed in the plane where the spectrum is focused before recombination. Each column of pixels would then correspond to a narrow spectral band. At any instant, wavelengths required in the shaped spectral distribution would be reflected back towards recombination, while other wavelengths would be reflected out of this path. Particularly suitable for such an application might be the long thin DMDs (7056×64 pixels) that have been developed for hard-copy devices.

4.3.5 Three-primary colorimeters

If the experimenter requires only colorimetric, and not spectroradiometric, versatility, a range of stimulators can be designed on the basis of three primaries. The latter may be provided by colour filters, by light-emitting diodes, or by lasers.

Economical to build are colorimeters of the Burnham type. A beam is drawn from a single source and is passed through an array of three colour filters (four filters may be used if the experimenter requires access to a larger area of colour space). The array is mounted in the beam and can be mechanically adjusted in its own plane so that more or less of each colour of light is passed. The light is then mixed to give the chromaticity desired. The mixing can be achieved by a glass integrating bar, as in Burnham's original design (Wyszecki and Stiles, 1967; p. 369), by an integrating sphere, or by lenses that form an image of the source. If suitable neutral filters are mounted over each colour filter, then the experimenter can arrange that all the chromaticities generated are of equal luminance. Stepping-motor driven stages may be used to adjust the position of the filter array under computer control. Anyone who contemplates preparing a colorimeter of the Burnham type will find it useful to consult the description of the La Jolla colorimeter by Boynton and Nagy (1981, 1982).

4.3.6 Computer-controlled colour monitors

A special case of the three-primary colorimeter, and perhaps the commonest instrument in colour research today, is the colour CRT (cathode-ray tube). It offers the experimenter great freedom to manipulate stimuli in chromaticity, in space, and in time. It also offers many pitfalls for the unwary.

Colour-mixing is often performed by combining the beams of a multichannel Maxwellian-view system; but after many hours of constructing such a system, the experimenter may be disappointed to discover that the stimulus field is marred by striations or patches of different hue, or that there is a gradient of hue from one side to the other. Tweaking individual components will be of little help in this situation. The writer offers the following solutions:

1. Work up your courage and remove all the components except the source; it will take surprisingly little time to put them back, especially if you record their axial positions carefully (if you are using traditional optical benches, you can use spare saddles as stops to mark positions). Using a laser beam running from the position of the subject's pupil to the source, reintroduce each component in turn, ensuring that the laser beam remains centred on the source. The last fraction of a millimetre counts. For general advice on the alignment of Maxwellian-view systems, see Chapter 1 by W. Makous.
2. Use a ribbon-filament lamp as the source rather than, say, a tungsten–halogen projector lamp. The coiled filament of the latter provides a spatially inhomogeneous source and will contribute to inhomogeneity in the Maxwellian field.
3. (If you can afford some loss of output luminance) form a secondary source by imaging the real source on to a small aperture at which you have mounted a piece of diffusing material, and focus the secondary source on the subject's pupil.

4.3.6.1 The colour CRT

A schematic diagram of a typical colour monitor is shown in Fig. 4.6. Within a vacuum tube, streams of electrons are emitted from the heated cathodes of each of three guns and are drawn towards the front screen of the tube, which is coated with red, green, and blue phosphor dots or stripes. Just before the screen is a 'shadow mask', pierced with apertures, which ensure—more or less—that a given electron beam excites only its proper subset of phosphor elements. The electron beams are deflected electromagnetically from left to right, and, in successive passes, from the top to the bottom of the screen, so that the modulation of the electron flux in time draws out a detailed two-dimensional image.

In the Trinitron design, introduced by Sony, the electron beams are drawn from three horizontally aligned cathodes and the corresponding perforations are nearly continuous apertures between vertical wires. The vertically strung wires of a Trinitron display need one or two (depending on screen size) horizontal wires for support, and these are visible as dark lines. The experimenter will normally be able to work round them, once their nature is understood.

The 'pitch' of a colour monitor is the width of each triad of phosphor dots or lines: a typical value for a 19-inch (48-cm) graphics monitor is 0.28 mm.

4.3.6.2 The input signal

The CRT may be driven either from a 'framestore' that allows picture elements (pixels) to be addressed individually, or from a specialized waveform generator that is dedicated to producing repetitive stimuli with high precision (see Chapter 3). In both cases, a stream of numerical values is transformed by three digital-to-analog converters (DACs) into three corresponding voltage signals, which specify the

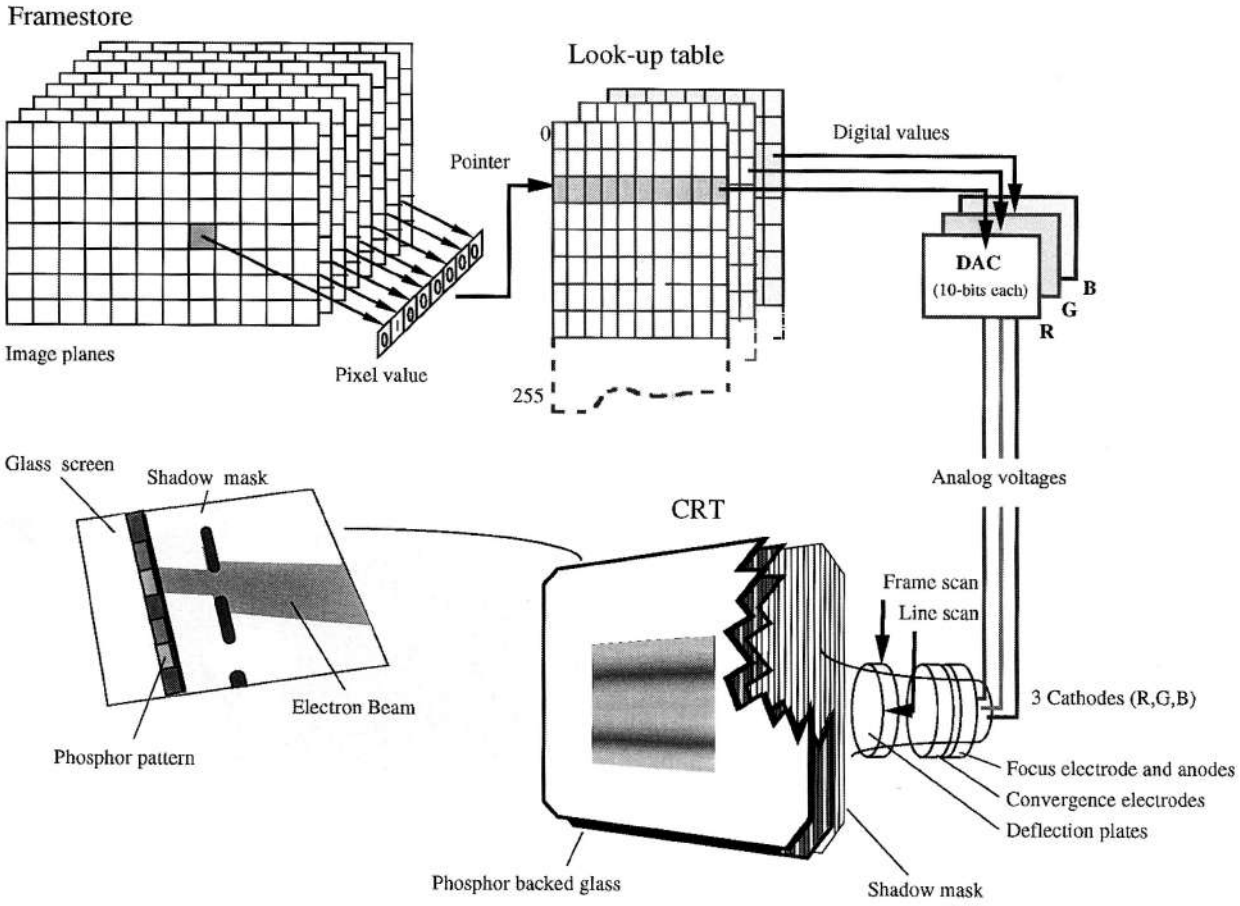

Fig. 4.6 Schema of a graphics board (above) and CRT monitor (below). In the framestore, the value stored for a particular pixel points to one entry in the look-up table. The look-up table depicted here has 256 entries (requiring 8 planes of framestore memory) and for each entry the gun voltage can be specified with a resolution of 10 bits. The monitor shown is based on the Sony Trinitron arrangement, in that the shadow mask or 'aperture grill' consists of vertical 'wires' and the phosphor pattern consists of vertical stripes.

required outputs of the red, green, and blue guns of the monitor. Any distortion of the temporal modulation of these input signals—introduced by the stimulus generator, the cables, the electronics of the monitor, or nearby equipment—will be faithfully translated into a spatial distortion on the screen. It is important to match the temporal properties of the monitor and of the graphics board, rather than buying them independently: more specifically, the 'pixel rate', the frequency with which the graphics board can send new voltages to the monitor, should be matched to the 'bandwidth' or 'video rate' of the monitor, the limiting frequency at which the electron flux of the guns can be modulated (for further details, see Chapter 3).

The upper left part of Fig. 4.6 schematically represents a framestore. The maximum number of individually addressable pixels in the image will be determined by the x, y dimensions of this store. The number of 'planes' of memory (i.e. the number of 'bits per pixel') will determine the number of colour/luminance combinations we can concurrently display. Many systems allow spatial resolution and palette size (number of colour/luminance combinations) to be traded off within a fixed memory array.

To sample colour space adequately we need to be able to control each of three guns with a resolution of 8 bits (giving us 256 different output levels) and if we wish to determine chromatic thresholds, we ideally need at least 10 bits per gun to avoid significant quantization in the measurements (Cowan, 1983). However, if you find yourself limited to 8 bits per gun there are four ways of securing the resolution needed to measure, say, chromatic thresholds on a white background:

1. Use 'dithering' or 'half-toning', gaining an extra bit by setting alternate raster lines to different values.
2. Reduce the contrast of the display by optically superposing it on a homogeneous background field (Cole *et al.*, 1993).
3. When setting up the monitor, turn down the contrast knob and turn up the brightness knob.
4. If the monitor has a high frame rate, then a temporal analog of (1) can be used, in that the same pixel can be set to alternate values on different frames (see Chapter 3.)

It is possible to have a framestore 24-bits deep (i.e. 24 bits per pixel) and thus directly to specify three gun values at each pixel with a precision of 8 bits per gun. But the more common arrangement in recent research has been that shown in the upper part of Fig. 4.6. What is stored for each pixel is not a set of three gun values but a pointer to an entry in a look-up table (Rodieck, 1983; see also Chapter 3). If there are 8 planes of framestore memory, then there will be only 256 possible entries in the look-up table, since this is the number of addresses that can be specified by 8 bits. However, each entry in the look-up table can specify with high precision (say, 10 or 12 bits or more) the signal to be sent to each gun from the digital-to-analog converters of the graphics board. Although the number of bits-per-pixel limits the number of colours concurrently displayed, this palette of colours can be rapidly changed between frames by changing just the entries in the look-up table. If a target area and its background field are linked to different positions in the look-up table, then the target can be made to appear and disappear according to whether the two entries in the look-up table contain the same set of gun values. Thus, although 24-bit graphics systems may recommend themselves for work with spatially complex stimuli (and would be necessary for work on natural scenes), many visual scientists will find it more useful to have a look-up table of limited size combined with DACs of high precision.

Notice that there is no automatic relationship between the pitch of a colour monitor (see previous section) and the number of pixels in a horizontal line: it is left to the user to determine what this relationship should be and to ensure that no aliasing occurs between the stimulus pattern and the pattern of phosphor triads.

4.3.6.3 Non-additivities and non-uniformities

In principle, if the experimenter knows the chromaticities of the three phosphors of a CRT and the gamma function for each gun (the relationship between the input voltage and the luminous output; see Chapter 3), then it is possible to calculate the gun values needed to generate any stimulus within the gamut of the monitor. However, this calculation depends on a number of assumptions, which need not hold.

A critical assumption is that of 'gun independence', the assumption that the luminous and chromatic output resulting from a given input to a given gun is independent of the signals sent to the other two guns. If we measure the X, Y, and Z outputs when given voltages are sent to each gun separately, we can determine whether the sum of the three sets of X, Y, and Z is obtained when the specified voltages are sent to the three guns simultaneously. This is a basic test to perform before paying for a new monitor. In the author's laboratory, good additivity has been found for Barco CD351, Mitsubishi HL20, and Sony GVM1400 monitors driven by a Cambridge Research Systems VSG board. However, failures of additivity may be observed in both cheap and expensive colour monitors, especially at high settings of contrast. In the case of cheap monitors, the non-additivity may arise from a low accelerating voltage present at the screen (Rodieck, 1983). In the case of modern expensive monitors, the culprit may be circuits deliberately introduced to ensure compliance with regulations on X-ray emission: these circuits place limits on total emission. Non-additivities of these kinds can often be eliminated or reduced by turning down the gain (the 'contrast') of the monitor (thus reducing the maximum luminance available from each gun). In principle, non-additivities may also arise in the graphics board.

An assumption that fails for all monitors known to the author is that of spatial uniformity of gun outputs across the screen: usually there is a falling off of luminous output near the edges of the screen and this may be as much as 20% (Mollon and Baker, 1995). However, Brainard (1989) found that most of the spatial variation could be described by a single scale factor applied to all gamma functions; so there is little change in chromaticity.

A third assumption is that the luminance and chromaticity from a given region of the screen are independent of the illumination of other regions. Mollon and Baker (1995) illustrate the complex ways in which this assumption can fail in research monitors. It is important for the experimenter to check thoroughly for artefacts of this kind, above all in experiments on colour constancy and simultaneous contrast, but I suspect that such checks are far too often neglected.

Another problem is that the time constants of the red, green, and blue phosphors are quite likely to differ. Vingrys and King-Smith (1986) have described how these differences may lead to detectable luminance artefacts when a temporal substitution is made between nominally equiluminant stimuli.

Finally, the nature of the raster scan means that different parts of a monitor screen are displayed at different times. It ought not to be necessary to remark on this. But I have refereed papers on electroretinogram (ERG) recording where the subject was placed 10 cm from a large monitor and a temporal red–green modulation was applied to the entire screen: under such conditions there will be a substantial phase shift between the signals generated in the inferior and superior retina.

4.3.7 Digital light processors (DLP)

This is the term that Texas Instruments have adopted for projection devices that incorporate their digital micromirror devices (see Section 4.3.4). In the single-chip version of the DLP, light from a high-intensity source passes through a rotating red–green–blue (RGB) colour filter and on to the mirror array. At any instant each

pixel—each individual mirror—is reflecting light either into or out of the aperture of a lens that images the array on a screen. The input to the currently available system is the SVGA standard used for computer monitors, and the intensities of the separate RGB signals at each pixel are achieved by controlling the proportion of the time that the corresponding mirror is in the 'ON' state as the red, green, or blue filter is passing through the beam. As a result of this method of manipulating the intensity of each primary, the response of the device is linear. The images are bright and sharp. In principle, the individual pixels are independent of their neighbours and each primary is independent of the other primaries; so two of the types of interaction that characterize CRTs (Section 4.3.6) should be absent. And, in contrast to the case of a CRT, the chromaticities of the primaries could be modified to suit particular experimental purposes.

4.3.8 Munsell papers

Our colour vision evolved to discriminate coloured surfaces and to detect coloured targets in variegated backgrounds. It is surprising therefore that surface colours are rather little used in modern colour research, except in the study of colour deficiency and of colour naming. In the field of colour constancy, many experimenters simulate coloured surfaces on CRTs but seldom discuss the (very relevant) issue of how well they achieved the illusion.

An internationally recognized collection of coloured papers is that of the Munsell system, available from Macbeth (or in Europe from D. G. Colour). The papers can be bought both in atlas form and as individual sheets of varying size. Specially chosen subsets are available for particular purposes: thus there is a useful small atlas of skin, hair, and iris colours. General information about the Munsell system is available on the web at http://www.munsell.com. The papers are specified in terms of perceptual dimensions—Hue, Chroma (saturation), and Value (lightness)—but colorimetric specifications (CIE x, y values) are tabulated and plotted in Wyszecki and Stiles (1982, p. 840 ff). Notice that these chromaticity coordinates are valid only for the specified illuminant (CIE Illuminant C, an approximation to natural daylight of colour temperature 6774 K).

In general, since the spectral stimulus at the eye depends on the product of (a) the reflectance spectrum of a surface and (b) the spectral power distribution of the illuminant, it is essential when working with surface colours either to use a standard illuminant or to measure instrumentally the chromaticities (or spectral fluxes) actually presented to the eye by the coloured surfaces. For any exact work that relies on a standard illuminant, not just the chromaticity but also the spectral power distribution of the illuminant must be standard: otherwise the product of the spectral reflectance of the surface and the spectral power distribution of the illuminant may not give the intended spectral flux at the eye.

The Macbeth easel lamp, which consisted of a standard tungsten lamp and a blue glass filter, used to be a practical means of simulating Illuminant C. It has been replaced by the Macbeth 'Sol Source', an adjustable desk lamp with a tungsten–halogen source and blue filter. The 'Sol Source' is sold as a simulation of the standard Illuminant D65, which is nowadays preferred by the CIE as an approximation to

natural daylight and which corresponds to a slightly lower colour temperature of 6500 K (Wyszecki and Stiles, 1982; p. 144). A cheap way of simulating Illuminant C, suitable for the routine administration of pseudoisochromatic plates and similar materials, is to illuminate the coloured materials with an ordinary tungsten source and to ask the subject to wear trial frames fitted with suitable colour filters: Pokorny *et al.* (1978) used a combination of Wratten filters nos 78B and 80B. Fluorescent lamps exist that approximate to D65 (e.g. Philips 'Colour 96' lamps): these would not serve for very exact work, since their spectral power distributions exhibit marked emission lines that are absent from D65, but they cost only a little more than ordinary fluorescent tubes and they could be useful where, say, experimenters wish to ensure that their subjects are in a known state of neutral adaptation.

4.3.9 Control of intensity when broad-band spectral stimuli are used

If broad-band colour stimuli, rather than strictly monochromatic beams, are used, then care is needed in the manipulation of intensity. 'Neutral' density filters will seldom attenuate all wavelengths by exactly the same factor; the extinction produced by crossed polaroids is very wavelength-dependent, typically being least at short wavelengths; varying the voltage or current available to a tungsten-filament lamp will change its colour temperature, making it redder the lower the current. A classical solution is physically to interrupt a known proportion of light, temporally with a rapidly rotating episcotister or spatially with a filter consisting of an array of small apertures; but often the best thing to do is to construct the required chromaticity by using monochromatic lights in the first place. Neutral-density filters can then be calibrated in the beam in which they are being used.

4.4 Measuring coloured stimuli

4.4.1 Spectroradiometry

If working in Maxwellian-view, the colour scientist can get a long way with a simple silicon photodiode that has been calibrated absolutely and spectrally. Any desired chromaticity can be constructed from individually measured monochromatic lights. A photodiode is also useful for measuring the gamma functions of the individual guns of a colour monitor.

A new generation of spectroradiometers are available in which the incoming light is dispersed on to an array of photodiodes and the energy in each spectral band is measured in parallel. Such devices especially recommend themselves where the stimuli are relatively broad-band. Instruments suitable for measuring surface reflectance are available for as little as $3000; the author has good reports of the Ocean Optics PS1000 instrument, which is small enough to be used in the field with a portable computer. Telespectroradiometers, suitable for measuring the spectral flux reaching an observer from a distant source, are more expensive. The author has had very satisfactory experience with a Photo Research 650 instrument, which has proved stable and accurate under the extreme conditions of a tropical rain forest.

When measuring CRT screens with radiometric devices it is important to remember that the stimulus is concentrated into brief, intense pulses, which may transiently saturate the detector and thus give you a distorted measurement. Some spectroradiometers have a special mode in which measurements are synchronized to the frame rate of a CRT (although I have found the Photo Research 650 unwilling to lock on to some monitors). Another factor to consider when calibrating CRTs is that the radiation emitted from the screen may be partly polarized and components in the detecting system may act as analysers.

To check the calibration of the spectroradiometer itself, the laboratory will need a standard light source and a very stable, constant-current power supply. A suitable system, after calibration at a standards laboratory, may cost $5000—more than some spectroradiometers.

4.4.2 Colorimeters

Electronic colorimeters directly estimate the CIE X, Y, and Z tristimulus values rather than first measuring the spectral power distribution of the stimulus. The colorimeter typically incorporates three detectors, which measure the light passed by three colour filters. The three detector/filter combinations have sensitivities that approximate to the CIE $\bar{x}$, $\bar{y}$, $\bar{z}$ functions, and such instruments will be accurate only to the degree to which they successfully simulate these functions. Some widely used instruments take a short-cut in simulating the $\bar{x}$ function: the detector/filter combination is unimodal in its spectral sensitivity (whereas $\bar{x}$ is bimodal) and the short-wave mode of $\bar{x}$ is simulated by adding the $\bar{z}$ signal, suitably scaled, to the long-wave lobe of $\bar{x}$.

The C1200 colorimeter made by Lichtmesstechnik was recommended in the survey of Berns *et al.* (1993). This instrument achieves high precision by the device of partial filtering: small pieces of secondary filter are placed in series with the primary filters to improve the approximation to the $\bar{x}$, $\bar{y}$, $\bar{z}$ functions. However, the price of such a colorimeter is comparable to that of a diode-array spectroradiometer, such as the Photo Research 650.

There is still a place for visual colorimeters, in which a match is made between the sample light and light drawn from a reference source and passed through calibrated glass filters. The Lovibond Tintometer (Tintometer Ltd) is an instrument of this type. Although such instruments are slow to use and the experimenter's vision may not be exactly that of the CIE standard observer, the calibrations are likely to be at least as valid as those made with a mid-priced electronic colorimeter, and it is certainly difficult to make gross errors with a visual colorimeter. For measurements in the short-wave corner of the CIE diagram, where colour matches become imprecise, Mollon and Baker (1995) recommend a method of triangulation using desaturated stimuli.

4.5 Screening tests for colour vision deficiencies

Some 8% of Caucasian males exhibit a hereditary anomaly or deficiency of colour vision. The experimenter will often wish to exclude such subjects from experiments in

which colour is a variable. The most sensitive, widely available screening test is offered by the Ishihara plates, available from Keeler and from many other suppliers. These plates are not suitable for classifying colour-deficient observers, but will efficiently exclude all but a tiny minority of anomalous trichromats (the *mini-malanomale Trichromaten* of Vierling). The plates must be presented under a standard illuminant. A computer-based test that quantifies the loss of colour discrimination is described by Regan *et al.* (1994) and is available to run on Cambridge Research Systems graphics boards. For a full treatment of tests for the diagnosis of colour deficiency, see Pokorny *et al.* (1979) or Birch (1993).

Acknowledgements

I am grateful to the editors and to C. R. Cavonius, G. Jordan, B. C. Regan, and Y. Tadmor for comments on the text.

References

Barlow, H. B. (1982). What causes trichromacy? A theoretical analysis using comb-filtered spectra. *Vision Research*, **22**, 635–43.

Baylor, D. A., Nunn, B. J., and Schnapf, J. L. (1987). Spectral sensitivity of cones of the monkey *Macaca fascicularis*. *Journal of Physiology*, **390**, 145–60.

Berns, R. S., Motta, R. J., and Gorzynski, M. E. (1993). CRT colorimetry. Part II. Metrology. *Color Research and Application*, **18**, 315–25.

Birch, J. (1993). *Diagnosis of defective colour vision*. Oxford: Oxford University Press.

Bonnardel, V., Bellemare, H., and Mollon, J. D. (1996). Measurements of human sensitivity to comb-filtered spectra. *Vision Research*, **36**, 2713–20.

Boynton, R. M. and Nagy, A. L. (1981). *The La Jolla analytic colorimeter: optics, calibrations, procedures and control experiments*, CHIP Report No. 109. Center for Human Information Processing, University of California, San Diego, CA.

Boynton, R. M. and Nagy, A. L. (1982). La Jolla analytic colorimeter. *Journal of the Optical Society of America*, **72**, 666–7.

Brainard, D. H. (1989). Calibration of a computer controlled color monitor. *Color Research and Application*, **14**, 23–34.

Cole, G. R., Hine, T., and McIlhagga, W. (1993). Detection mechanisms in L-, M- and S-cone contrast space. *Journal of the Optical Society of America*, **10**, 38–51.

Cowan, W. B. (1983). Discreteness artifacts in raster display systems. In *Colour vision: physiology and psychophysics* (ed. J. D. Mollon and L. T. Sharpe), pp. 145–53. Academic Press, London.

Derrington, A. M., Krauskopf, J., and Lennie, P. (1984). Chromatic mechanisms in lateral geniculate nucleus of macaque. *Journal of Physiology (London)*, **357**, 241–65.

Holtsmark, T. and Valberg, A. (1969). Colour discrimination and hue. *Nature*, **224**, 366–7.

Ives, H. E. (1915). A precision artificial eye. *Physical Review*, **6** (*New Series*), 334–44.

Jacobs, G. H., Neitz, J., and Deegan, J. F. (1991). Retinal receptors in rodents maximally sensitive to ultraviolet light. *Nature*, **353**, 655–6.

Kaye, G. W. C. and Laby, T. H. (1973). *Tables of physical and chemical constants* (14th edn). Longman, London.

MacLeod, D. I. A. and Boynton, R. M. (1979). Chromaticity diagram showing cone excitation by stimuli of equal luminance. *Journal of the Optical Society of America*, **69**, 1183–6.

Mansfield, R. J. W. (1985). Primate photopigments and cone mechanisms. In *The Visual System* (ed. A. Fein and J. S. Levine) pp. 89–106. Alan R, Liss, New York.

Mollon, J. D. and Baker, M. R. (1995). The use of CRT displays in research on colour vision. In *Colour vision deficiencies XII* (ed. B. Drum), pp. 423–44. Kluwer Academic, Dordrecht.

Mollon, J. D. and Cavonius, C. R. (1986). The discriminability of colours on c.r.t. displays. *Journal of the Institution of Electronic and Radio Engineers*, **56**, 107–10.

Mollon, J. D. and Cavonius, C. R. (1987). The chromatic antagonisms of opponent process theory are not the same as those revealed in studies of detection and discrimination. In *Colour vision deficiencies VIII* (ed. G. Verriest), pp. 473–83. Martinius Nijhoff/Dr W. Junk, Dordrecht.

Mollon, J. D. and Jordan, G. (1988). Eine evolutionäre Interpretation des menschlichen Farbensehens. *Die Farbe*, **35/36**, 139–70.

Palmer, D. A. and Whitlock, C. A. (1978). An improved monochromator using a single graded interference filter. *Journal of Physics, Section E: Scientific Instruments*, **11**, 996–7.

Pokorny, J., Smith, V. C., and Lund, D. (1978). Technical characteristics of 'Color-Test Glasses'. *Modern Problems in Ophthalmology*, **19**, 110–12.

Pokorny, J., Smith, V. C., Verriest, G., and Pinckers, A. J. L. G. (1979). *Congenital and acquired color vision defects*. Grune and Stratton, New York.

Regan, B. C. and Mollon, J. D. (1997). The relative salience of the cardinal axes of colour space in normal and anomalous trichromats. In *Colour vision deficiencies XIII* (ed. C. R. Cavonius and J. D. Moreland), pp. 261–70. Kluwer, Dordrecht.

Regan, B. C., Reffin, J. P., and Mollon, J. D. (1994). Luminance noise and the rapid determination of discrimination ellipses in colour deficiency. *Vision Research*, **34**, 1279–99.

Rodieck, R. W. (1983). Raster-based colour stimulators. In *Colour vision: physiology and psychophysics* (ed. J. D. Mollon and L. T. Sharpe), pp. 131–44. Academic Press, London.

Rushton, W. A. H. (1972). Pigments and signals in colour vision. *Journal of Physiology*, **220**, 1–31P.

Smith, V. C. and Pokorny, J. (1975). Spectral sensitivity of the foveal cone photopigments between 400 and 500 nm. *Vision Research*, **15**, 161–71.

Stiles, W. S. (1955). 18th Thomas Young Oration. The basic data of colour-matching. In *Physical Society Year Book*, pp. 44–65. *The Physical Society, London.*

Stromeyer, C. F., Cole, G. R., and Kronauer, R. E. (1985). Second-site adaptation in the red–green chromatic pathways. *Vision Research*, **25**, 219–37.

Turner, W. H. (1973). Photoluminescence of color filter glasses. *Applied Optics*, **12**, 480–6.

Vingrys, A. J. and King-Smith, E. (1986). Factors in using color video monitors for assessment of visual thresholds. *Color Research and Application*, **11**, 557–62.

Vos, J. J. (1978). Colorimetric and photometric properties of a 2 deg fundamental observer. *Color Research and Application*, **3**, 125–8.

Webster, M. A. and Mollon, J. D. (1995). Colour constancy influenced by contrast adaptation. *Nature*, **373**, 694–8.

Winderickx, J., Lindsay, D. T., Sanocki, E., Teller, D. Y., Motulsky, A. G., and Deeb, S. S. (1992). Polymorphism in red photopigment underlies variation in colour matching. *Nature*, **356**, 431–3.

Wyszecki, G. and Stiles, W. S. (1967). *Color science*. Wiley, New York.

Wyszecki, G. and Stiles, W. S. (1982). *Color science* (2nd edn). Wiley, New York.

Further reading

An indispensable piece of equipment in any colour laboratory is a copy of *Color science* by G. Wyszecki and W. S. Stiles (published by Wiley), which offers authoritative detail on the

topics introduced above. For many purposes, the first (1967) edition is the better. Also very useful, and free, are the catalogues distributed by the larger suppliers of optical apparatus (for example, the *Monochromators and light sources* volume from Oriel): it has become the custom for these catalogues to contain detailed tutorials on several of the topics covered in the present chapter. Kaye and Laby's standard reference work *Tables of physical and chemical constants* (Longmans, London) contains useful sections on optics, radiation, and colorimetry.

Appendix: Addresses of suppliers

Beuth Verlag, Burggrafenstraße 4–10, Berlin, Germany.
Cambridge Research Systems, 80 Riverside Estate, Sir Thomas Langley Rd, Rochester, Kent ME2 4BH, UK (www.crsltd.com).
CRL, Dawley Rd, Hayes, Middlesex, UB3 1HH, UK.
D. G. Colour, 138 Greenwood Ave, Salisbury, Wilts SP1 1PE, UK.
Keeler, Clewer Hill Rd, Windsor, Berks SL4 4AA, UK.
Lichtmesstechnik, Helmholtzstr. 9, D-10586 Berlin, Germany.
Macbeth, Little Britain Rd, PO Box 230 Newburgh, NY 12550, USA; Macbeth House, Pacific Road, Altrincham, Cheshire WA14 5BJ, UK (www.munsell.com).
Melles Griot, 1770 Kettering St, Irvine, CA 92714, USA (www.mellesgriot.com).
Ocean Optics, 1237 Lady Marion Lane, Dunedin, FL 34698, USA (www.oceanoptics.com).
Oriel, PO Box 872, Stratford, CT 06497–0872, USA; 1 Mole Business Park, Leatherhead, Surrey K22 7BR, UK.
Photo Research, 9330 DeSoto Ave, PO Box 2192, Chatsworth, CA 91313–2192, USA (www.photoresearch.com).
Stage Electrics, Cofton Rd, Marsh Barton, Exeter EX2 8QW, UK.
Texas Instruments Inc., PO Box 655012, M/S 6, 13532 North Central Expressway, Dallas, Texas 75265, USA (www.ti.com).
Tintometer Ltd, Waterloo Rd, Salisbury, SP1 2JY, UK.

Notes

1. Although radiations outside these limits can normally be neglected with respect to their effect on human colour vision, the cautious experimenter has at least three reasons to exclude them from the stimulus: (a) ultraviolet (UV) radiation may be a health hazard, especially if an arc lamp is the source (see Chapter 1); (b) any extraspectral radiation may contaminate radiometric calibrations (see below); and (c) UV radiation may be re-emitted as visible fluorescence/luminescence by glass components such as colour filters and by the lens and other components of the eye. Turner (1973) gives examples of how glass colour filters may emit significant visible light when exposed to a xenon lamp emitting UV radiation.
2. In more formal treatments, X, Y, Z are the amounts of three imaginary reference lights that, when mixed together, produce the same sensation as the test colour.
3. In addition to the values x and y, we can also define $z = Z/(X + Y + Z)$.The latter does not appear directly in the diagram, but can be readily derived since the nature of the definitions mean that the sum of x, y, and z is always unity.