38 Bedford Square
London, WC1 3EL
UNITED KINGDOM

t: +44 (0)20 7631 4905
f: +44 (0)20 7631 4906
www.filmlight.ltd.uk

FilmLight
DIGITAL FILM TECHNOLOGY

Technical Note **FL-TN-00-002**

Standard Colour Spaces

Richard Kirk

Creation Date	09 January 2004
Last Modified	10 November 2004
Revision No	1.65

Summary

In 1931, the Commision Internationale d'Éclairiage (CIE) recommended a system for colour measurement. This system allowed the specification of colour matches using the CIX XYZ tristimulus values. In 1976, the CIE recommended the CIE LAB and CIE LUV colour spaces for the measurement of colour differences, and colour tolerances. These colour spaces, and their more modern variants are the basic tools for modern colorimetry.

CIE colour spaces are not widely used for images. Scanners and printers usually communicate in their own device-dependent RGB. Video has its own colour standard. Densitometers use Status A and Status M colour spaces. Truelight uses all of these spaces to build up a colour transform from film expose units to display RGB.

You can use Truelight without knowing all about CIE colour spaces. However, if you wonder why the XYZ and the L*a*b* calibration for the same monitor look different, you may find your explanation here.

Whites are dealt with in a separate section. Most of us know what white paper and white paint is. We might think we know what white light is too. A full discussion of what is and what is not 'white' is much too big a task for this small note, but we introduce a few basic issues.

FilmLight
DIGITAL FILM TECHNOLOGY

1 Matching a light source

Suppose we have three LEDs, which give out red, green, and blue light. If we vary the current through an LED, the light output intensity will vary in proportion, but the spectrum will remain unchanged.

We could use these three LEDs to estimate the colour of an unknown light source using an apparatus like this:

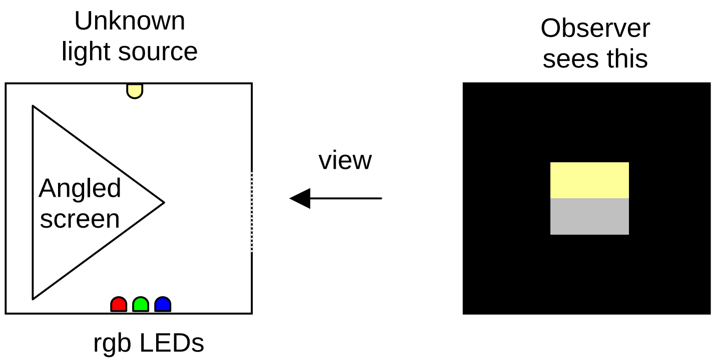

We adjust the current though the LEDs until the two sides of the screen appear the same colour. We then measure the current through the three LEDs in milliamps. This gives us three rgb numbers that will identify the unknown light source. Of course, our rgb figures would be meaningless to someone who did not have our apparatus.

1.1 Matching a colour

Suppose we halved the signal going through our three LEDs. Our LEDs would be outputting the same ratio of red, green, and blue light, so the LED illuminated screen would appear the same colour as our unknown light source, but it would be half as bright.

Let us make a new **RGB** = rgb/(r+g+b).

Our RGB colour space has three dimensions. Our new **RGB** colour space has only two independent dimensions, because **R** + **G** + **B** = 1. Our scaling has taken out the brightness factor. This does not mean all the colours will appear the same brightness to us. That will depend on our original choice of RGB primaries. But it will mean every colour is represented with a unique brightness.

The diagram opposite shows the triangle of colours we get from a particular set of **RGB** primaries. If you plot colours like this, the colours you get by mixing any two colours will always lie on the straight line joining those two colours.

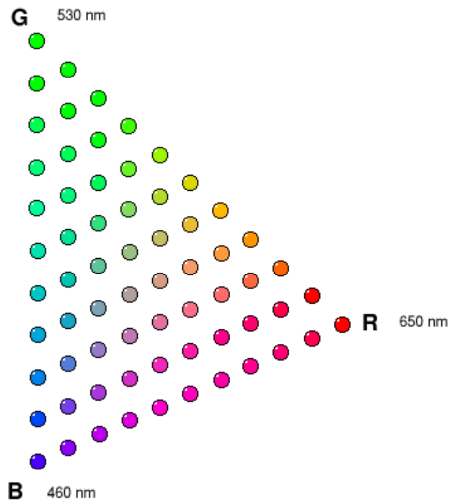

FilmLight
DIGITAL FILM TECHNOLOGY

1.2 Matching an out of gamut colour

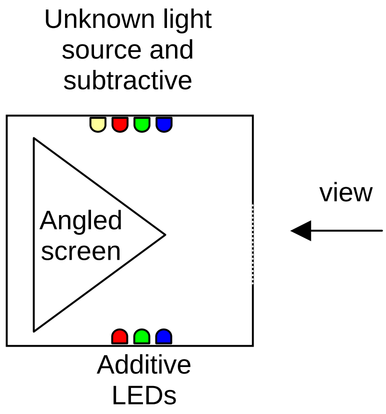

Our **RGB** colour space in section 1.1 has a triangular gamut. There may be colours outside this triangle, but we cannot get a match for them without using negative **RGB** values. We cannot get negative RGB illuminations from our LEDs, but we can modify the apparatus in section 1:

We cannot add negative amounts of blue light, but we can now add light to the other side of the balance, which has the same effect. Now we can have negative RGB values, we can extrapolate outside the **RGB** triangle.

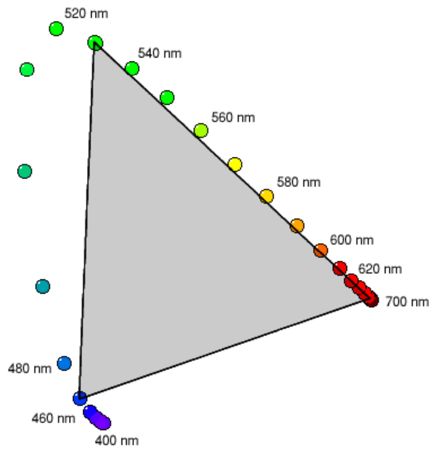

In the diagram opposite, we have plotted out the spectrum at 10nm intervals to the same scale as the diagram in section 1.1. The grey triangle represents the gamut of positive **RGB** values.

In the diagram in section 1.1 we had primaries of 650nm (red), 530nm (green) and 460nm (blue). The original CIE researchers for the 1931 standard used white lights with coloured filters. These experiments were repeated when suitable lasers came available.

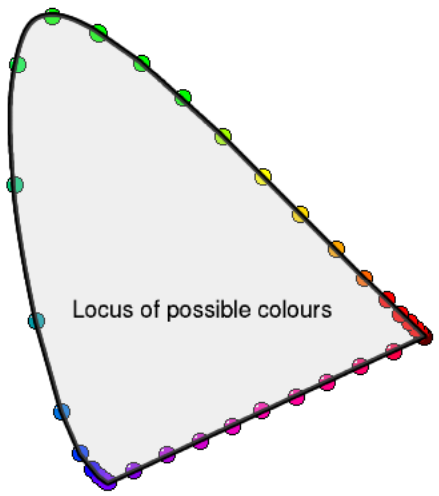

This locus of all the possible colours is the grey region shown in the diagram opposite. The convex curved edge is the locus of all spectral colours. The straight line closing the bottom of the spectral locus is the locus of all colours made by mixing the shortest (violet) and the longest (red) visible wavelengths.

2 CIE XYZ Tristimulus

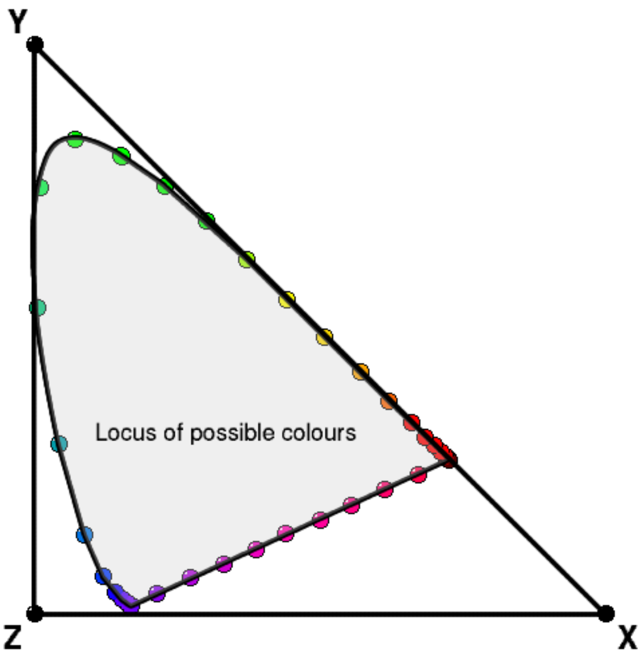

Here, we have taken the diagram from section 1.2, and added three new points, labelled X, Y, and Z. These three colours are the CIE XYZ primaries, normalised in the manner of RGB in section 1.2.

We have not chosen our XYZ points to form a nice, right-angled triangle. Rather, we arranged our earlier plots, so things would turn out this way. The CIE reasons for choosing their XYZ primaries are more cunning.

The 1931 CIE researchers used primaries of 700nm (red), 546.1nm (green) and 435.8nm (blue) to measure their spectral locus. These primaries did not have equal brightness, so they could choose an extrapolated X and Z with zero visual brightness. They could do this because XYZ are not real colours.

Y is the only primary with brightness. It is the Y brightness value in colour spaces such as Yxy, and Yu'v'.

NB: Y is not the Y in video colour spaces such as YIQ and YCrCb.

The XY line was chosen to be a tangent to the spectral locus.

The ends of the spectral locus are the limits of human vision. These would seem to be unique points in the visual colour space. However, it is difficult to locate these points reliably, because the eye has very little sensitivity to these wavelengths. The ends of the spectral locus do not have any significant values in XYZ. The violet point does not lie on the **ZX** line

The point **X** = **Y** = **Z** is a white-ish colour. It has no particular significance.

2.1 Measuring tristimulus values

We can measure absolute light intensities in milliwatts using a bolometer. This is a cool black body with a temperature sensor such as a thermocouple embedded in it. All radiation landing on it, irrespective of wavelength, converts to heat. Bolometers are not very sensitive devices. Usually bolometers are used to calibrate other more sensitive but non-absolute light-measuring devices, such as photomultipliers or photodiodes.

We can calculate the relative spectral contributions from the graph of section 2.

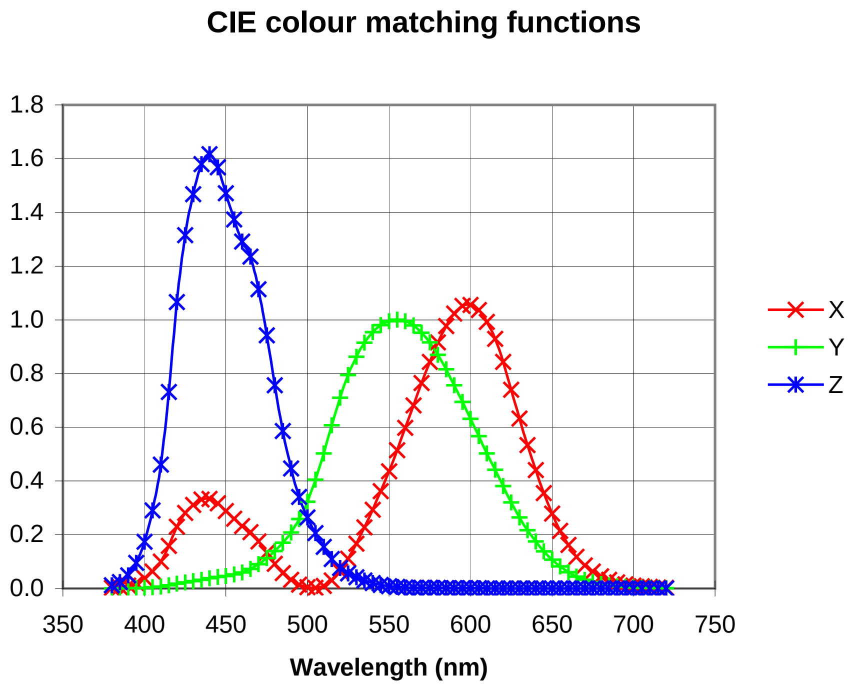

The CIE published the original standard observer colour matching functions in 1931. There have been several revisions since then as better monochromatic light sources have become available. This graph shows the CIE 1931 observer as modified by Vos (1978). The main changes have been in the blue-violet region (350-450 nm), increasing the bend at the violet end of the spectral locus.

This can be important: not all CIE XYZ measurements are quite the same. You can get slightly different values from the same spectrum, depending on the standard observer functions. Once you have converted to XYZ, there is no way to correct for the differences in standard observers unless you can get back the original spectra. There is a lot of published work based on the original 1931 observer, which is why it is still used, even though more accurate values exist.

Actually, for most practical work, there is little difference between the different CIE colour matching standards. Some phosphors have a violet emission line, but most artificial light contains little violet, and the eye is not particularly sensitive to violet anyway. Often, the differences between different colour matching functions are smaller than the experimental or calibration errors you get with real colour measuring instruments. Truelight does not export or import XYZ images. If you use the same instrument for all your XYZ measurements, any systematic errors should cancel themselves out.

2.2 XYZ units

XYZ measurements are usually in foot-lamberts or candelas/sq.m.

The SI unit of luminous intensity is the candela. A uniform point light source that emits $4\pi/683$ watts of light with a wavelength of 555nm will have a luminous intensity of one candela. If you look back at the graph in section 2.1, you will see the peak of the Y colour matching function lies at 555 nm. We scaled the Y-axis to make the colour matching value at this point 1.0. We can calculate the luminous intensity for wavelengths other than 555nmn by multiplying by this Y colour matching value.

If you are measuring the light coming from a screen, you will be measuring the luminous intensity per unit area, hence the SI unit of candelas/sq.m. One foot-Lambert is equal to 3.426 candelas/sq.m. Here are some typical values:

Condition	Foot-Lamberts	Candelas/sq.m
Scene in bright sunlight	1500	5000
Optimum for human eye	600	2000
White paper under typical reading light	150	500
SMTPE video standard white	30	100
SMTPE cinema standard white	16	55
Mesopic threshold (rod cells saturating)	0.3	1
Photopic threshold (limit of colour vision)	0.03	0.1

There are many more photometric units, but these are probably the only ones you will need for Truelight.

The SI unit of illumination is the **lux**. If we illuminate a white screen with 1 lux, it will have a luminance of $1/\pi$ candelas/sq.m. Typical cinema low lighting is about 1 lux. If someone in the audience is wearing white, they will reflect a luminance of about 0.3 candelas/sq.cm. This is less than 1% of the film white.

3 Luminance-chrominance co-ordinates

3.1 Yxy co-ordinates

Yxy is the simplest luminance-chrominance space. It is derived from XYZ using…

$$x = \frac{X}{X+Y+Z}, \quad y = \frac{Y}{X+Y+Z}$$

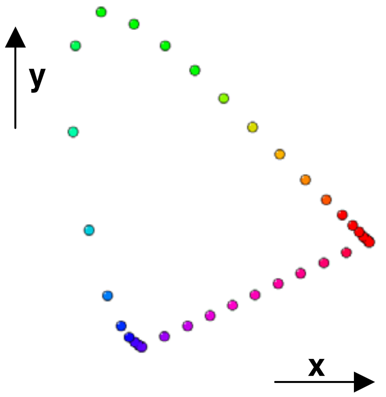

Y is the luminance measured in foot-Lamberts or candelas/sq.m, and the x and y co-ordinates are dimensionless. The inverse transform is…

$$X = \frac{x \cdot Y}{y}, \quad Z = \frac{(1-x-y) \cdot Y}{y}$$

The xy colour space is simple, but it is uneven. A small difference in x and y in the blue-violet corner is more visible than the same difference in the green corner.

3.2 Yu'v' co-ordinates

Yu'v' is a uniform luminance-chrominance space. It is derived from XYZ using…

$$u' = \frac{4X}{X+15Y+3Z}, \quad v' = \frac{9Y}{X+15Y+3Z}$$

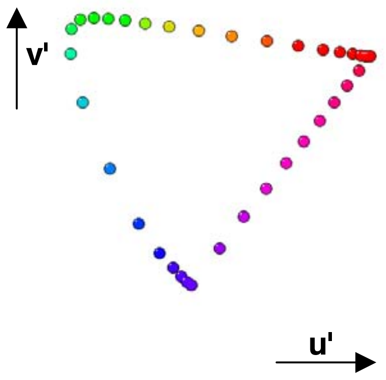

Y is the luminance measured in foot-Lamberts or candelas/sq.m, and the x and y co-ordinates are dimensionless. The inverse transform is…

$$X = \frac{9Y \cdot u'}{4v'}, \quad Z = \frac{3Y \cdot (4-u')}{4v'} - 5Y$$

The blue-violet corner has been stretched out and the green corner has been compressed. It is not perfect, but it is much better than Yxy. Additive mixtures still lie on straight lines.

Yuv was an earlier attempt at a uniform colour space that was replaced by Yu'v' in 1975. There is no reason to use Yuv today.

4 Perceptually uniform colour spaces

Our eyes adapt to different illumination conditions. An object will seem to have the same colour to us under very different lighting conditions. Our adaption is not perfect, but most of the time and over a range of lighting conditions, we see colour as a property of an object, and not a property of the illumination.

This suggests that any 2:1 luminance contrast should look the same to us, irrespective of the absolute luminance. The contrast between 20 and 10 ft-lamberts ought to look the same to us as the contrast between 200 and 100 ft-lamberts, or 2.0 and 1.0 ft-lamberts.

There are limits to this. The table in section 2.2 suggests we may get dazzled if our illumination goes much over 2000 ft-lamberts, and we will not see colour at 0.2 ft-lamberts. We would not see a contrast between 2.0 and 1.0 ft-lamberts in an average luminance of 100 ft-lamberts, because of stray light within the eye. However, there should be a luminance range where equal contrast ratios gave equal eye stimuli.

Suppose we have two grey patches with luminances of 90 ft-lamberts and 10 ft-lamberts. A patch of 50 ft-lamberts would be halfway between the two in luminance units. A patch of 30 ft-lamberts should appear to be halfway between the two, because it has a 3:1 contrast ratio either way.

4.1 CIE L*a*b*

The CIE L*a*b* is defined in terms of the stimulus XYZ and the reference white $X_nY_nZ_n$ using a function f as follows:

$$L^* = 116 \cdot f\left(\frac{Y}{Y_n}\right)$$

$$a^* = 500 \cdot \left[f\left(\frac{X}{X_n}\right) - f\left(\frac{Y}{Y_n}\right)\right]$$

$$b^* = 200 \cdot \left[f\left(\frac{Y}{Y_n}\right) - f\left(\frac{Z}{Z_n}\right)\right]$$

$$f(w) = \begin{cases} w^3 & w \geq 0.008856 \\ 7.787w + 0.1379 & w \leq 0.008856 \end{cases}$$

The f function has a cubic region that approximates to the constant contrast response we anticipated in section 4, and a linear region for very deep shadows where we expect the constant contrast model to break down. The function is smooth and continuous about the value 0.008856.

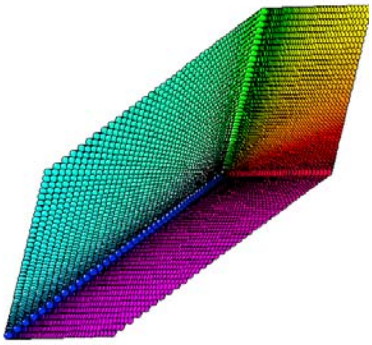

A video RGB colour gamut plotted in CIE XYZ

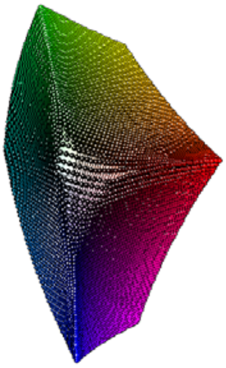

A video RGB colour gamut plotted in CIE LAB

For these diagrams, we made a list of all the RGB points with values 0,5…255 on the surface of the ITU rec709 video gamut, calculated the tristimulus values, then plotted them out in XYZ and L*a*b*. In both views, we are looking at the black corner of the gamut. The colours should be correct when viewed on a monitor with a gamma of 2.0.

The XYZ gamut has straight edges, because the red, green, and blue phosphors are independent. The light points are spaced out, but the dark points are all crammed together. This is a result of the ITU monitor gamma, rather than a property of the XYZ space. The L*a*b* gamut contains the same colours as the XYZ gamut, but we have expanded the dark regions where our colour resolution is good, and compressed the cyan region where our colour resolution is poor.

The luminance co-ordinate is L*. L* values can go from zero (black) to 100 (white). Negative L* values are errors. L* values beyond 100 should not happen if the reference white is chosen properly.

The chrominance co-ordinates are a* and b*. The reference white has zero a* and b*. You can get values of a* and b* greater than +100 and less than -100 for very saturated colours. The exact gamut depends on the reference white. The a* value is positive for reds, and negative for greens. The b* value is positive for yellows and negative for blues.

4.2 Delta E

ΔE is the difference between two colours as given by…

$$\Delta E^*_{ab} = \sqrt{\left((L^*_0 - L^*_1)^2 + (a^*_0 - a^*_1)^2 + (b^*_0 - b^*_1)^2\right)}$$

CIE L*a*b* space is based on measurements of just perceptible differences between colours. If any two colours are separated by a ΔE of 1.0, then you should just see the contrast between them under certain idealised viewing conditions. This is not wholly accurate, but it is close enough to be a useful measure of perceived contrast.

Because ΔE is based on just perceptible differences, it does not necessarily follow that it is a reliable measure of large differences. Two colours separated by a ΔE of 30 from a given colour may not necessarily seem to be equally close.

4.3 CIE L*u*v*

The CIE have another colour space called L*u*v*. It uses the same luminance value L as in L*a*b* (section 4.1) and the chrominance co-ordinates from Yu'v' (section 3.2). It also has its own ΔE measurement, which is similar, though not identical, to the ΔE of section 4.2. It does not assume a white point, and it is marginally easier to calculate. However, it is less good at predicting colour shifts than L*a*b*, and it can predict colours outside the spectral locus.

There are no good reasons for using L*u*v* these days.

4.4 Colour appearance modelling

We have described some of the 'classical' CIE colour spaces. These give us the ability to measure the differences between colours, and, to a limited extent, to predict the effects of changing the white point or the overall luminance. This is enough to allow us to simulate the appearance of a cinema image on a monitor in a darkened room with a similar brightness and surround. They cannot really be used to compare, say, the image we see in a cinema against the original daylight scene, which may have a thousand times the intensity, and no clear surround.

There are more general colour appearance models that cover these conditions. The most recent CIE colour appearance model at the time this was written is called CIECAM02. This came out in 2002, replacing CIECAM97. It is widely recognised that there are still significant aspects of colour appearance phenomena that are not described well. The colour differences are not always accurate, and spatial and temporal effects are not included. People have proposed ICAM, or Image Colour Appearance Modelling as an even more complex spatial and temporal colour standard.

The general colour appearance models are very complex, and they are still evolving. There are calls for the adoption of simpler standards such as ZLAB, that are less general but easier to use.

Truelight is used for comparing displays under similar viewing conditions. In the most critical grading conditions, the display and the reference image should be viewed in a darkened room, and both images should have similar brightness and white point, and the images are bright enough for normal colour vision. If we restrict ourselves to these sorts of viewing conditions, then Truelight has no real need to use colour appearance models. Truelight has a simple flare correction for less critical pre-grading work under office lighting conditions, but this is not intended for use when high colour accuracy is important.

In the future, we may use Truelight to match the appearance of the film image to the original scene. If we do this, we may have to cope with dramatic differences in brightness when rendering daylight scenes.

5 The Luminous White Point

The reflection colour white can be defined. An idealised white reflects all visible wavelengths. Practical whites reflect a similar fraction of all visible wavelengths.

There is no corresponding definition for the whiteness of a light source. We can see reflection colours by daylight, incandescent light, fluorescent lamps, or arc lamps. All of these light sources may seem 'white' to us when viewed in isolation, though they may seem coloured when compared to other lights. Our eye-brain system somehow identifies the dominant light source, and corrects our vision to compensate for the differences in illuminants.

We can see many colours as white given the right conditions. However, most illuminants that we normally regard as white lie close to the set of colours we get from incandescent sources, even though the physics behind the light sources are very different. The D55 and D65 daylight standards have very different spectra to the 5500K and 6500K Plank's law for thermal spectra, but their colour values are very similar.

Here is the thermal locus plotted relative to the spectral locus we saw in 1.2. In this diagram, we plot the colours with 6500K as neutral. Any point down to 2000K can appear white.

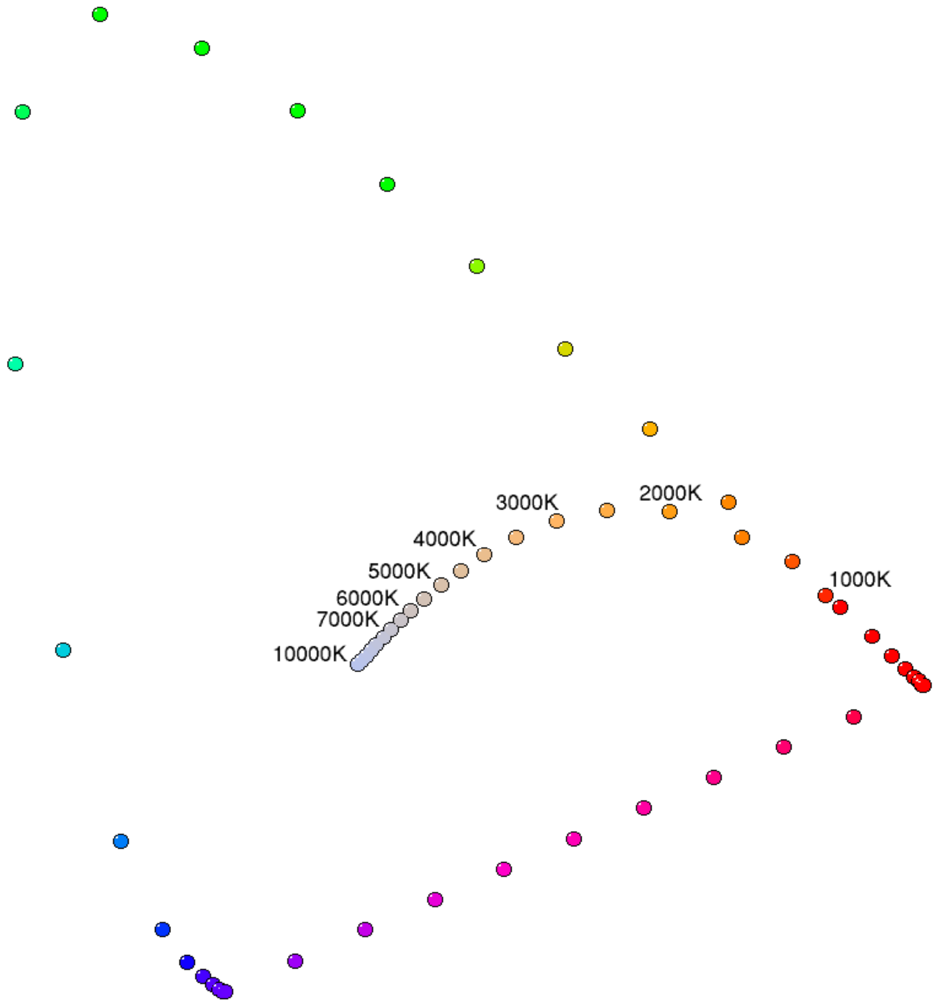
3000K
2000K
4000K
5000K
1000K
6000K
7000K
10000K

5.1 Colour temperature

Because most 'white' useful light sources lie close to the blackbody curve, it is common to quote light source colours in terms of the temperature of the nearest thermal point. Here are some typical values…

Light source	Temp
Old monitor white	9300K
Old monitor white	7200K
CIE D65 daylight standard	6504K
Video ITU rec 709 white	6500K
German cinema standard	5500K
Kodak slide viewer recommendation	5500K
Kodak Cineon view recommendation	5400K
US Cinema standard	5300K
European cinema standard	5200K
ICC print viewing D5000	5000K
Emerging art museum viewing standard	3700K
Steenbeck (incandescent)	3000K
CIE illuminant A	2856K

People often quote colour temperatures. This reduces our two colour co-ordinates to a single easily remembered physical value. Unfortunately, it does this by discarding the offset in the green-pink direction perpendicular to the thermal locus, so it is only meaningful when this offset ΔE is less than one (see section 4.2).

Colour temperature is also not a perceptually uniform measure. At the beginning of this section we plotted out points in 500K intervals. The 2000K point is well separated from its neighbours, while the 9500K point is overlapped by its neighbours. An infinite thermal temperature has a sensible blue-white colour according to Plank's law.

Sometimes colour temperatures are measured in mireds, or micro-reciprocal-degrees. A colour temperature of 2000K is 500 mireds. Measurements in mireds are more perceptually uniform. You can also get colour correcting filters with positive or negative mired values, which will, when added to the mired value of any thermal source, give the mired value of the filtered light. For example, a -100 mired filter will change our 2000K 500 mired source to a 2500K 400 mired source.

5.2 Thermal spectra and CIE daylight illuminants

A hot body will give off radiation. The spectral power distribution of this radiation will normally depend on the shape of the body, and the material that the body is made of, as well as the body's temperature. However, if we have a large cavity at a uniform temperature, and we look at the radiation through a small hole so we do not disturb the radiation field significantly, then the spectrum will depend on the body temperature alone. These spectra are variously called blackbody spectra, thermal spectra, incandescent spectra, cavity spectra, or Plankian spectra…

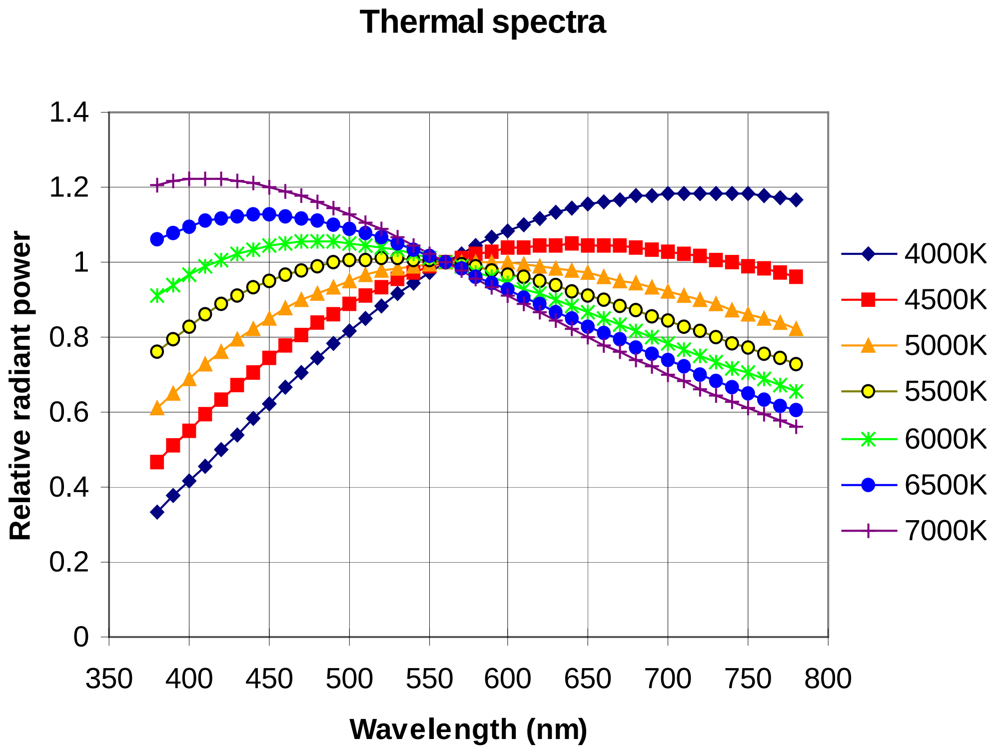

The CIE standard illuminant A is the thermal spectra of a 2856K body. Tungsten has a melting point above 3650K, so you cannot get any of the spectra in the graphs above directly from a solid body, though you can get good approximations using the mired filters described in the last section.

The CIE standard daylight illuminants are based on measurements of daylight spectra. The daylight spectrum was measured under a variety of conditions, and reduced to a mean spectrum, and two eigenvectors, which can then be used to calculate a whole family of spectra. These spectra are usually named according to the approximate colour temperature. D65, for example, is a Daylight illuminant with an approximate colour temperature of 6500K.

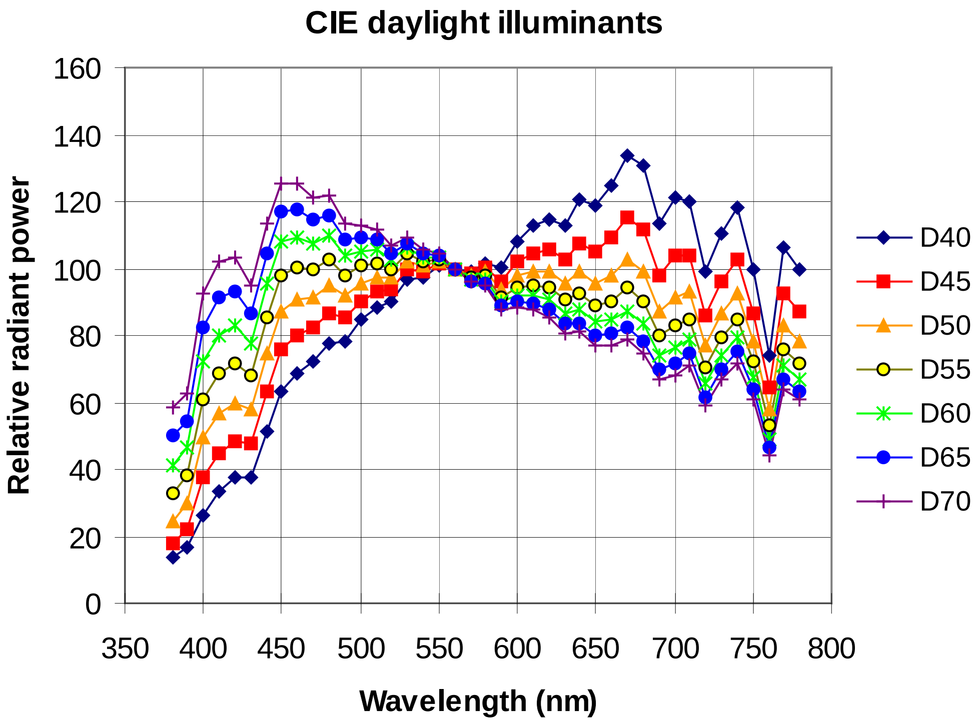

You cannot get a standard light source with any of these spectra: the complex and jagged spectral distribution is hard to reproduce in an artificial lamp. These CIE daylight simulation illuminants only exist as a theoretical ideal for calculations such as the colour rendering index described in section 5.4. When people describe a fluorescent tube as being D65, they probably mean the colour is matched to D65.

There is no reason why the CIE daylight illuminants should fall on the thermal locus. The points come close, but there are differences…

Thermal	x	y	CIE	x	y
4000K	0.3820	0.3792	D40	0.3823	0.3838
4500K	0.3620	0.3656	D45	0.3621	0.3709
5000K	0.3460	0.3532	D50	0.3457	0.3587
5500K	0.3330	0.3421	D55	0.3325	0.3476
6000K	0.3224	0.3324	D60	0.3217	0.3378
6500K	0.3137	0.3239	D65	0.3128	0.3292
7000K	0.3063	0.3164	D70	0.3054	0.3216

5.3 Psychophysical phenomena

In 1941, A.A.Kruithof published a graph (plotted below), which summarised the relationship between colour temperature, intensity, and the 'pleasant' quality of a light source. For example, a space will appear white when illuminated with 6000K light at 100 ft-lamberts, but will appear grey and gloomy when illuminated at 20 ft-lamberts. This is the sort of effect that makes a dimmed fluorescent tube look 'grey', even though 'grey light' should not be possible.

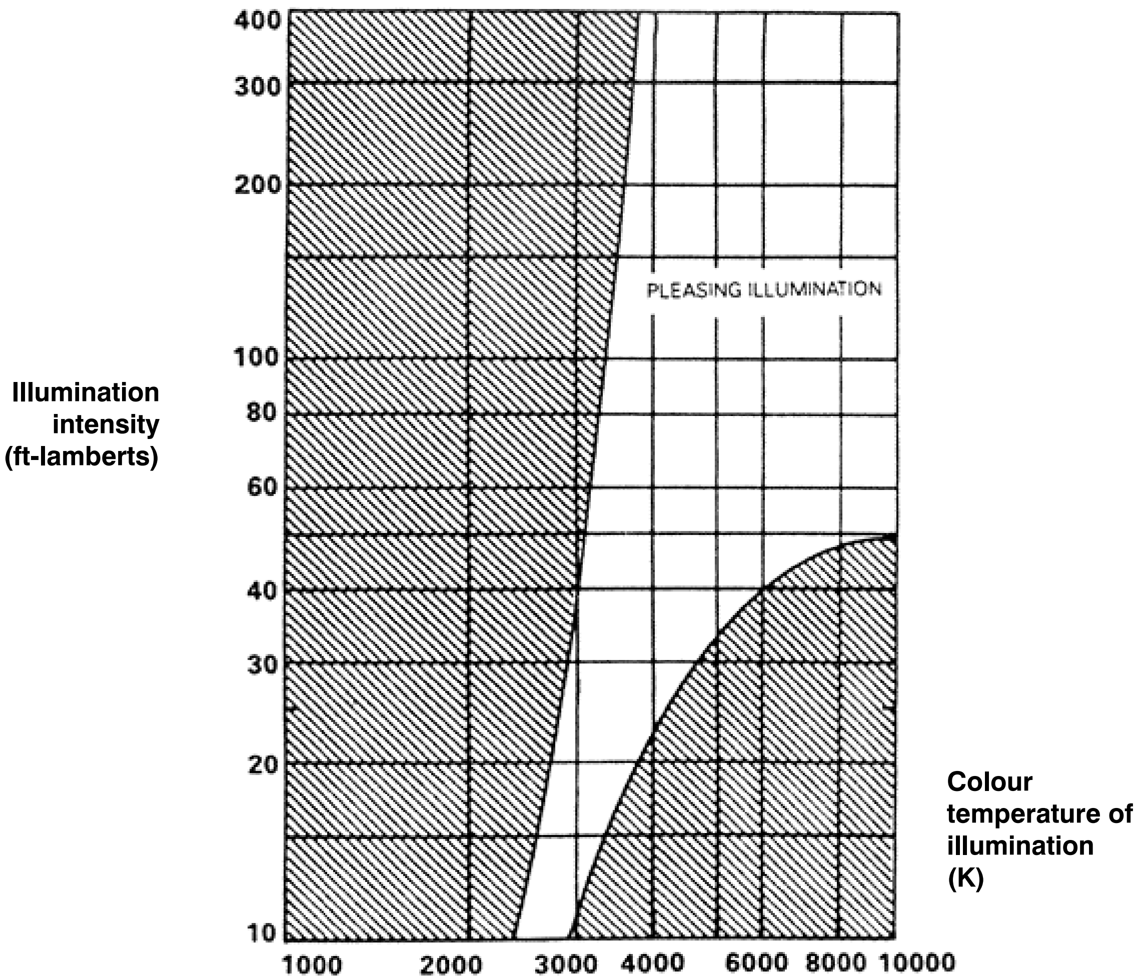

A cinema screen will typically have a colour temperature of about 6000K, and an open gate intensity of 16 ft-lamberts. If we assume the open gate white corresponds to the colour of a matt white object in the original scene (not necessarily true, but probably not too far out for many scenes on real films) then our view in a cinema would be like looking at the scene lit to 16 ft-lamberts. The Kruithof curve should predict that a typical cinema scene looks really gloomy and unpleasant.

A typical cinema scene does not look gloomy and unpleasant. Our eye-brain system may associate certain colour temperatures with certain illumination levels, but it is sophisticated enough to learn to reject these associations under special circumstances. The Kruithof curve is not wrong, but we have learned to override it when we are in a cinema.

There is a similar effect when viewing CRTs under office lighting. You could match a monitor to typical 3500K office lighting using instruments, but it would appear as orange as the 3500K dot in the diagram. Even standard 5000K lighting looks strange. We seem to know that CRTs have their own white - about D65 - and we do not expect them to fit the ambient white. I have not found this effect described in the literature, but it was at one time well-known in printing, where it was common to view images on monitors and compare them to printed images under standard lighting. As monitors are increasingly set up to the video D65 white, we might expect this effect to become stronger in the future.

There is not much we can do about these effects. They are hard to measure objectively, and they may be different from one person to another. We cannot consistently correct for the degree to which we 'know' that a CRT does not use ambient light, or 'know' that a cinema image is not a real scene. Like optical illusions, the best we can do is to know they exist, and to keep alert.

5.4 Colour Rendering Index

The D65 lighting standard was designed to match noonday ambient light. Typical D65 tubes will use five phosphors to give a similar, though not identical, power distribution from red to violet. A cheaper D65 light might only use three narrow-band phosphors. 6500K is beyond the melting point of tungsten, but we could match D65 using a negative mired filter. We could even match the white of D65 using just two wavelengths - for example red and blue-green. All of these illuminants would look the same colour; whites and greys would look the same; but colours would look very different. In the most extreme case - the two-wavelength white - a full colour image would appear to have only two primaries.

The CIE have defined a number, called the Colour Rendering Index. This is a measure of how well the light source will match an ideal incandescent light source. A perfect match (another incandescent source) would have a colour-rendering index of 100. A good illuminant such as our five-phosphor tube or a high pressure Xenon arc should have a colour-rendering index of 90 or more. A poor illuminant such as a high-pressure sodium lamp would have a colour-rendering index of 25 or less.

The colour-rendering index was designed to test the ability of the light source to render arbitrary colours. It is not suited to choosing a bulb for a projector that will pass light through three fixed film dyes. A lamp with a high colour-rendering index will probably behave a bit like a Xenon arc lamp, which also has a high colour-rendering index. However light sources with lower colour rendering indexes can actually make better projector lamps. A lamp with dark bands between blue and green, and between green and red will give brighter, more saturated colours with typical film dyes. Nevertheless, if you want to match the current industry standard, then a high colour-rendering index is probably a good thing.

6 Densitometry

A densitometer consists of a light source and a detector. A transmission densitometer will have the light source on the opposite side of the film from the detector. The instrument measures the transmittance of the film in density units:

$$density = -\log_{10}(transmittance)$$

A clear film has zero density; a film that transmits 10% of the light has unit density, and so on. The densitometer zero is set by making a measurement with no film.

Typical films scatter light, as well as absorbing it. In an ideal world, we would choose our densitometer optics to match our application. If we wanted to match the appearance of a typical f/2.0 projection system, we would want a f/2.0 light source and detector. If we wanted to measure the appearance of a film on a light box, we would want a diffuse light source and a collimated detector. In practice, many densitometers use a 90-degree (f/0.5) input and exit cones. This design makes the densitometer insensitive to light scattering artefacts such as fingerprints, but it does mean the densitometer will be picking up some scattered light that would not be seen in a projected image.. Fortunately, most colour film stocks do not scatter much light, so this does not affect the measurement much. Film greys made from silver grains with a hard high-contrast edge will scatter much more light, so you may have to take extra precautions when measuring black and white film.

The first film densitometers measured black and white material. Film greys made from silver deposits will generally have the same transmittance for all wavelengths, so we would expect different light sources to give similar density readings. However, the scattering can give a yellowish tints to silver images. If you wanted to measure the density of the image to estimate how it would print, then you would use a bluish light, because orthochromatic print stock is more sensitive to blue light.

Colour film can absorb different amounts of red, green, and blue light. We can measure the density to red, green, and blue light using three monochromatic densitometers with different light sources. Usually, it is more convenient to use a single instrument with one light source and three filtered detectors.

Before making colour measurements, the densitometer should be zeroed. If we are taking absolute density measurements, the densitometer is zeroed without film. If we are taking measurements relative to the film base, then the densitometer is zeroed on a bit of clear film base.

There are three ISO standard sets of filters, known as Status A, Status M, and Status T.

6.1 Status A colour densitometry

Colour film has cyan, magenta, and yellow dyes. A densitometer with status A (Analytic) filters is used for measuring how much of these dyes are present. They are filters with well-separated sharp peaks…

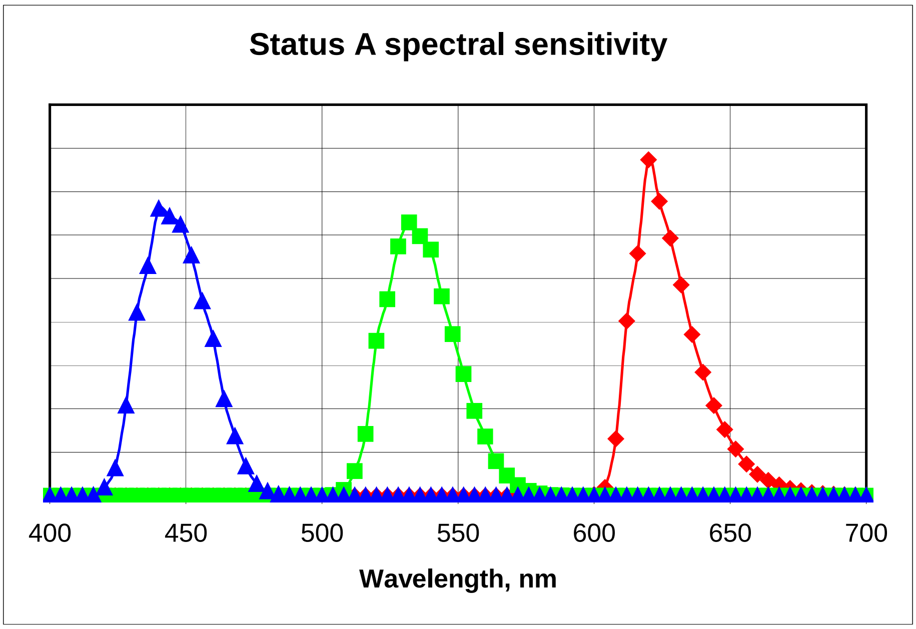

Status A measurements are designed for measuring combinations of cyan, magenta and yellow dyes. If you measure other materials, you may get misleading results. A material that absorbed 500nm (blue-green) or 600nm (orange) would give the same status A measurements as clear film, but might look very different. A neutral bit of colour film might have the same status A measurements as a piece of black and white film, but it would look lighter because combinations of cyan, magenta, and yellow dyes have transmission peaks at about 500 and 600nm.

Status A measurements can be made with a heat-absorbing filter. This makes the spectral sensitivity cut off faster as you approach the infrared. Under most circumstances this should make little difference to your measurements.

Status T filters are used to measure positive print densities for colour separations. The peaks are narrower, and more separated.

6.2 Status M colour densitometry

Status M measurements are useful for predicting how negatives will print. Modern colour film printing is done using an additive lamp house. An additive lamp house will have three light sources matched to the red, green, and blue sensitivities of the colour print stock. A densitometer that is matched to print conditions should have similar spectral peaks to the additive lamp house.

The status M spectral sensitivity curves are like this…

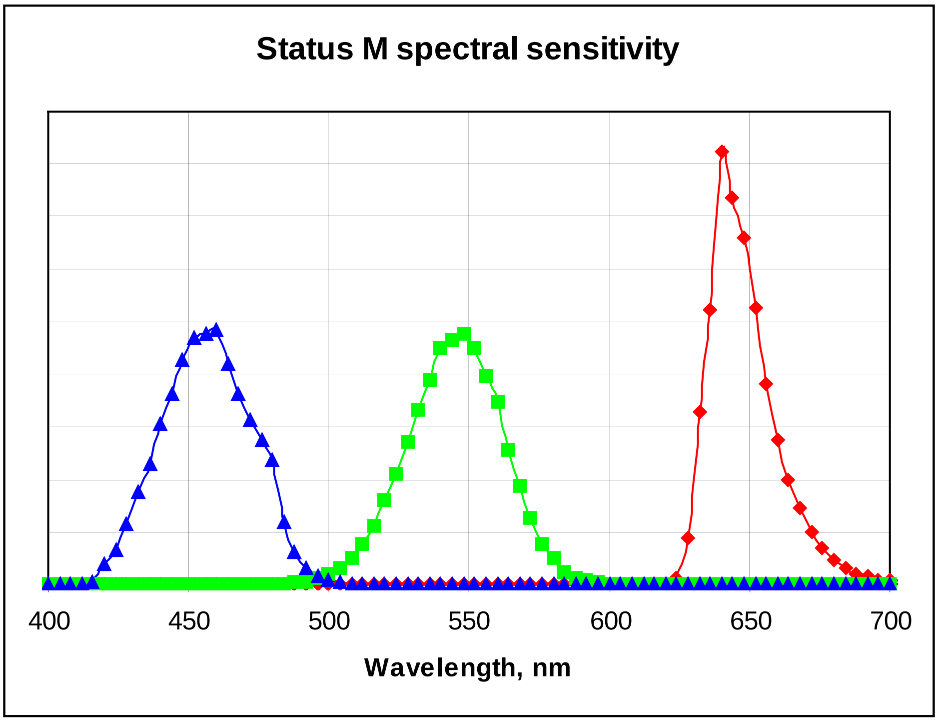

The ISO status M standard is used for most negative work, including motion pictures, still pictures, and reversal prints. The SMTPE RP 180 recommended practice paper describes a similar densitometry standard that is aimed specifically at the motion-picture printing industry. Unfortunately, there is not yet any equivalent to the tried and tested Xrite TR310 that supports RP 180.

7 Cineon Exposure Space

'Cineon' is a Kodak trademark for an image processing package, and an image file format. The image file format has a standard way of representing negative densities in 10-bit data that is been widely used in the film industry. For the rest of this section, we shall use the term 'Cineon' for this colour space.

You may think it very strange to define an image in terms of negative densities. There are many ways of printing an image. You can completely change the appearance of the image by changing the printing stock or the exposure. You can even get different densities on the negative when you use a different film recorder, as the Cineon calibration only is done along the neutral axis. However, Cineon space has useful properties for calibration, and it is consistent enough for some people to exchange images.

The most popular Cineon file format packs one pixel's RGB values into 10 bits of data. In this note we shall describe Cineon values as 10-bit integers from 0 to 1023.

The graph shows the relationship between Cineon value and negative density for an aim gamma of 1.0.

The graph is a straight line. Each Cineon unit corresponds to a density shift of 0.002 * aim gamma. Low Cineon values are light on the negative and dark on the print. Cineon RGB images viewed without colour correction on a monitor look like a washed-out version on the print.

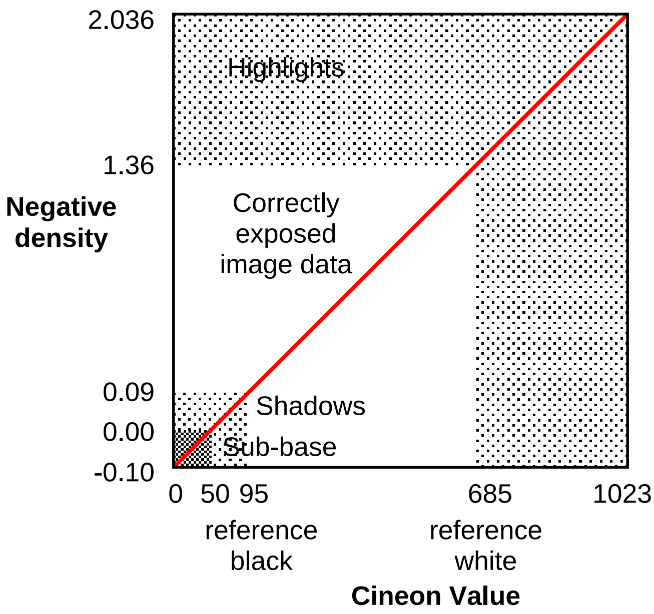

Kodak defines absolute RGB 'Dmin' densities for each of their intermediate stocks. In the Kodak documentation this is described as the blackest black that can be recorded onto that stock. It is approximately the black you get when you take a picture of a 1% black card. They assign the Cineon reference black value 95 to this Dmin density.

The Kodak formula for the absolute density is…

$$D_{absolute} - D_{min} = (\text{value} - 95) * 0.002 * \text{aim\_gamma}$$

We show the film base density as having a base density of 50. If you stick to the Kodak definition, then the actual Cineon base value will depend on the actual choice of film stock. The actual Cineon value of 5242 base stock is 42.

In practice, you may not have Dmin values if you are not using Kodak stock, or you may prefer to measure negative densities relative to base. We suggest a nearly equivalent alternative formula…

$$D_{absolute} - D_{base} = (\text{value} - 50) * 0.002 * \text{aim\_gamma}$$

7.1 Reference White and Reference Black

Cineon space covers all the useful density range for conventional negatives. A Cineon image usually contains more information than you can see in a single print. You can vary the exposure in Truelight and get more detail out of the highlights or shadows.

There is a subset of the Cineon space, with values between 95 and 685 that is commonly used for 'correctly exposed' image data. Kodak documents call 95 the 'reference black', and 685 the 'reference white'. Fixed conversions from Cineon to display RGB used to clip values outside this range. Old tube based film recorders also used to limit their output range to the 'reference white' to preserve tube life.

You are not restricted to the 'correctly exposed' range. Nevertheless, there are reasons for keeping most of your image within these limits…

- The negative exposure curve has its 'toe' beneath the reference black. Down at the toe, accurate interpolation is difficult, and the shape of the tone curve can vary with processing chemistry or the age of the negative film stock. You can avoid both these problems by adding 50 to all your Cineon values, which reduces the transmission of your negative by 20%, then increasing the exposure by 20% when you print.
- Most print stocks have a tone curve with the 'shoulder' above the reference white. This print shoulder has all the same problems as the negative toe. You can avoid the problems in a similar way.

7.2 Aim Gammas

If our Cineon step corresponded to a shift in negative density of 0.002 in all three channels, then a set of patches with equal values in all three channels would give a neutral negative. Unfortunately, a neutral negative wedge will produce a print with the darker colours looking distinctly bluish on Vision stock.

Kodak tabulate a set of aim values in their H-387 document "Kodak Digital LAD test image". If you plot out these values, they lie pretty close to a straight line. Kodak worked on the RP-180 densitometry standard mentioned briefly at the end of section 6.2, so aim gammas may have been a correction for the differences between RP-180 and status M measurement. Other people have used aim gammas as a correction for the colour shift of the last paragraph.

The Kodak aim gammas for Vision film stock are 0.966, 1.063,1.087. These are typical values, They are close to 1.0 and the red value is lower than the others.

7.3 Calibrating a Recorder

Usually, film recorders have to be calibrated before taking Cineon data. The recorder is loaded with the appropriate film stock. A set of neutral test patches is output. These test patches should cover at least the range from reference black to reference white. The densities of these patches are measured on a densitometer such as an Xrite TR-310. The measured density values are used to make a look-up table that converts Cineon values to the value to produce the correct density on a neutral patch.

Unfortunately, we need these look-up tables to generate the neutral patches on the test strip. The first time a recorder is calibrated, the usual approach is to generate a test wedge using some supplied set of look-up tables. These are unlikely to give a wholly neutral test strip, but the calibration will be an improvement on the original tables. After a couple of cycles, the test strip should be accurately neutral.

7.4 Setting up a Scanner

If you are setting up a scanner to produce Cineon data, you will have to chose your density origin and your gamma values. It might seem sensible to set up the scanner and the recorder so you can record out an image to film, scan it back in, and get back the same values. In practice, this is rarely a useful thing to do.

More often, the scanner will be scanning camera film. If you are digitally grading the whole film, then it is not important to match your scanner to your recorder, as any differences will be corrected in the digital grade. If, however, your final film will cut between camera original and recorded film, then you will want a print of your camera original to match a print of your output. If you have the Kodak recommended aim gamma values for your camera stock and your recorder stock, then you should use both sets of values: this should compensate for any systematic errors in densitometry due to the differences in dye and base spectral properties.

If you are not using the Kodak recommended aim gammas, then use the same aim gammas in your scanner and recorder. The differences between the two film stocks will probably be small.

The density offset is not an obvious choice. Our recorder base Cineon value may be about 50 but we can chose to avoid using values between 50 and 95 (see section 7.1). Some recorders cannot write data in this range because they are unable to lay down low enough exposures. The camera film, on the other hand, can have zero exposure, and may have very low exposure values in the shadows. Users often set their scanner up so the film base is the Cineon reference black 95 so their shadow detail will print without grading.

Switching Cineon 95 from Dmin to film base density may seem arbitrary. However…

- There will be a jump in base density anyway when cutting between camera and recorder negative.
- There is no equivalent for the grey slate image on the camera negative.
- The camera exposure was probably only accurate to half a stop, or about 4 printer points.
- There is usually plenty of headroom in the Cineon space for overexposures. Moving the whole image up by 45 Cineon values to give more room in the shadows is unlikely to destroy useful information in the highlights.

7.5 Exposure Units

Film labs have controls on their printer lamp house that set the RGB exposure. These controls are often marked in units of printer points, with 12 printer points to a stop. Increasing the printer point setting will increase the exposure through the negative, which will make the print darker. Sometimes people quote a figure of 8 printer points to a stop, instead of the usual 12. This is about right when referring to stops on the camera that took the original film, and assuming a typical camera film gamma.

Laboratories work in absolute printer points. The absolute printer point setting to get a given print result on a given stock will be different for each laboratory, and will change with the condition of the chemicals. People requesting prints from a laboratory will often ask for relative exposure shifts in printer points, and leave the absolute settings to the laboratory.

A shift of -1 printer point is equivalent to….
An exposure filter factor of 0.9439
An increase of +0.025 in the negative density.
A shift of +12.5 Cineon values
A shift of about -0.08 in typical mid-tone print densities

8 Video

The original video standards were for analogue equipment. They defined how analogue luminance and chrominance signals relate to red, green, and blue voltages, and how these voltages relate to light levels in the video image. The new, digital standards retain a surprising amount of this analogue heritage. Digital video engineers use digital equivalents of the traditional analogue tools, such as PLUGE tests and vectorscopes. Digital RGB signals can be converted to luminance-chrominance YCrCb, with luminance sub sampling to save bandwidth, just like the old analogue signals, then converted back again at the other end. In a long video pipeline you may have several such conversions, and the digital errors can accumulate.

8.1 Digital to Voltage conversion

In the standard, the red, green, and blue voltage signals are called R', G', and B'. The voltage is linearly related to the value in the digital standards. R', G' and B' have these ranges:

Limit	signal	V	8-bit value	10-bit value
Black point	0 mV	0	16	64
White point	700 mV	1	235	940

Digital video values can go outside this range. Video that is clipped to this range is known as 'legalized video'. Values outside the range 1-251 (4-1004) can give overflow or underflow errors depending on the hardware, particularly when converting to luminance-chrominance values (see section 8.4).

Values outside the legal range are not usually allowed on video media, but computer graphics images usually uses the whole 8- or 10-bit range. Some video hardware can compress the range of computer graphics image data when transferring to video, and expand the range of video data when transferring to computer graphics images. This re-ranging is often combined with the luminance-chrominance transformation described in section 8.4. At the time of writing, there are no clear standards on when data is re-ranged and when it is not. If there is room for doubt, you should check the values of a test image in the video and computer graphics formats.

8.2 Voltage to Luminance conversion

The voltage V ranged 0 to 1 is related to the corresponding luminance L for each of the R'G'B' channels using the formula:

$$V = 1.099 \times L^{0.45} - 0.099 \quad \text{for} \quad 0.018 \le L \le 1$$

$$V = 4.5 \times L \quad \text{for} \quad 0 \le L \le 0.018$$

This function is known as the 'opto-electronic transfer characteristic' or the 'gamma correction'.

In the standard, the voltages are called R'G'B', and the luminance values after the gamma correction are called RGB. Outside the standard, the luminance values are rarely referred to, and the voltages are often called 'RGB'.

This formula is simplistic. Over most of the range it approximates to the sort of power law function you get with CRTs, and at the dark end there is a linear section to keep the maths well behaved, a bit like the CIE function in section 4.1). Unfortunately, the linear section at the bottom end does not correspond with real monitor behaviour.

Real monitors are often set up with a PLUGE signal. A pluge signal has bars of –4%, 0%, and +4% grey. A digital 8-bit pluge image might have three bars with neutral values of 7, 16, and 25. The operator adjusts the monitor controls until the 7-16 edge cannot be seen, but the 16-25 edge just can. The 16 value is the video black value. If the monitor had the standard transfer characteristic, then this would set the luminance zero in the right place.

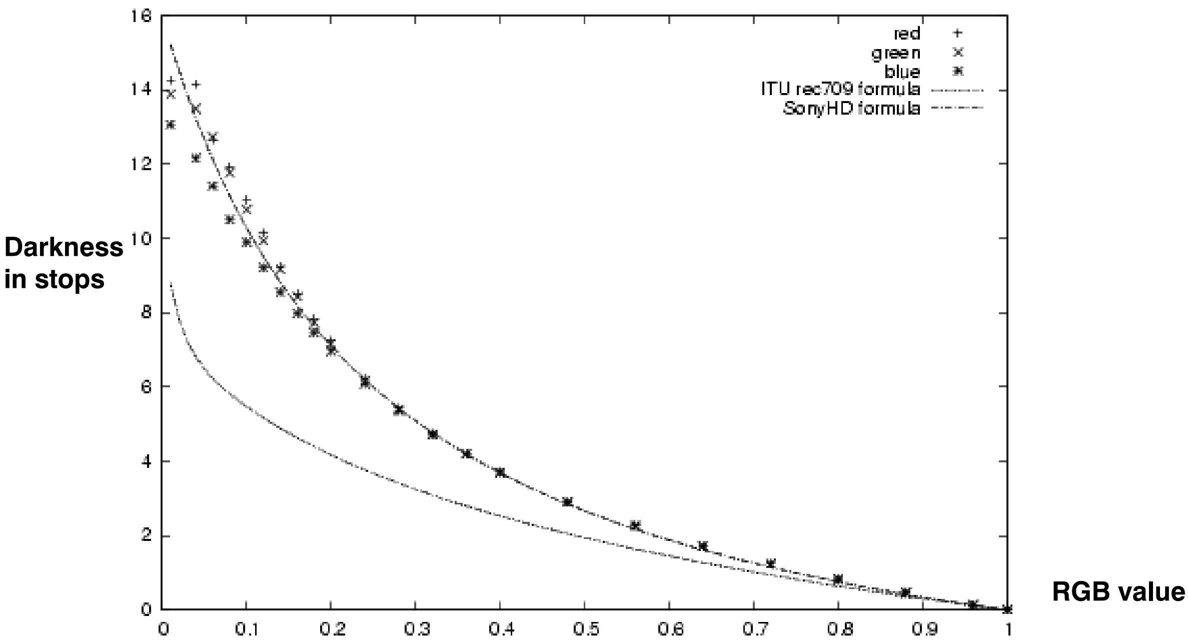

This plot shows experimental data measured off a Sony HD monitor calibrated in a darkened room using the PLUGE test image, the tone curve function taken from the ITU rec709 standard, and a fixed function we use in Truelight to represent Sony HD monitors.

Our measured values are typical of HD monitors. A new HD monitor will give a smooth curve that gets steeper towards the shadows. An older monitor, such as this one, has a background glow that stops our shadow range going beyond about 14 stops.

We might expect a calibrated HD monitor would fit the ITU standard. In fact, we find the ITU standard gives a curve that is very different to the measured one. We could increase the contrast on the monitor, and stretch the ITU curve so it fitted the top half of the experimental data, but the shape at the shadows is still very different. You might not see the difference if you are viewing with ambient lighting, but you will see the difference in a darkened room.

8.3 RGB colours

The standard defines the CIE chromaticity coordinates for the RGB primaries, and the R=G=B white:

Colour	x	y
R	0.640	0.330
G	0.300	0.600
B	0.150	0.060
White	0.3127 (D65)	0.3290 (D65)

The white chromaticity sets the relative brightness of the R, G, and B signals. If we normalize our white tristimulus Y - not video Y or Y' - to 1.0, then XYZ and RGB are related by…

$$\begin{pmatrix} R \\ G \\ B \end{pmatrix} = \begin{pmatrix} +3.242979*X & -1.538336*Y & -0.498920*Z \\ -0.968998*X & +1.875492*Y & +0.041545*Z \\ +0.505668*X & -0.204117*Y & +1.057698*Z \end{pmatrix}$$

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = \begin{pmatrix} +0.412135*R & +0.357675*G & +0.180357*B \\ +0.212507*R & +0.715350*G & +0.072143*B \\ +0.019319*R & +0.119225*G & +0.949879*B \end{pmatrix}$$

8.4 Luminance-Chrominance Coding

Within the standard, the luminance signal is called Y' and the chrominance signals are called C'_R and C'_B. Outside the standard, these signals are often just called Y, C_R and C_B.

There are two standards for luminance-chrominance coding in common use. The SD (standard definition) standard specified in SMTPE-125M is…

$$\begin{pmatrix} Y' \\ C'_R \\ C'_B \end{pmatrix} = \begin{pmatrix} +0.299 * R' + 0.587 * G' + 0.114 * B' \\ +0.511 * R' - 0.428 * G' - 0.083 * B' \\ -0.173 * R' - 0.339 * G' + 0.511 * B' \end{pmatrix}$$

When the current HD (high definition) standard was defined, these values were replaced for most HD standards by another set defined in SMTPE-274M…

$$\begin{pmatrix} Y' \\ C'_R \\ C'_B \end{pmatrix} = \begin{pmatrix} +0.2126 * R' + 0.7152 * G' + 0.0722 * B' \\ +0.5114 * R' - 0.4646 * G' - 0.0468 * B' \\ -0.1172 * R' - 0.3942 * G' + 0.5114 * B' \end{pmatrix}$$

There is a 1250 line HD standard that uses the old SMTPE-125M equations. Almost all the other HD standards, including all the popular ones, use the SMTPE-274M ones.

Note that both standards are calculating the luminance and chrominance from the R'G'V' voltage signals, not from the gamma-corrected RGB luminance signals. This was an engineering compromise in the analogue days that made the separation into luminance and chrominance inaccurate in some parts of the colour space. The C'_R and C'_B values are zero for R'=G'=B' neutrals, but for some colours the chroma will vary slightly with the Y' value.

The digital Y' values are ranged like the R'G'V' values of section 8.1…

Limit	Voltage	Y'	8-bit value	10-bit value
Black point	0 mV	0	16	64
White point	700 mV	1	235	940

The chrominance signals C'_R and C'_B are signed, and have slightly different ranges for the two standards. The digital values are also ranged slightly differently…

Limit	Voltage	C'_R, C'_B	8-bit value	10-bit value
Black point	0 mV	-0.5114	16	64
White point	700 mV	+0.5114	240	960

FilmLight
DIGITAL FILM TECHNOLOGY