https://doi.org/10.2352/issn.2169-2629.2019.27.38

Chromaticity coordinates for graphic arts based on CIE 2006 LMS with even spacing of Munsell colours

Richard A. Kirk
FilmLight Ltd, 14-15 Manette Street, London W1D 4AP

richard@filmlight.ltd.uk

Abstract

We construct redness-greenness (r,g) coordinates to fit the spectral locus into a triangle in the normalised CIE 2006 LMS plane. The reflection spectra for the Munsell patches for blackbody illuminants from 5000 to 6500 K appear as near circles in this space, suggesting that equal steps in (r,g) space may correspond to equal perceived colour contrasts within the gamut of reflective colours. We fit a matrix to convert from XYZ to LMS for reflective colours.

The motion picture industry needs a new colour space to support new high dynamic range and high colour gamut monitors. This space is intended for use with camera images and realistic scenes rather than simple test patterns. It should allow us to scale brightness and saturation without changing hue, so we can fit our images to our current display. It should feel perceptually uniform. Ideally it should work with any white point in the range 5000K to 6500K. It should be simple to implement, yet be flexible enough to allow the creative artist to exploit the new colours.

Any measured spectrum may be converted to cone cell response LMS values by a weighted sum. The LMS weights correspond to the response of the long-, medium-, and short-wavelength sensitive cone cells within the human eye. Two different spectra that give the same LMS sum should give the same cone stimulus, and hence look the same to a standard observer. This note uses the LMS weights for the CIE 2006 standard observer[1]; a similar construction should work for any similar observer model.

For typical photopic vision with reflecting objects in a scene, two lights with the same L:M:S ratio appear the same hue. This does not hold for all light levels, or for very simple, experimental scenes with few colour patches and no clear neutral stimulus to adapt to. This ratio hue matching is a part of ‘colour constancy’: a complex property where illuminated objects seem to possess ‘colour’, and the illumination is largely ignored[2]. The colour space described in this note is intended for use with realistic scenes, so we shall assume the seen chroma is determined by the L:M:S ratio, while recognising that this may not be so in all cases.

This ‘colour constancy’ is only possible when the scene can be interpreted, and an illuminant can be inferred. This happens in the higher visual centres, very far from the original LMS stimuli. The coordinate system in this note may model our ‘colour constancy’, but not our actual LMS signal processing.

If we convert our LMS values to normalised *lms*, where *lms* = LMS/(L+M+S) and $l+m+s = 1$, we may represent our colour gamut by a barycentric plot in an equilateral triangle, as in refs [3][4] and elsewhere. The CIE 2006 standard observer weights have a peak wavelength response of 1.0. This scaling has no particular significance when we look at scenes with broadband colours. It gives a small S value for common white points: for example D65 white with CIE 1931 Y = 1.0 has CIE 2006 LMS values of 1.070, 0.916, and 0.588. We could have chosen some other LMS scale, such as normalising LMS to equal-energy white, but we have not.

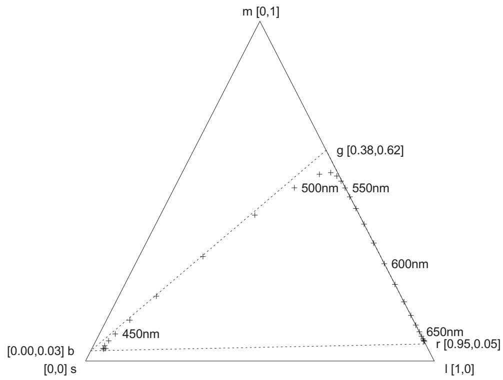

Fig. 1: The spectral locus at 10 nm intervals plotted in normalised [l,m].

In Fig. 1, the spectral locus between 400 nm and 700 nm is plotted in normalised *lms*. The fitted r,g,b triangle is not unique. The *r-g* edge is fitted closely by the long wavelengths. The *b-g* edge is a tangent to a curve. The *b-r* edge is determined by mixtures of the visible spectrum limits.

We take the visual spectrum limits to be 400nm and 700nm in this note. We see wavelengths above 700 nm but the locus reverses direction[5], so we get few extra colours. We can wavelengths below 400 nm, but a deep violet that was intense enough to see in a normally lit scene might do us harm. The colour coordinates change little with wavelength at the limits, so we may approximate a stimulus of 380 or 720 nm by a smaller stimulus of 400 or 700 nm.

The $[rgb]$ and $[lms]$ coordinates are related by…

$$\begin{pmatrix} l \\ m \\ s \end{pmatrix} = \begin{pmatrix} 0.95 & 0.38 & 0 \\ 0.05 & 0.62 & 0.03 \\ 0 & 0 & 0.97 \end{pmatrix} \begin{pmatrix} r \\ g \\ b \end{pmatrix} \quad (1)$$

We also normalise *rgb* in our barycentric *rgb* triangle, so $r+g+b = 1$. The $[rg]$ and $[lm]$ coordinates are then related by…

$$l \quad 0.95 \cdot r \quad +0.38 \cdot g \quad m \quad 0.02 \cdot r \quad +0.59 \cdot g \quad +0.03 \tag{2}$$

$$\begin{matrix} r & +1.0671 \cdot l & -0.6873 \cdot m & -0.02062 \\ g & -0.0362 \cdot l & +1.7182 \cdot m & -0.05155 \end{matrix} \tag{3}$$

The spectral locus fills almost 94% of the *rgb* triangle.

We could use the *l* and *s* primaries instead of having separate points for the bottom of the triangle. To fit the spectral locus, we must move the third corner to (0.345, 0.683). The *r-g* and *b-g* edges are now the LMS opponents, and the transform is simpler. However the spectral locus now fills less than 90% of the *rgb* triangle. We shall use our original three points that gave the best fill, but note we could get similar results with this construction.

The *rgb* values are linearly related to the *lms* values. We can use the same linear relationship to construct RGB spectral weights from the original LMS weights.

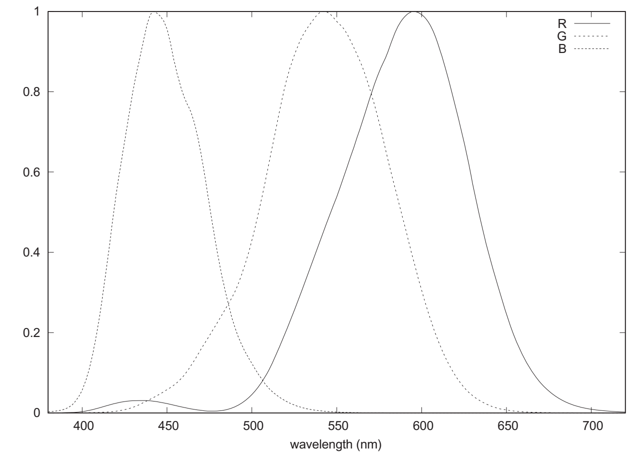

Fig. 2: The RGB primaries corresponding to the corners of the spectral locus bounding triangle.

These *RGB* weights are shown in Fig. 2. The ‘R’ weights have a small second peak in the short wavelengths, similar to the X weights of CIE 1931 XYZ but much smaller. The ‘G’ weights are very similar to the ‘M’ weights but with a sharper cut-off in the short wavelengths. The ‘B’ weights are identical to the ‘S’ weights. The *RGB* weights are similar to the raw RGB spectral response curves of some cameras.

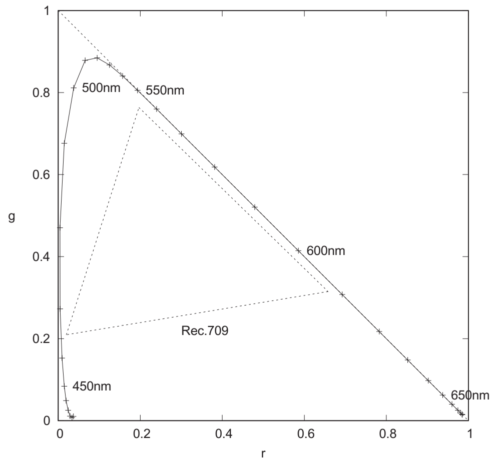

Fig. 3: The spectral locus plotted in (r,g)

We have not plotted (*r,g*) in Fig. 3 on an equilateral triangle, as we did with (l,*m*) in Fig. 1. Instead, the *r* and *g* axes are mapped on to the cartesian x and y. The diagonal dotted line is the *r-g* gamut limit: points above and to the right of this line are not real colours.

We have also plotted the triangle corresponding to the Rec.709 gamut. We are not particularly interested in video, but it is a fair match to Prof. Pointer’s gamut of real reflection colours [6], and it is less confusing when we plot it on top of other data.

The Munsell colour system[7] is a colour space with three parameters: hue, value (lightness) and chroma (colour purity or saturation). The Munsell Colour book has a set of coloured chips at regular intervals in hue, value and chroma, as judged by a human observer under normal daylight conditions. In Munsell’s experiments, ‘daylight’ may have been the now obsolete CIE standard Illuminant C, which approximates to 6700K. We probably get a good approximation to Munsell’s experiments if we viewed the Munsell Colour Book under 6500K illumination.

Munsell may be less perceptually uniform than other spaces such as OSA-UCS or CIECAM02, but it has a simple empirical basis, and it has been widely used in the graphic arts industries. Here we develop a simple colour space for graphic arts. We are aiming for the overall perceptual ‘feel’ of the Munsell space, rather than exact uniformity: if Munsell points have a similar circular arrangement in our space, this may be good enough for our immediate needs.

The University of Finland has a database of reflection spectra for the Munsell Colour Book[8]. We downloaded this, and calculated the CIE 2006 LMS values we would get for a 6500K blackbody illuminant.

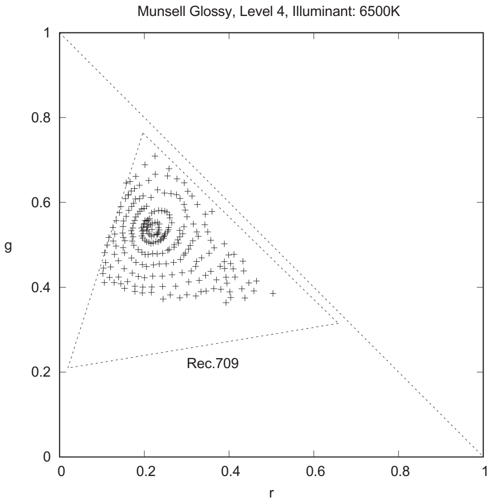

Fig. 4: Munsell patches with a 6500 K illuminant plotted in (r,g)

In Fig. 4, the reflected LMS for the Munsell level 4 – the level with the largest range of colours – is plotted in the new (r,g) space. The circles are the sets of colours with equal Munsell chroma. The different colours are distributed fairly uniformly about these circles. The other levels have less colour values, but give similar results. This suggests that equal spacings in (r,g) will approximate to equal colour contrasts for a human observer for at least the subset of colours covered by the Munsell Colour Book.

If we change to a 5000K illuminant, then the circles are slightly rounder. If we go to a typical incandescent white of 3200K (Fig. 5) then the circles are squashed the other way, suggesting that we are seeing less yellow-blue contrast, which is what we might expect.

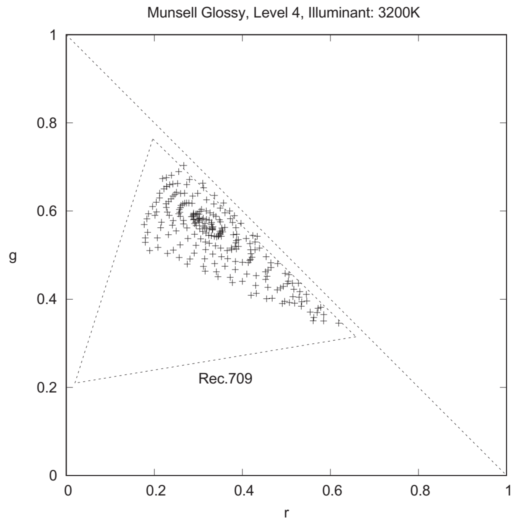

Fig. 5: Munsell patches with a 3200 K illuminant plotted in (r,g)

The (r,g) space is not intended as a renormalisation of the Munsell hue, value, chroma space. We are only using the Munsell data to show that the space is fairly uniform with respect to Munsell hue and chroma for typical viewing white points, from D50 standard viewing booths, through D60 cinema, to D65 displays.

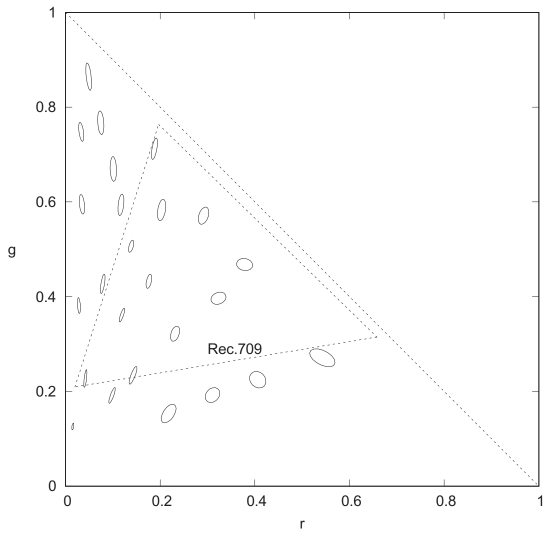

Fig. 6: MacAdam ellipses plotted in (r,g) at 5x scale

If the (r,g) space is perceptually uniform to a viewer adapted to a daylight-like white, it is unlikely to be uniform to just noticeable differences measured in an unadapted eye, as in MacAdam's experiments. When we plot the MacAdam 1942 ellipses[9], in Fig. 6 we see non-linear distortions, particularly along the r-axis. We do not know the exact spectra that MacAdam used, but I have calculated the smoothest spectra that fitted his data for this plot. As our space is intended for use with realistic scenes, and MacAdam's data is based on simple, static images viewed in a darkened room, it is better to fit the Munsell data than the MacAdam data.

Why do the Munsell points look equally spaced in (r,g) space? We designed (r,g) to compactly contain the spectral locus, but we did not explicitly design for this. There were several ways of representing the gamut. We could have used a barycentric plot, which displays the gamut corners equally. We could have plotted (r,g) on rectangular axes, and keep the origin for the out of gamut green. Plotting in (r,g) was one of several easy options, and it gave the longest side to the r-g colour opponent, which seemed right. We noted earlier that the scaling of the LMS values was also arbitrary. It seems this combination of decisions gives a useful set of coordinates.

The CIE-2006 standard[1] has a luminance value given by…

$$Y \quad 0.68990272 \cdot L + 0.34832189 \cdot M \tag{4}$$

With this, we can make a three-component Yrg space, similar to the CIE 1976 UCS Yu'v' space. We can transform from LMS to Yrg like this…

$$
\begin{array}{cc}
Y & 0.68990272 \cdot L + 0.34832189 \cdot M \\
a & L + M + S \\
l & L/a \\
m & M/a \\
r & +1.0671 \cdot l - 0.6873 \cdot m + 0.02062 \\
g & -0.0362 \cdot l + 1.7182 \cdot m - 0.05155
\end{array} \quad (5)
$$

The inverse transform is…

$$
\begin{array}{cc}
l & 0.95 \cdot r + 0.38 \cdot g \\
m & 0.02 \cdot r + 0.59 \cdot g + 0.03 \\
a & Y/(0.68990272 \cdot l + 0.34832189 \cdot m) \\
L & l \cdot a \\
M & m \cdot a \\
S & (1.0 - l - m) \cdot a
\end{array} \quad (6)
$$

The Y coordinate is not perceptually uniform. The CIE 1976 L*a*b* space was designed to be perceptually uniform. It's Luminance L* is generated from the CIE 1931 Y value. We can try the same CIE function on our CIE 2006 Y, and see how they relate to the Munsell value.

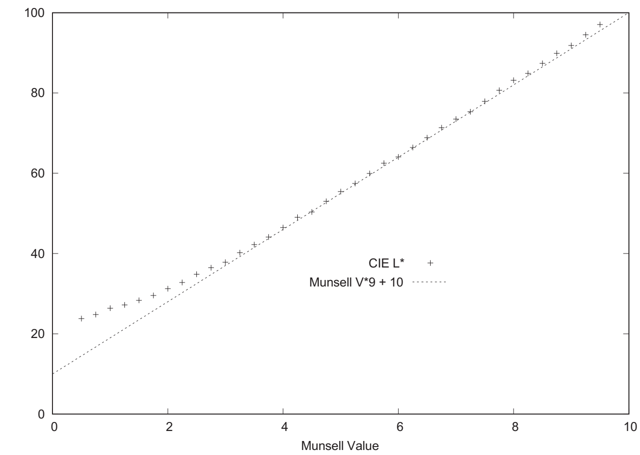

Fig. 7: A simple relation between CIE Luminance and Munsell Value.

In Fig. 7, a simple linear relationship maps the Munsell Value onto CIE L* over most of the range. There is a deviation at the dark end, for reflectances below 5%. This may be because the Munsell Colour Book is a set of reflectance colours, and we learn techniques for discerning gradations in shadow detail in real scenes and reflection prints, which may not apply to displayed images.

We are still not trying to renormalise Munsell space. If we were, we must reduce the L* value to less than a quarter of its current value for steps in L* to match the visual contrast of equal steps in (*r,g*). It is not clear whether this is useful, as L* and Y are familiar parameters. We could specify colours in L**rg* relative to some white luminance, or Y*rg* in absolute units.

These spaces are potentially useful to the graphic arts industries, but most people still work in CIE 1931 XYZ, and CIE 2006 LMS is unfamiliar to them. There can be no 3x3 matrix that converts from XYZ to LMS, as XYZ and LMS are based on different sets of spectral weights. However, we may be able to make a matrix where the worst errors lie away from the locus of reflective colours.

We have calculated a set of artificial reflection spectra. These are the smoothest spectra from minimising the sum of the squares of consecutive values, where the values are constrained to lie between 0.0 and 1.0, and to fit the target CIE 1931 (x,y) for the given illuminant. We calculated these spectra at intervals to span the set of real colours. With these we made an independent set of LMS and XYZ values, assuming a 6500K illuminant. We fitted a matrix to these values by minimising the RMS error between transformed XYZ and LMS. This gave large errors close to the spectral locus, particularly in the blues and violets. This is not surprising, as the CIE 1931 XYZ weights do not agree with more modern measurements of the eye's sensitivity to violet wavelengths. In general, the CIE 1931 XYZ standard observer is good enough for everyday use, but if extreme blues and violets are important, you should use CIE 2006 LMS.

We need a matrix that is optimised for the reflective colours, so we dropped the points from my calculation where the Y value was less than a fraction between 1%, 2% … 20% of the white point. The exact threshold did not make a lot of difference above a few percent. The 1% threshold gave good results...

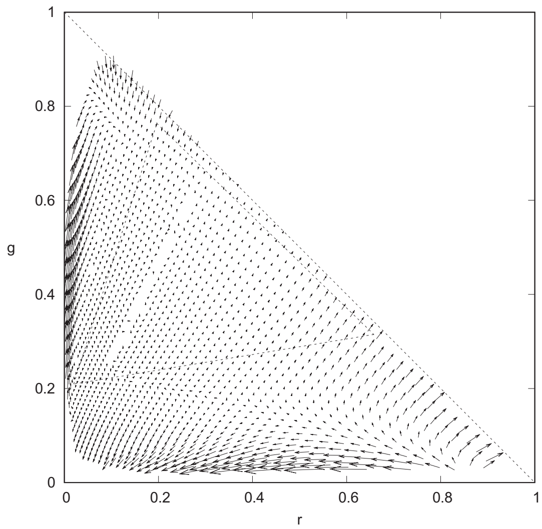

Fig. 8: XYZ to RGB matrix errors plotted in (r,g).

In Fig. 8 we plot the errors in our matrix calculation for our smooth, reflective colours with a 6500K illuminant. The tail of the arrow is the LMS was generated from the calculated CIE 1931 XYZ using the matrix, and the head is the calculated CIE 2006. We can see the errors are small in the locus of reflective colours as represented by the Rec.709 triangle. The matrix used for this plot was...

$$\begin{pmatrix} L \\ M \\ S \end{pmatrix} = \begin{pmatrix} +0.257085 & +0.859943 & -0.031061 \\ -0.394427 & +1.175800 & +0.106423 \\ +0.064856 & -0.076250 & +0.559067 \end{pmatrix} \begin{pmatrix} X \\ Y \\ Z \end{pmatrix} \quad (7)$$

We repeated the calculation using white points between 6500K and 3200K, and also used the Munsell glossy spectra, and did not see any significant errors. This is not surprising: if we replace one smooth spectrum illuminant with another, the reflection spectra are still smooth. There are no large errors on the r-g diagonal, so the 3200K white does not shift our points to a bad place. The major errors happen for very dim red-blue colours ($g < 0.1$) and very saturated blue-green colours ($r < 0.02$).

We got similar results when using the smooth spectra with the daylight D65 illuminant, which is not smooth. Again, this is not that surprising as the reflection spectra were smooth.

We cannot expect this sort of match for every illuminant and every colour. We can construct extreme examples where the illuminant and reflection spectra do not overlap at all. But if our illuminant and reflection spectra are fairly continuous, as they often are in real scenes, we may be able to live with the errors.

Conclusion

We have constructed a simple Y*rg* luminance-chroma space that fits well with visual colour contrasts as represented by the Munsell colour book for illuminants in the D50-D65 range. It has a triangular clip that approximates to the locus of real colours. The Y*rg* space may be a useful replacement for spaces such as Yu'v', and the approximate XYZ to LMS matrix may help the transition from 1931 CIE XYZ.

We need to extend our experiments beyond the reflection colours of the Munsell colour book. We could extrapolate points in (r,g) on a high-gamut display and see how far this Munsell-like uniformity goes.

Most motion picture grading uses RGB images normalised to a particular white: luminance-chroma spaces are only used in the user interface. Motion picture graders could use our *RGB* primaries normalised to a working white point for image processing, and Y*rg* for the user interface. The real colours in a normalised *RGB* space all have positive *RGB*, and the real colours fill almost 94% of the positive *RGB* space.

Y*rg* and *RGB* are based on CIE 2006 LMS, while most industrial colour measurement is still based on 1931 CIE XYZ. We have shown most reflection colours can be fitted well with a matrix transform from XYZ to LMS, with the largest errors close to the spectral locus. This can be used to transfer most images from XYZ-based formats. Most cameras and film stocks record broadband illuminated scenes well, but are not good at recording monochromatic sources. Only high gamut displays may need calibrating in LMS.

Thanks to Andrew Stockman for helpful suggestions when writing this.

References

[1] CIE TC 1-36 (CIEPO06). See also //www.cvrl.org.

[2] For a review of Colour Constancy see D. H. Foster, Vision Research, 51(7): 674 (2010)

[3] G. Wyszecki, W. S. Stiles *Colour Science 2nd Edition* (Wiley Classics, 2000) p.121

[4] D. I. A. MacLeod, R. M. Boynton, J Opt. Soc. Am., 69(8): 1183 (1979)

[5] G. S. Brindly, J. Physiol. 130: 35 (1955), graph also in ref. [3], p.435

[6] M. R. Pointer Colour Research & Application, 5(3): 145 (1980)

[7] //munsell.com/about-munsell-color/Munsell.com

[8] //www.uef.fi/en/web/spectral/munsell-colors-glossy-all-spectrofotometer-measured

[9] D. L. MacAdam, J Opt. Soc. Am, 32(5): 247 (1947), (x,y) plot in ref. [3], p. 308.

Author Biography

Richard Kirk has worked since 1985 on colour calibration and imaging technology, with Crosfield Electronics for the printing industry; with Canon for colour photocopiers, cameras, and desktop devices; with Framestore-CFC for film post-production; and finally for FilmLight, where he developed the Truelight colour calibration system, the standard for the film previewing and visualization. He holds over 30 patents on colour and imaging techniques, and an AMPAS Scientific and Technical Award (Oscar) for the Truelight system.