Gamut Mapping in Perceptual Colour Space

Lindsay W. MacDonald
Crosfield Electronics, Ltd., Herts, United Kingdom

Abstract

The accurate reproduction of coloured images requires not only an understanding of device behaviour but also of the differences in colour appearance between the source and destination media. These differences can be compensated by transforming the image data into a perceptual colour appearance space, but the problem still remains of how best to convert the gamut of colours in the original to make best use of the colour gamut of the reproduction device. Some shortcomings of existing gamut mapping algorithms are described and a five-stage transform for optimum image reproduction is proposed.

1. Introduction

The CIE system was originally defined to measure the visual equivalence of coloured lights, and was later applied to the measurement of surface colours and assessment of the colour differences between two colourant mixtures on similar media. It has proven very successful for comparing, for example, two dyes applied to a similar fabric or two paints applied to a similar substrate. Provided that similar constraints are imposed, the CIE system can be equally effective af0a metric for assessing the accuracy of a colour match in image reproduction.

Unfortunately there are several factors that can cause problems in the use of CIE methods to establish an exact colorimetric match:[1]

a) White point equivalence
b) Media differences
c) Measurement geometry
d) Gamut mapping

A reproduction will look exactly the same as the original image **only** if both have the same *XYZ* tristimulus values for the white point, the two media have similar surface characteristics and are observed under similar viewing conditions, and the reproduction medium can produce all the colours present in the original. In practice these factors normally vary considerably from one medium to another, making it impossible to achieve better than equivalent reproduction.[2]

A particular problem in the reproduction of colour images is that no photographic medium or device, whether colour film or phosphor display or ink printer or other, can produce the full range of colours that can be perceived by the human visual system. Every real device is limited to the gamut determined by its primary colourants. This gamut generally lies within the gamut of real-world surface colours[3] but differs for each device class: there are some colours that can be displayed but not printed and vice versa, and there are some real-world colours that can be neither displayed nor printed.

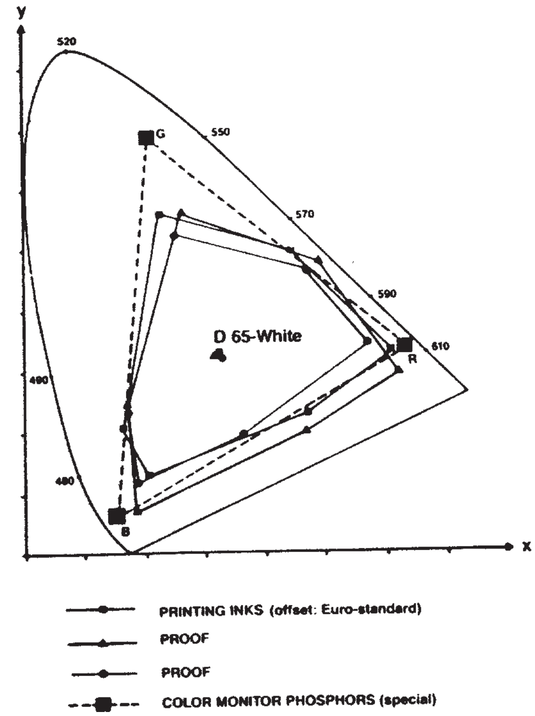

Figure 1. The colour gamuts of a typical phosphor display (triangle), printing inks and two colour proofing materials do not coincide.

Despite many years of colour reproduction practice there is little published data about how the compression of colour gamut from source to destination device should be quantified and in which colour space it should be performed. This paper describes some of the work we have done to clarify the issues and develop an improved basis for gamut compression in the reproduction of images.

2. Gamut Mapping

As a general rule, for good reproduction of colour a large gamut is better than a small one. The more saturated the primary colourants the greater the range of colours that can be produced by a device. Ideal colourants would be intense monochromatic lights for a display or 'block dyes' for printing, which have chromaticity coordinates near to the spectral locus, though in practice these are not achievable. It is also necessary, of course, that there are no 'gaps' within the gamut boundary, i.e. colours that ought to be producible by the device but for some reason such as quantising or dithering patterns cannot be achieved.

In practice, the gamut of common output devices such as display monitors and printers is considerably less than that of photographic transparencies or real-world scenes that might be captured through a digital camera.

The vital question is thus: how best can one modify out-of-gamut colours in an image to bring them into gamut? Various strategies may be followed:[4]

- Do nothing — The device clips each signal to its maximum value.
- Truncate — Colour replaced by nearest value on gamut boundary.
- UmbrellaLinear — scaling of all colours to fit inside gamut.
- Non-linear — Scale colours near boundary more than those inside.
- Image dependent — Let the user adjust controls to achieve best result.

A simple pragmatic method for monitor display is to move the chromaticities of the three phosphors toward the white point, simulating the use of a monitor with desaturated phosphors or a higher level of ambient light. Such a procedure is loosely justified for 'cheap and cheerful' applications by the argument that people often view images on different monitors anyway, without being aware that they differ in colour.

A more effective approach involving a three-stage transform was developed by Maureen Stone at Xerox PARC. Her objective was to define a sequence of transformations that preserve the image appearance whilst fitting the source (display) image gamut to the destination (print) image gamut. The process follows a number of widely accepted graphic arts and psychophysical principles:

(1) Grey axis of the image should be preserved
(2) Maximum luminance contrast (black to white).
(3) Few colours lying outside the destination gamut
(4) Hue and saturation shifts should be minimised
(5) Increase of colour saturation is preferable

The Xerox algorithm consists of converting each pixel of the image to *XYZ* tristimulus values, then mapping the grey axis via translation, scaling and rotating from source to destination device, and finally reducing colour saturation by an 'umbrella folding' of values around the neutral axis.[5]

A weakness of the Stone approach is that the gamut mapping calculations are performed in the *XYZ* colour space, which is not perceptually uniform. This means that for most devices (those that have convex or concave gamut boundaries) the compression may be higher than necessary at some hue angles. The relativity of compression is thus maintained but may not produce the optimum visual result, and indeed the authors accept this and suggest that improvements may be obtained by using a more uniform colour space such as CIELAB.

3. Colour Appearance

Two samples with different colourants may have the same *XYZ* values under one illuminant and therefore give a visual match, but under a different illuminant they may have different *XYZ* values and therefore not match, a phenomenon known as metamerism. The CIE system is an excellent tool for defining the similarity between two stimuli under the same illuminant, but though it can tell us that the colours no longer match as the illuminant changes it does not describe what the colours actually look like.

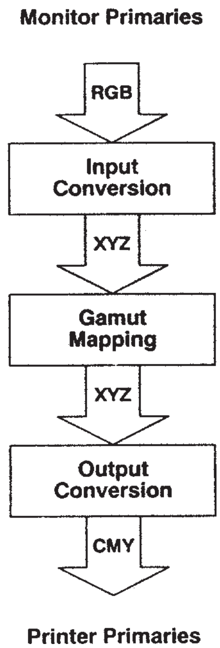

Figure 2. Stone's three-stage transform to convert images from monitor signals to printing ink values.

Surround also plays an important part in the appearance of colour. The simultaneous contrast effect is well known for a single coloured patch on a coloured background,[6] but applies also for a coloured border around a complex image. The effect of a dark border, for example, is to make the image appear lighter and less contrasty, whereas a light border makes the image appear darker and also less colourful. On a display monitor there are four concentric zones surrounding the central image area, all of which have an effect on the colour appearance of the image.[7] The result of this is that two of the same colour patches presented on identical media and with identical viewing geometry, even though they are measured to have identical *XYZ* values, will usually not match visually if the surrounds are different.

Over the past six years we have been engaged in a series of experiments to gather data on how colour appearance changes in different media and under different viewing conditions.[8] The media studied included displays, surface colours, photographic prints, dyed textile swatches, backlit large transparencies and projected 35mm transparencies. Variations in viewing conditions included a range of white reference values (from 2800K to 6500K), high and low luminance levels, white, grey and black backgrounds and white and black border surrounds. In each viewing mode, the observer was asked to estimate the appearance of the test colour in terms of the three perceptual attributes of lightness, hue and colourfulness.

The results of the experiments, which yielded nearly 100,000 observations, have been used to develop and refine a sophisticated mathematical model of colour appearance, which predicts the perceived lightness, hue and colourfulness of a coloured area from a knowledge of the *XYZ* colour stimulus and parameters defining the viewing conditions.[9]

The validity of various colour models was also tested by comparing each model's predictions with the mean visual result of the observers. The models selected in addition to the Hunt model were CIELUV, CIELAB, CMC and Nayatani, the first three of which are most commonly used by the colour industry to estimate colour difference. The Hunt and Nayatani models, on the other hand, were designed to estimate the change in colour appearance under different viewing conditions, and so were able to take into account all the parameters studied in the experiments.[8]

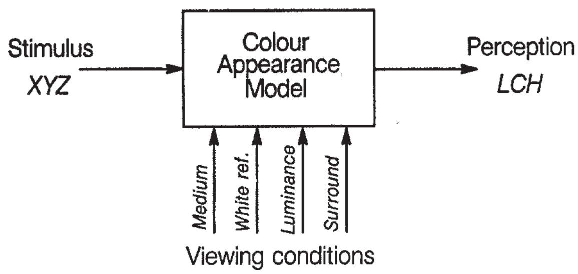

Figure 3. A colour appearance model converts the colour stimulus to perceptual coordinates, taking viewing conditions as parameters.

The performance of each colour model indicated that for the hue results Hunt's model is superior to Nayatani's model, with a mean coefficient of variation (CV) of 8 over all the experimental phases. For the lightness results, the Hunt model was clearly the best, and superior to the CIE L^* scale (ranked number two) by around 50%. All models gave a poorer fit to the results with white surround condition than the other types, implying that the perceived lightness against a white surround is quite different from that of the other surrounds studied. For the colourfulness results, the CMC and Hunt models performed slightly better than the others.

4. A Gamut Compression Experiment

The initial approach for gamut mapping was based on principles derived from experience in the graphic arts,[10] working in the perceptual coordinates of lightness, hue and colourfulness (chroma):

A. Maintain hue unchanged;
B. Compress lightness linearly, ideally taking into account the surround conditions;
C. Compress colourfulness (chroma) linearly.

Studies based on the above criteria, however, indicate that the approach has some limitations, in that it works well for some images but not others. It can be improved by mapping the distribution of colours in a particular image towards a centre point on the lightness axis (i.e. a neutral grey), obtained by averaging the minimum and maximum lightness values in the print.[11] The success of this method depends on the perceptual uniformity of the colour space in which the mapping is performed and on the similarity of the shapes of the colour gamut solids. It works best when the gamut solids have similar shapes and the print gamut is wholly contained within the original gamut. If parts of the print gamut extend beyond the original gamut, however, then the algorithm can be generalised to provide gamut expansion to take advantage of the more colourful printing inks at particular hues.

An experiment was devised to find out how graphic arts scanners and their operators typically perform optimum gamut mapping in image reproduction. A Kodak Ektachrome Q60A transparency[12] was taken to five different sites for reproduction, using two Crosfield scanners, two Hell scanners and one Scitex scanner. Two of these sites were in the UK and three were in the USA. Each operator was asked to scan the transparency twice, the first time using his 'standard' parameter setup and the second time with additional colour edits to optimise the reproduction.

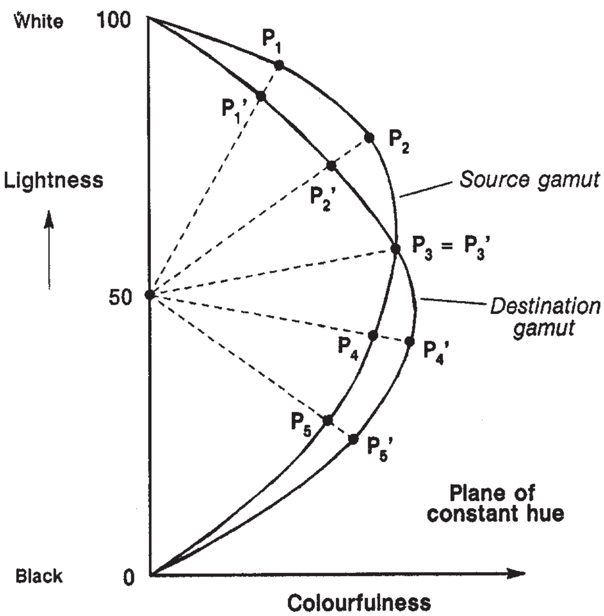

Figure 4. A method of gamut mapping in lightness and colourfulness, keeping hue constant.

The colour of each patch on the Q60A transparency was measured with a Bentham telespectroradiometer and the corresponding CIELAB coordinates computed under the D50 illuminant for the 2-degree observer. The ten resultant proofs were measured using an X-rite spectrophotometer and the same CIELAB coordinates computed and plotted against those of the original transparency. It was found that more meaningful results were obtained when the reference white was taken to be the lightest step of the grey scale, rather than the perfect diffuser.

The results of the experiment are complex and still being assessed, but preliminary analysis suggests that gamut compression in the graphic arts industry follows a generally similar set of rules, regardless of the type of equipment used for colour reproduction or the operator doing the work. It is clear that when the shape of the original gamut is not similar to that of the reproduction the above rules do not generally provide the desired result. For better reproduction the requirement to maintain constant hue in the mapping must be relaxed and the lightness of the chromatic colours must be modified differently to that of the neutrals.[13]

5. The Five-Stage Transform

Unless two images are viewed on identical media under identical viewing conditions it is unreasonable to expect anything more than equivalent colour reproduction (see Sec. 2). In general, an image that has been processed to look the same as it does on a different medium under different viewing conditions will have quite different *XYZ* colour stimulus values for each pixel. It is the overall visual effect that matters.

It follows that the simple three-stage transform from source device signals into a colorimetric colour space (*XYZ* or *LAB)* and thence to the destination device signals is limited in its applicability. In order to preserve the appearance of an image, the processing chain must also include a transformation into a colour appearance space, such as Hunt's *LCH.* In this space the gamut mapping can be performed in an optimal manner.

The resultant processing chain is shown in Figure 5, with the scenario of converting an image from *RGB* monitor signals to *CMYK* printing ink values. First the *RGB* image is converted into *XYZ* using the monitor profile data obtained from a prior characterisation. Then the *XYZ* data is transformed via the colour appearance model into perceptual *LCH* coordinates, using the parameters that define the monitor viewing conditions. Note that in this form the image is completely device-independent and represents the appearance of the picture on the monitor. The *LCH* image can then be modified, if necessary, to take account of the differences in colour gamut between monitor and printer, using one of the techniques described in Sec. 4 together with the both sets of device profile data. The modified *LCH* image is then transformed via the inverse colour appearance model back into *XYZ*, using the parameters that define target print viewing conditions. Finally, the *XYZ* image is converted to *CMY* ink values using the print profile data and a black ink (*K*) value is generated from *CMY* by means of rules controlled by parameters such as the degree of grey-component removal.

6. Acknowledgements

Thanks are extended to Tony Johnson at Crosfield and to Dr. Ronnier Luo at the Design Research Centre at the University of Derby for their encouragement in the writing of this paper and for their excellent work over the past six years in the two collaborative colour research projects carried out between Crosfield Electronics and Loughborough University. The experimental work was performed under the CARISMA project, supported by the U. K. Department of Trade and Industry under contract IED-1201.

References

1. Johnson A. J., Device Independent Colour ... Is it Real?, *Proc. TAGA Conf,* Vancouver, April 1992.
2. Hunt R. W. G., *The Reproduction of Colour in Photography, Printing and Television,* 4th Edition, Fountain Press, Tolworth U. K., 1987, pp. 177-197.
3. Holub R. and Xu L., Three Component Color Representations for Graphic Arts, *Proc. TAGA Conf,* Minneapolis, 1993.
4. Robertson P. K., Visualising Color Gamuts, *IEEE Comp. Graphics & Applns,* Vol. **8**, No. 5, Sept., 1988, pp 50-64.
5. Stone M. C., Cowan W. B. and Beatty J. C., Color Gamut Mapping and the Printing of Digital Color Images, *ACM Trans. on Graphics,* Vol. **7**, No. 4, Oct., 1988, pp. 249-292.
6. De Grandis L., *Theory and Use of Colour,* Blandford Press, Poole U. K., 1986, pp. 103-151.
7. MacDonald L. W., Luo M. R. and Scrivener S. A. R., Factors Affecting the Appearance of Coloured Images on a Video Display Monitor, *J. Phot. Science,* Vol. **38**, No. 4, 1990, pp. 177-186.
8. Luo M. R. et al, Quantifying Colour Appearance, Parts I and II, *Color Res & Appln,* Vol. **16**, No. 3, 1991, pp. 166-197.
9. Hunt R. W. G., Revised Colour-Appearance Model for Related and Unrelated Colours, *Color Res & Appln,* Vol. **16**, No. 3, 1991 pp. 146-165.
10. Johnson A. J., Perceptual Requirements of Digital Picture Processing, *Proc. IARIGAI Conf,* 1979.
11. Laihanen P., Colour Reproduction Theory based on the Principles of Colour Science, *Proc. IARIGAI Conf,* Eisenstadt, Vol. **19**, 1987, pp. 1-36.
12. Maier T. O. and Rinehart C. E., Design Criteria for an Input Color Scanner Evaluation Test Object, *Proc. TAGA Conf,* 1988, p. 469.
13. Johnson A. J., *Private Comm.,* Jan., 1993.

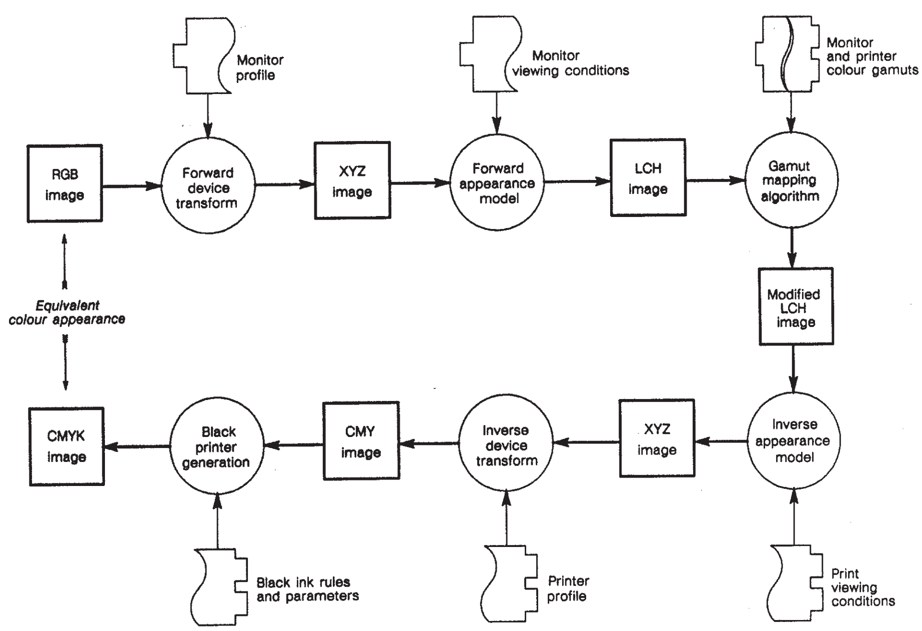

Figure 5. The five-stage transform is a comprehensive means of processing images to achieve equivalent colour appearance between two media.