Colour Normalisation and Spatial Comparison: A Robust Way to Integrate the Light Spectrum

Alessandro Rizzi

Dipartimento di Tecnologie dell'Informazione, Università degli Studi di Milano, via Bramante 65, 26013 Crema (CR), Italy
Email: alessandro.rizzi@unimi.it

In the physical world no colours exist, only light energy distributions with various spectral compositions. It is well known that there is a direct relationship between spectral power distribution and perceived colour when colour is considered out of context. If a colour is perceived in context, this relationship no longer holds. To integrate the rich amount of spectral information, in order to obtain the triplet necessary to generate colour sensation, the human retina integrates the spectrum by means of cone sensitivity curves. Colorimetry does the same using colour matching functions. Considering colour out of context, small changes in the integration curves result in corresponding changes in colour. In the case of colour in context this does not apply, since human vision colour normalisation mechanisms play an important role. This paper presents experiments that show how colour normalisation and spatial computation can compensate for changes in the integration curves. In other words, the human visual system has developed ways to adjust to changing lighting composition and distribution. Experiments prove that these mechanisms can also compensate for differences in cone sensitivities, or more generally in every set of spectrum integration curves.

Colour: Design & Creativity

Introduction

Our vision system receives colour stimulus from the detection of light energy distributions in their various spectral compositions. Light spectral power distribution (SPD) ranges over a wide set of wavelengths, from about 380 to 780 nm. In such a wide range, the number of possible SPD combinations is very large.

To deal with the extremely high light signal variability, our vision system integrates the light spectral information into a simple triplet, referred to as *tristimulus*, which corresponds to the three different types of retinal cones, at high light levels. Each cone type has different sensitivities, and integrates the spectrum with three different sets of weightings.

To obtain a final colour for colorimetry purposes, a set of colour matching functions (CMFs), or alternative integration curves, converts continuous or stepwise spectral information into three chromatic values. Much research in the field of colorimetry has been undertaken and several CMFs developed [1]. Tristimulus values vary considerably from one CMF to another. But what motivated this study was to find out if these changes are evident if the context in which colours occur is considered. With the test presented, the aim has been to analyse from a different perspective the robustness of colour perception, something that is usually considered only from the colour constancy point of view.

Moreover, from a physical point of view, the colour signal detected by our vision system derives from the interaction between spectral light distribution and surface reflectance, the percentage of incident light that, due to the surface characteristics, is reflected at each wavelength by objects. The colour information detected by our vision system is always the mix of the two: a white paper under yellow light is characterised by the same signal as a yellow paper under white light. It is impossible to separate these two components by relying only on the spectral signal at each single point [2].

To deal with this ambiguity, our vision system adopts spatial mechanisms and normalising behaviours. An important debate that is ongoing concerns how normalisation is performed by our vision system and whether it is based on global or spatial/local mechanisms [3]. In this paper, both these mechanisms are considered in the experiments presented. The basic goal of colour normalisation and spatial computation is to lower the effects of changes in light spectral composition and geometric distribution in the scene. Thus, interrelationships among colour signals appear to be more important than their absolute values.

Colour Perception Process in Context

CMFs have been proposed that consider colour as an isolated stimulus. Given that the spatial distribution of cones in the human retina varies greatly among different subjects [4], we would expect a corresponding difference in a subjective colour perception, but this has not been noticed. Observers with indistinguishable normal colour vision were shown to have very different numbers of long-wave and middle-wave cones. The ratio of long vs medium wavelengths cones can differ by more than 40% among different subjects without causing significant differences in colour perception and discrimination [4]. This suggests that some mechanisms operate to compensate for this variability, as suggested by many studies on visual perception [5-9].

A second element that underlines the role of compensation mechanisms in colour perception is the evidence that spectral sensitivities of the cones are very different from colorimetric CMFs. Cone sensitivities are highly overlapped and uneven [1]. The kind of colour triplets that can be synthesised in this way results in colours that have a poor contrast [10]; the maximum signal-to-noise ratio obtainable is about 2 and results in low colour saturation. However, this does not correspond to our everyday visual experiences when looking at scenes.

It is important to understand how the interaction of the spectral integration of cones (or CMFs) with colour normalisation or spatial computation results in perceptual robustness. Seen from another point of view, this study addresses the following question: how can colour normalisation (or spatial colour computation) decrease the effect of CMF variations on the final colour appearance? To answer this question, an experimental setup was devised to test the relationships between tristimulus triplets obtained using different CMFs and the change of these relationships when normalisation or spatial colour correction is applied. It was not the intention of the study to consider or compare different CMFs, or colour normalisation methods.

Testing Robustness

To test the robustness of tristimulus values against variation in the spectral integration curves, a multispectral image was synthetically generated to mimic the spectral light distribution of a real scene. Two scenes were computed with different light sources using a spectral ray tracer, which is a program that is able to reproduce the dispersion of light rays in a simulated environment. Then a set of different CMFs were applied in order to generate triplets of stimuli for each point, in the same way as our vision system (or an electronic device) obtains a colour signal at each point. Thus, the triplet differences through the various CMFs were compared with and without colour normalisation or spatial colour computation methods. This is described below and a fuller discussion of the theory behind the work has been published elsewhere [11,12].

Multispectral Image Generation

This study employed a synthetic scene, similar to the Cornell box [13,14], which contained a simplified Macbeth-like colour checker, with the same colour patches as the original. To simplify its geometrical description this was characterised by the same reflectances, but without the grey background separation.

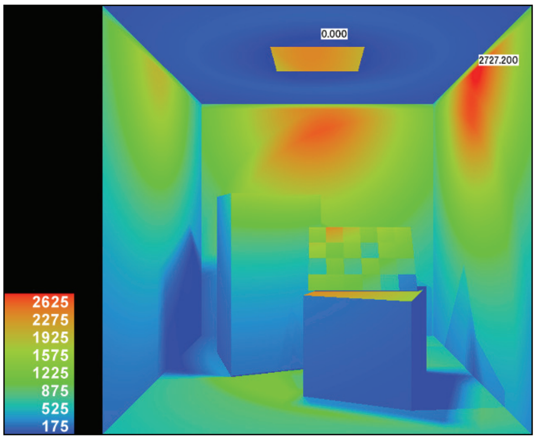

Figure 1 Luminance values in false colour of the synthetically generated multispectral image of a Cornell box with a Macbeth-like colour checker

Two sets of illuminants were selected: the first consisted of two D65 light sources, while the second comprising of one A and one C illuminant (subsequently referred to in this paper as A/C). In both cases the positions of the sources were the same: one in the centre of the ceiling and one in the top left corner, pointing to the opposite corner. In Figure 1, false colours demonstrate an indication of the luminance values of the scene.

To generate the synthetic images (608 × 608 pixels) we used the photometric ray tracer, developed by Rossi *et al.* mentioned above [15]. This sampled spectral luminance values in 80 frequency steps ranging from 380 to 775 nm at increments of 5 nm for each pixel. A global illumination algorithm was adopted to compute the interaction between spectral light distributions and surface reflectances, generating a spectral luminance distribution for each pixel.

Image Computation Pipeline

This approach to colour rendition was organised in the following sequence. To use the multispectral values to determine the corresponding tristimulus values, each pixel of the multispectral image was converted using spectral integration curves. Due to the intensity of the chosen light sources (Figure 1), the computer graphic method generated a high-dynamic-range (HDR) spectral luminance distribution. By applying different CMFs, the multispectral image was converted into a set of HDR RGB images. The last stage in the sequence was the application of a tone-mapping operator that converted the HDR image to the available dynamic range of the output device (a monitor).

Colour Matching Functions Chosen

Three different CMFs were analysed:

1. CIE RGB 1931 curves [16]
2. The CMFs proposed by Stiles and Burch [17]
3. A set of curves defined by Thornton [18].

The CIE RGB 1931 curves were chosen because they are the most used standard for RGB conversion. Stiles and Burch's curves are of interest because they contribute to the standard 10° CIE 1964 observer [19]. Thornton's CMFs were relevant to this study because they are a recent, highly debated example of a method of improving CMFs. A useful property of the chosen CMFs is that they produce RGB-like tristimulus values. Thus they can be used to compare and visualise differences between tristimulus values without the need for any further transformation. In Figure 2 the shape of the three sets of CMFs is presented.

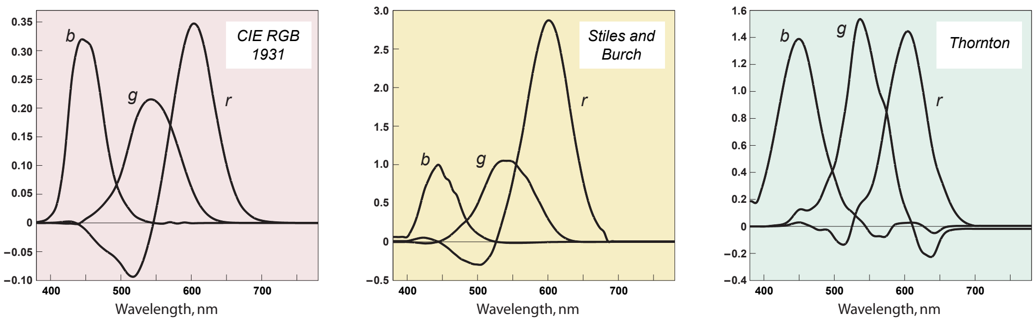

Figure 2 CMFs tested in the current study

Final Colour Processing

At the end of the image generating sequence two sets of results were generated (one for each illuminant configuration, D65 and A/C) for three different HDR RGB images (one for each CMF). These six images have been mapped into standard 8-bit RGB values using three different processing methods:

1. No normalisation or spatial computation
2. A global colour normalisation
3. A global and local spatial computation.

The global colour normalisation method was based on that of von Kries [9], while a retinex algorithm [7] was chosen for its characteristic of global and local spatial colour correction. There are many variants of the retinex model and multiple implementations [20]. In this instance, an original Land and McCann version was used with reset and Brownian path approximation by saccadic-like (i.e. related to the fast movement of the eye) steps, to compute ratios [21].

Methods 2 and 3 were subjected to a normalised output. In order to make comparisons possible, in method 1, in which no colour correction was performed, the triplets of unclamped floating point values were normalised with a simple logarithmic mapping into the same 8-bit range.

Logarithmic outputs were chosen, following the Weber-Fechner law. This choice was not critical, as a previous experiment with linear scaling reported similar results [22]. This was possible since no judgements on the colorimetric results or on the normalisation algorithms were performed in this experiment, as described in the next subsection.

Test Measures Goal

The objective was to test the robustness of the final colours against changes in the integration curve for each of the final mapping methods. More precisely, it was important to check how triplet differences, induced by modifying the CMFs, changed within each final processing result. To this end, for each of the two sets of images (illuminants D65 and A/C), the tristimulus values of the resulting low dynamic range (LDR) images, measured separately within each mapping method, were measured. No comparison between mapping methods, no single colour judgement and no comparison between CMFs were made. Figure 3 shows the test setup.

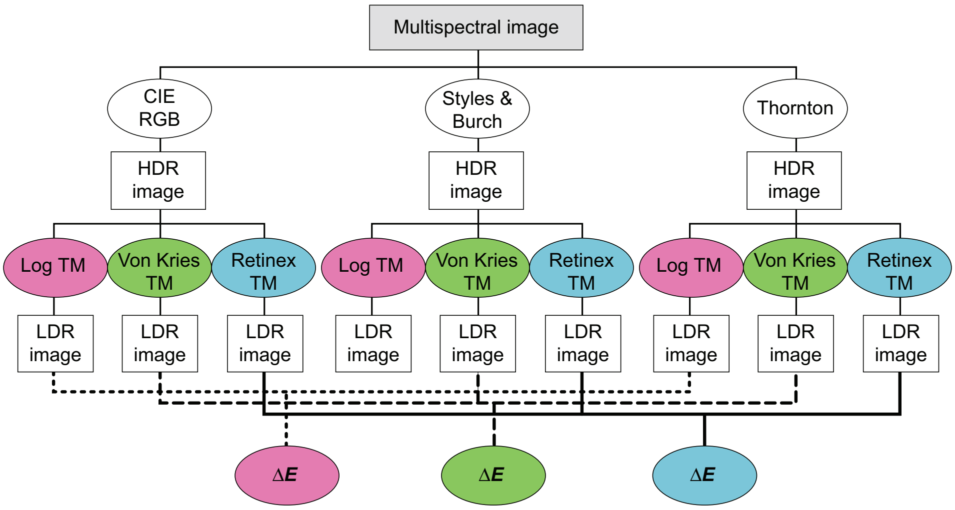

Figure 3 Experimental plan

As a measure of the difference between tristimulus values, we used Δ*E* in CIELAB space, computed considering each triplet as a point in the standard sRGB space, used in this case as a common conversion space. No colorimetric judgement was attempted and, since the purpose was to investigate the relationships between tristimulus values, this same computation was performed identically for all the images. To obtain only one number per image, Δ*E* was computed as the average distance of each corresponding pixel between two images. Separate measures for each colour-checker patch were taken, basically confirming the averaged data trend, published elsewhere [11,12,22].

Test Results

Figure 4 summarises the test results in a single diagram showing relationships, averaged over a whole image, between the three CMF pairs, for the two images under D65 and A/C illuminants. As the intention is to discuss the experiment in qualitative terms, only the overall results are presented; detailed data can be found elsewhere [11,12,22].

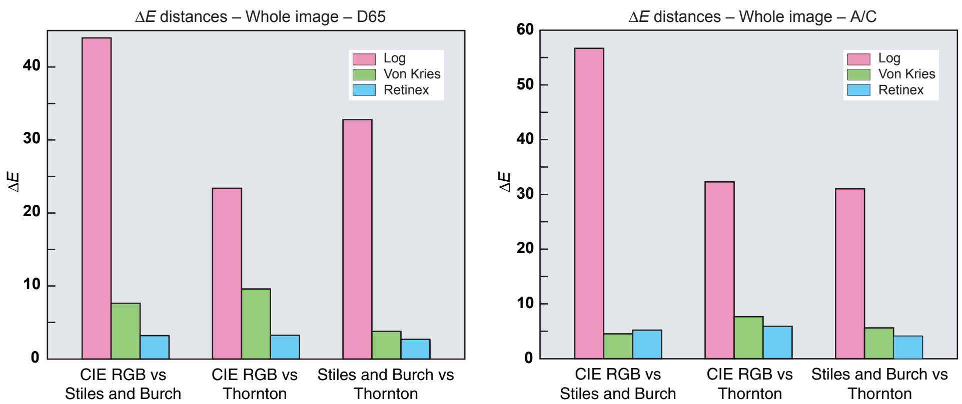

Figure 4 Overall test results showing relationships between the three CMF pairs

The principal evidence, shown in Figure 4, is that the logarithmic mapping, based on single pixels, exhibits a considerably higher differences with the two colour-in-context correction methods.

This difference was clearly visible in the resulting images [11], with the ΔE measures produced giving a quantitative idea of the size. It is important to emphasise that ΔE values cannot be interpreted as a classic chromatic difference, since they represent only differences between the CMFs, measured in a common RGB space that has no colorimetric meaning for the data considered. In fact, it is worth restating that the purpose of this study was not to compare the effectiveness of the methods nor of the CMFs to recreate a colour or its appearance.

These results show that colour normalisation or spatial colour corrections reduce differences between tristimulus values obtained using different spectral integration curves. This results in an increased robustness of the colour vision system in relation to changes in CMFs or cone sensitivities.

From another point of view, it is also evident that CMFs still play a role since the difference never reached zero.

Concluding Remarks

In everyday life we rarely observe colours isolated from their context, and the effects of the scene on colour appearance are well known and have been extensively studied. This is not considered in classical colorimetry that deals with colour in aperture mode, which is a colour that is observed through a hole, separated from any effect of the context in which it occurs.

Context plays an important role in everyday vision, yet context varies considerably, but in a complex and unpredictable way. The human vision system has developed a certain degree of robustness to deal with these changes. This robustness is not characteristic of a colour-measuring device that, on the contrary, has the need to keep the context stable and controlled.

In response to the question posed in the introduction, the rationale beyond all the existing CMFs, it is suggested that the concept of robustness is adopted as a key word by which we can analyse the performance of our vision system.

This article describes a preliminary test. Further investigations are required for a more general comprehension of the robustness of spectral integration in our vision system. However, results suggest that the mechanisms (and their robustness) of our vision system are strongly interconnected and can be analysed from the point of view of each one of the parts of which they are composed. The final perception is the interaction of all its components and not just the sum or their juxtapositions.

In other words, colour in context can be better explored if considered in its entire framework. If all the components are dissected and studied separately, and then reconstructed, this would result in a sort of a Frankenstein's creature. Context is not just something that changes colour appearance, but is the fundamental information through which colour is robustly extracted.

References

1. G Wyszecki and W S Styles, *Color Science: Concepts and Methods, Quantitative Data and Formulae* (New York: John Wiley & Sons, 1982).
2. J J McCann, S P McKee and T H Taylor, *Vision Res.*, **16** (1976) 445–458.

3. A Rizzi and J J McCann, Proc. IS&T/SPIE Electronic Imaging, San José, CA (2007).
4. H Hofer, J Carroll, J Neitz, M Neitz and D R Williams, *J. Neurosci.*, **25** (42) (2005) 9669–9679.
5. J Albers, *Interaction of Colors* (New Haven, CT: Yale University Press, 1963).
6. S Zeki, *A Vision of the Brain* (Oxford: Blackwell Scientific, 1993).
7. E Land and J J McCann, *J. Opt. Soc. Amer.*, **61** (1971) 1–11.
8. G Buchsbaum, *J. Franklin Inst.*, **310** (1) (1980) 1–26.
9. J von Kries in *Sources of Color Science*, Ed. D L Mac Adam (Cambridge, MA: MIT Press, 1971) 109–119.
10. E Land, *Sci. Amer.*, **237** (6) (Dec 1977) 108–128.
11. A Rizzi, D Gadia and D Marini, *J. Electronic Imaging*, **15** (2006) 041202.
12. A Rizzi, D Gadia, D Marini, M R Antognazza, S Perissinotto and G Lanzani, Proc. IS&T/SPIE Electronic Imaging, San José, CA (Jan 2008).
13. C M Goral, K E Torrance and D P Greenberg, *Computer Graphics* (Proc. SIGGRAPH 84), **18** (3) (1984) 213–222.
14. The Cornell Box. http://www.graphics.cornell.edu/online/box/
15. M Rossi, D Marini and A Rizzi, *J. Cultural Heritage*, **5** (3) (2004) 291–300.
16. *CIE Proceedings 1931* (Cambridge: Cambridge University Press, 1932).
17. W S Stiles and J M Burch, *Optica Acta*, **6** (1959) 1–26.
18. W A Thornton, *Col. Res. Appl.*, **17** (1992) 79–122; **17** (1992) 162–186; **17** (1992) 240–262; **22** (3) (1997) 189–198; (with H S Fariman) **23** (2) (1998) 92–103; **23** (4) (1998) 226–233.
19. CIE Proceedings 1963 (Vienna Session), Committee Report E-1.4.1, B (1964) 209–220.
20. E Provenzi, L De Carli, A Rizzi, D. Marini, *J. Opt. Soc. Amer. A*, **22** (12) (2005) 2613–2621.
21. D Gadia, D Marini and A Rizzi, Proc. CGIV04, IS&T 2nd European Conf. on Colour in Graphics, Imaging and Vision (2004) 326–331.
22. A Rizzi, D Gadia and D Marini, *Electronic Imaging* (SPIE International Technical Group Newsletter), **15** (1) (2004) 6–9.

