

Surface chromaticity distributions of natural objects under changing illumination

Yazhu Ling, Milena Vurro and Anya Hurlbert, Institute of Neuroscience, Newcastle University, Newcastle upon Tyne, UK

Abstract

The problem of colour constancy is ill-posed. In order to extract surface reflectance accurately from the received colour signal, the visual system must rely on pre-imposed constraints based on properties of the natural world. Here we investigate the surface chromaticity distributions of 7 natural objects under 3 illuminations (D65, CWF and F), using a characterized Nikon D70 SLR camera. We find that these object surfaces exhibit intrinsic chromatic textures and provide a large number of reflectance samples on their own. The information may thereby be utilized to improve colour constancy over that achievable with artificial surfaces possessing single or limited chromaticities. By analyzing the pattern of the chromaticity distributions under changing illumination, we find that the distributions of within-surface cone contrasts for given objects form distinct signatures in cone-contrast space. These signatures transform predictably under changes in illumination. We suggest that this feature may be utilized to aid colour constancy.

Introduction

The problem of colour constancy is ill-posed, because object surface reflectance and illuminant spectral power distribution are not uniquely separable [1,2]. In order to extract surface reflectance accurately from the received colour signal, the visual system must rely on pre-imposed constraints based on properties of the natural world [3].

Empirical studies of colour constancy typically employ 'Mondrian' patterns with a limited number of uniform colour patches. These stimuli are not representative of natural surfaces, which often possess intrinsic chromatic and luminance texture [4]. For example, a banana is neither uniformly coloured nor uniformly bright, whether it is ripe yellow or unripe green. Thus, colorimetrically, the banana's surface may be highly heterogenous, with substantial variation in its chromaticity and luminance at different spatial locations, even though we might perceive its surface as largely uniform.

Computational models of colour constancy demonstrate that the estimation of the illuminant spectral power distribution improves as the number of distinct surface reflectance samples increases [5]. Therefore, a single natural surface with intrinsic chromatic texture may provide a large number of reflectance samples on its own, and thereby undergo improved colour constancy relative to a surface with a single chromaticity.

To test the hypothesis above, we developed a method to record accurately and analyse the chromaticity distributions of natural objects, under various illuminations. We found that the distribution of within-surface cone contrasts for a given object forms distinct signatures in three-dimensional cone-contrast space. These signatures transform predictably under changes in illumination, and this behaviour may be utilized to aid colour constancy.

Methods

We employed a Verivide Colour Assessment cabinet, containing 3 independent, stable illumination sources (D65, CWF and F), and a Nikon D70 SLR camera with a 18-70mm kit lens, characterised to capture the chromaticity values of objects under each illumination. The characterisation process had two aims: (1) to obtain a set of RGB values measured by the camera and their corresponding XYZ values as measured by a spectroradiometer, under each illumination, and (2) to construct a characterisation model which summarised the relationship between the two sets of values. The details of the characterisation process are given below.

Before the measurements, we optimised the camera's settings under each illumination by placing a Macbeth Digital ColorCheck SG within the viewing cabinet, and adjusting the settings of the camera to optimise the histograms of the images. In practice, only the white balance and shutter speed settings required changing for each illumination (as shown in Table 1); all other settings were kept constant.

Table 1. The white balance and shutter speed settings under all illuminations.

	D65	CWF	F
White balance	Cloudy	Cloudy	Incandescent+3
Shutter Speed	1/13	1/6	1/6

We measured 125 colour patches in total, 124 from the Munsell colour book and one black patch from the Macbeth Digital Color Check SG because it approximates a pure black better than the darkest colour in the Munsell colour book. Of these, 88 patches were used as the training data set for obtaining the characterisation model, and the remaining 37 patches as the testing data set to test the model's performance. Each set individually spanned the full Munsell hue circle at varying values and lightnesses. Their descriptions are given in Table 2.

Each colour patch was then positioned at the same location within the viewing cabinet, and a PR 650 spectroradiometer measured its surface reflectances, while the camera recorded its RGB values from the same location. The patches' surface reflectances were measured only under D65; we then measured the spectral power distributions of the 3 test illuminations to compute the patches' XYZ values under all 3 illuminations. The process was repeated until all the colour samples' RGB and XYZ values were recorded under all 3 illuminations.

We then developed an individual characterisation model for each illumination, using a fifth order polynomial regression model. Under each illumination, we obtained a 3x35 polynomial regression fit which transforms the RGB values into XYZ values for the 88-sample training data set [6]. We then assessed the performance of the regression model by computing the colour difference between the measured XYZ values and the model-predicted XYZ values, for both the training and testing data sets.

Table 3 illustrates the model's performance under all 3 illuminations. For the training data sets under 3 illuminations, the mean colour differences between the model-predicted XYZ and the measured XYZ are below $1.5 \bullet E_{uv}$ units, just slightly above one JND (just noticeable difference). For the testing data sets, the mean colour differences between the model-predicted and measured colours are all below $2.5 \bullet E_{uv}$ units. Note that the polynomial regression weights were obtained only from the training data sets. The good performances we obtain for the testing data sets therefore indicate that the characterisation model predicts well the XYZ value for any RGB input, whether or not it was included in fitting the model. We therefore conclude that the characterisation models will adequately predict chromaticity values for RGB values in any image, under the 3 test illuminations.

Table 2. Munsell notations for the training and testing data sets.

88 training data set
5R: 3/2, 3/6, 5/2, 5/6, 5/10, 5/14, 7/2, 7/6, 7/10 5YR: 3/2, 5/2, 5/6, 7/4, 7/8, 7/12, 8/8 5Y: 3/2, 5/2, 5/6, 7/4, 7/8, 7/12, 8.5/4, 8.5/8, 8.5/12, 9/8 5GY: 3/2, 5/4, 5/8, 7/4, 7/8, 8.5/2, 8.5/6, 8.5/10 5G: 3/2, 5/4, 5/8, 7/2, 7/6, 7/10, 8/6 5BG: 3/2, 5/4, 5/8, 7/4, 7/8, 8/4 5B: 3/2, 3/6, 5/4, 5/8, 7/4, 7/8, 8/4 5PB: 3/4, 3/8, 5/4, 5/8, 5/12, 7/4, 7/8, 8/6 5P: 3/4, 3/8, 5/2, 5/6, 5/10, 7/4, 7/8, 8/4 5RP: 3/2, 3/6, 5/4, 5/8, 5/12, 7/4, 7/8, 8/6 N: 2, 3, 4, 5, 6, 7, 8, 9, 9.5 and Color checker black
37 testing data set
10R: 2.5/2, 5/8, 8/6; 10YR: 3/2, 6/8, 8/10; 10Y: 4/2, 7/4, 8.5/10; 10GY: 5/6, 8/6, 9/4 10G: 4/8, 6/4, 9/2; 10BG: 3/4, 6/6, 8/2 10B: 3/6, 6/10, 9/2; 10PB: 2.5/6, 4/6, 7/8 10P: 3/4, 5/10, 7/6; 10RP: 5/8, 7/4, 9/2 N: 2.5, 3.5, 4.5, 5.5, 6.5, 7.5, 8.5

Table 3. The model's performance for both training and testing data sets under all 3 illuminations, reported as the mean and standard deviation of all CIE ΔE_{uv} values between model-predicted and measured CIE XYZ coordinates.

	Training data set		Testing data set	
	Mean	Std dev	Mean	Std dev
D65	1.4888	0.8172	2.2760	0.8560
F	1.1602	0.7742	2.4669	1.5986
CWF	1.1017	0.6240	1.8723	2.1107

To measure the chromaticity distributions of natural objects under changing illumination, we selected 7 fruits and vegetables: banana, carrot, apple, clementine, plum, strawberry and cucumber, all common objects in daily life, each possessing a distinct surface colour and texture. Each object was placed in the viewing cabinet, at the same location where the training and testing colour samples were measured. The Nikon D70 camera then captured an image of the object for each of the three test illuminations, under each corresponding setup (as shown in Table 1). We therefore obtained 21 object

images in total. By isolating all coloured pixels within the object's contour, and computing their chromaticities using the camera characterization model, we obtained the chromaticity distributions of all 7 objects under 3 illuminations.

Results

Figure 1 demonstrates the chromaticity distributions of one object: a clementine, under 3 illuminations, illustrated in the CIE x, y chromaticity plane. (The luminance distribution (Y) of the object is excluded from this analysis, as the major sources of luminance variations for these objects are extrinsic, i.e. due to 3D shading and illumination geometry, rather than intrinsic, i.e. due to surface properties.) To exclude surface chromaticities which appear only very sporadically, and to improve the speed of the computation, only those surface chromaticities which occur at least 5 times within the object's surface are presented in Figure 1 and are analysed further in the following sessions. (Note that these are incidence distributions, in that each value is represented equally provided it has occurred at least 5 times.)

Figure 1 illustrates that a clementine possesses a large number of distinct surface chromaticities. Moreover, the chromaticity distributions of clementine vary substantially under changing illumination - not only do they reside in different areas of the chromatic diagram under different illuminations (see D65 vs. CWF in Figure 1), they also form clusters with distinct features (the shape of clusters under illuminant F is notably different from the shape of clusters under D65 and CWF). The same conclusions can be drawn for all the other objects, although only the results for the clementine are shown.

The finding that surface chromaticity distributions vary under changing illumination is not surprising. The chromaticity of a surface is affected by both the surface's reflectance and the spectral composition of the illumination. We would thus expect that the chromaticities of the same object change while the illumination alters. In fact, the fundamental challenge of colour constancy is to establish how we do maintain a relative stable perception of an object under changing illumination..

In this paper, we attempt to examine the distribution of an object's surface chromaticities as a whole, and investigate whether the distribution itself exhibits any constant characteristic which may be used to aid colour constancy. To achieve this goal, we first transform the original chromaticity distributions into cone contrast space [7], to simulate the early stage adaptation and encoding undergone by the visual system. Here we assume that the visual system is adapted to the grey background around the object (see red solid dots in Figure 1). (In this space, the RG cone contrast is calculated as the contrast in the (L-M) excitation with respect to the background (L-M) excitation, where L denotes the long-wavelength-sensitive cone type and M the middle-wavelength-sensitive type. BY cone contrast is the S-(L+M) contrast, where S denotes the short-wavelength-sensitive cone type.)

Figure 2 illustrates the RG and BY cone contrast distributions of a clementine under 3 test illuminations. Visual inspection of Figure 2 suggests that the cone contrast distributions of the object under changing illumination are more constant than the chromaticity distributions in Figure 1. The banana's RG and BY cone contrast distributions form a distinct pattern in cone contrast space, for which not only the size, but also the shape and position remain relatively constant, even though the illumination has changed.

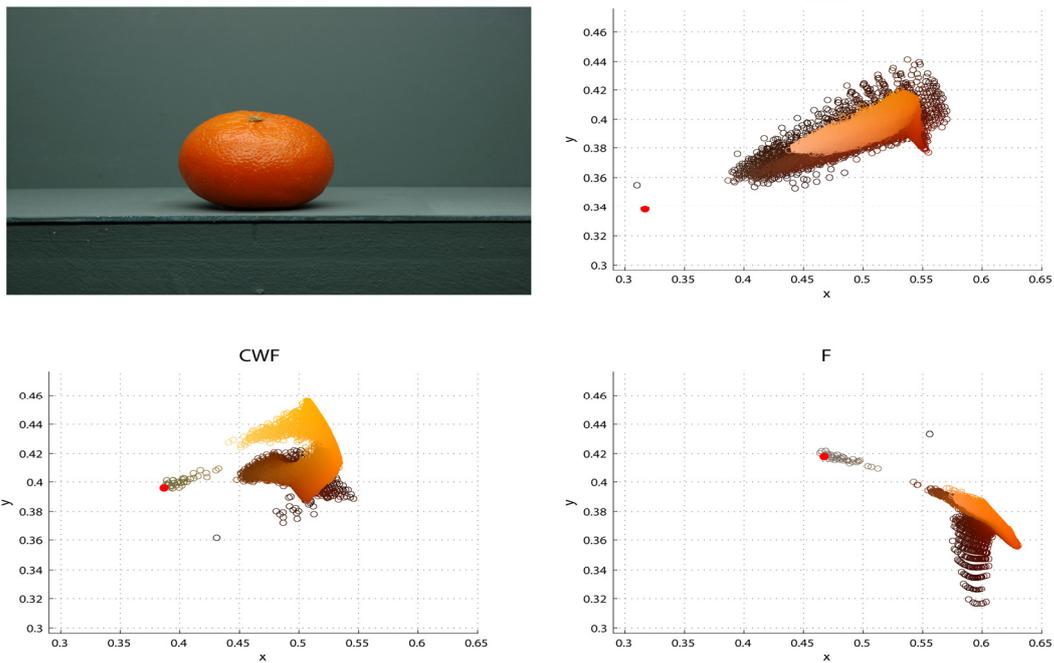


Figure 1. CIE xy chromaticity distributions of a clementine under 3 test illuminations. The coloured circles approximate the RGB colours in the original digital image. The red solid dots indicate the chromaticity of the background surface. Top left: original clementine image under D65; top right: the chromaticity distribution of clementine under D65 illumination; bottom left: the chromaticity distribution of clementine under CW illumination F; bottom right: the chromaticity of clementine under F illumination.

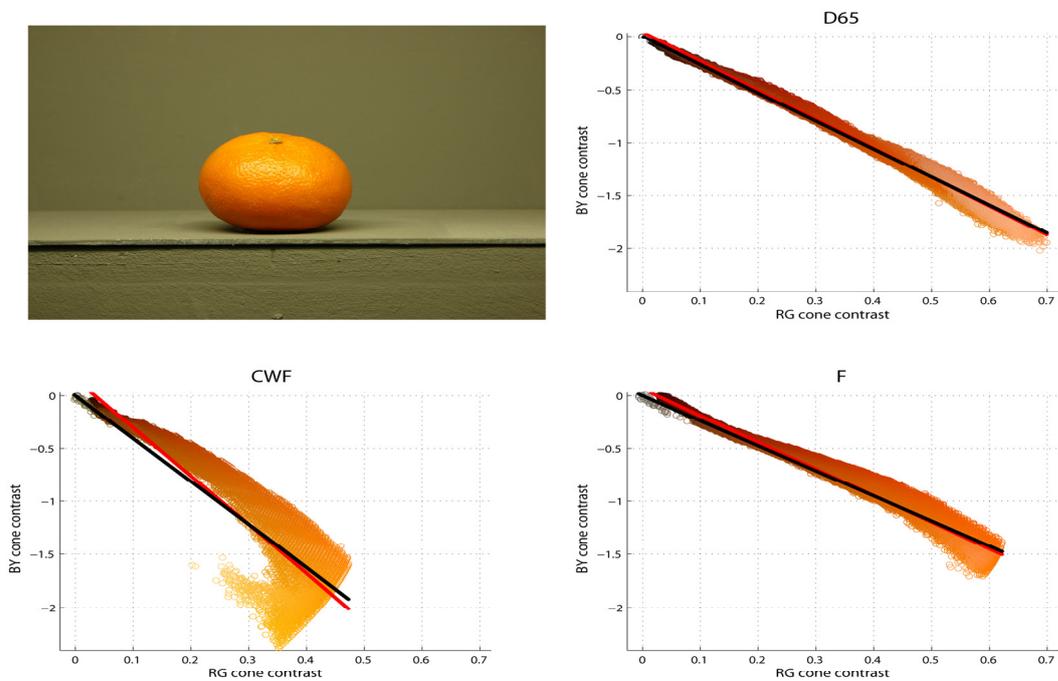


Figure 2. The RG and BY cone contrast distributions of a clementine under 3 test illuminations. The coloured circles illustrate the RGB colours in the original digital image. The black solid lines indicate the fitted main hue angles; the red solid lines indicate the fitted major distribution axes. Top left: original clementine image under CWF; top right: the cone contrast distribution of clementine under D65 illumination; bottom left: the cone contrast distribution of clementine under CWF illumination; bottom right: the cone contrast distribution of clementine under F illumination.

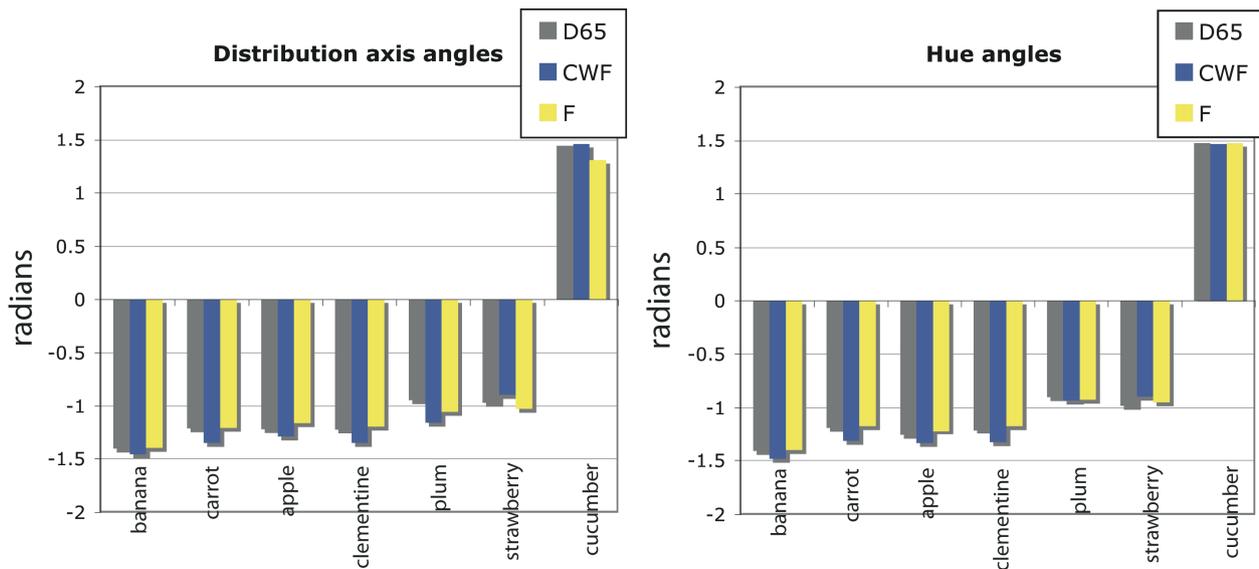


Figure 3. The major distribution axis angles and hue angles for all 7 objects, under 3 test illuminations.

We thus attempt to extract a simple numeric signature – a quantitative representation of the basic features of the distributions. We note that for almost all 21 images, the objects' cone contrast distributions form simple narrow shapes which are well represented by vectors describing the major and minor axes of the shape. In addition, the mean cone contrast of each distribution remains roughly constant. We therefore test two candidates for signatures, each with its own perceptual implication.

The first signature we test is the angle of the major axis of the distribution in cone contrast space, which characterises the primary direction of the cone contrast distribution. This signature is obtained by finding the best fits for p_1 and p_2 in a first order linear model, as shown in Equation 1:

$$BY = p_1 * RG + p_2 \quad (1)$$

where RG represents the RG contrasts of the distribution, and BY the BY contrasts of the distribution. The signature is then computed as the arctangent of p_1 , where p_1 and p_2 are the best fits to Equation 1 for a surface distribution.

The second signature is the mean hue angle of the distribution, which represents the mean vector of the distribution with respect to the neutral point ($[0, 0]$ in cone contrast space). This signature corresponds roughly to the definition of hue, and may correlate better with the perceptual hue observed for the object. The signature is obtained by taking the arctangent of the best-fitting p_1 for Equation 2:

$$BY = p_1 * RG \quad (2)$$

The solid red lines in Figure 2 demonstrate the fitted signature 1 (the distribution major axis angle) for the clementine, under 3 illuminations. The solid black lines in Figure 2 show the fitted signature 2 (the hue angles) for the clementine, under all illuminations. Although the red and black lines in Figure 2 are similar, they do not overlap completely, indicating the difference between measures. Nevertheless, across illuminations, the directions of the fitted vectors do not vary much, for both signatures.

Figure 3 illustrates the measured distribution axis and hue angles for all objects, under 3 illuminations. Despite changes in

illumination, both signatures remain relatively constant, and are not significantly affected by illumination (Two way ANOVA analysis: distribution axis angles: $F(2, 20)=1.9$, $p=0.1912$; hue angles: $F(2,20)=3.49$, $p=0.07$). We therefore hypothesize that these constant signatures may be utilized by our visual system to aid colour constancy.

Discussion

Empirical studies of colour constancy typically employ simple colour stimuli which contain limited numbers of chromaticities. In this paper, we have shown that in the natural world, even simple objects such as fruit or vegetables possess capacious amounts of surface information. Therefore, the visual system is consistently processing complex chromaticity distributions, rather than individual surface chromaticities. The theoretical implications of this characteristic of natural surfaces has largely been overlooked by previous studies, but it may play an important role in colour constancy.

In this paper, we consider the chromaticity incidence distribution of a natural object as a whole, and investigate how the distribution changes under varying illumination. We find that the cone contrast distribution of an object remains nearly constant under changing illumination. This finding prompts us to extract simple numeric signatures to represent the fundamental characteristics of a complex distribution, and discuss the possibility that these signatures may play a role in colour constancy.

This paper is a preliminary study to propose a feasible candidate for colour constancy. We do not intend here to verify the proposed signatures, or to examine when or how are they employed by the visual system. Nevertheless, given that we often perceive familiar objects as largely uniform in colour, and assign unitary "memory colours" to them [8], it is likely that the visual system employs some kind of algorithm to extract simple representations from complex chromatic distributions.

Whichever algorithm the visual system might use, its purpose should be the same, that is, to extract constant descriptors from variable signals. Yet the method is likely to be much more complex than the simple signatures we propose here. For example, to simplify computations here, we have

disregarded several types of information, such as the luminance incidence distribution, the spatial patterns of chromaticity and luminance, and the chromaticity and luminance frequency distributions. We have also excluded from consideration those features which would differentiate objects from one another under a single illumination (e.g. the length of the major axis, or the major/minor axis ratios). The contributions of these additional sources of information to colour constancy is beyond the scope of the current paper, and will be of interest for future studies.

References

- [1] A.C. Hurlbert, "Computational models of colour constancy." In: Walsh V, Kulikowski J (eds). *Perceptual Constancy: Why things look as they do*. Cambridge University Press. pp 283-321 (1998).
- [2] H. E. Smithson, "Sensory, computational and cognitive components of human colour constancy." *Philosophical Transactions of the Royal Society B*. 360, 1329-1346 (2005).
- [3] D.H. Brainard, P. Longere, P.B. Delahunt, W.T. Freeman, J.M. Kraft, B. Xiao, "Bayesian model of human color constancy" *Journal of Vision*. 6, 1267-1281 (2006).
- [4] A. Hurlbert, "Colour vision: Is colour constancy real?" *Current Biology*. 9, 558-561 (1999).
- [5] L.T. Maloney, "Evaluation of linear models of surface spectral reflectance with small number of parameters." *Journal of Optical Society of America A*. 3, 1673-1683 (1986).
- [6] L. Macdonald, W. Ji, "Colour characterisation of a high-resolution digital camera." *Proceeding of the First European Conference on Color in Graphics, Imaging and Vision*. 433-437 (2002).
- [7] R.T. Eskew, J.S. McLellan, F. Giulianini, "Chromatic detection and discrimination." In *Colour Vision: From Genes to Perception*, 345-368 (1999).
- [8] T. Hansen, M. Olkkonen, S. Walter, K.R. Gegenfurtner, "Memory modulates color appearance." *Nature Neuroscience*. 9(11), 1367-1368 (2006).

Author Biography

Yazhu Ling received her BS in Computer Science from the Southwest Jiaotong University (China) in 2000. She then obtained a MSc in Colour Imaging from University of Derby in 2001 and a PhD in human colour perception from Newcastle University (UK) in 2006. Since then she has worked as a research associate in the Institute of Neuroscience at Newcastle University. Her research interests include color preference and color appearance for real natural objects..