

Single Surface Colour Constancy

Graham Finlayson and Gerald Schaefer°
School of Information Systems
University of East Anglia, Norwich, United Kingdom

Abstract

There are two broad classes of colour constancy algorithms: statistical and physics-based. The former attempt to correlate the statistics of the colours in an image with statistical knowledge about light and surfaces. If there is good colour diversity in a scene then the statistical approach often works well. The latter, physics-based algorithms are founded on an understanding of how physical processes such as specularities and interreflection manifest themselves in images. The theory behind physics-based algorithms is both elegant and powerful. Indeed, colour constancy becomes possible even in scenes with as few as two surfaces. Unfortunately the theory rarely translates into practice; most physics-based algorithms do not work outside the lab.

In this paper we combine both statistical and physical knowledge in a new colour constancy algorithm. First, we observe that, statistically, the chromaticities of most illuminants are tightly clustered around the Planckian locus. Second, we make use of the physics-based dichromatic model of image formation. This model predicts that the chromaticities corresponding to a single convex, uniformly coloured, surface fall along a line in chromaticity space. This line is spanned by the body chromaticity (the colour of the surface) and the interface chromaticity (the colour of the shiny part of the surface i.e. the colour of the light). Simply by intersecting the dichromatic chromaticity line with the illuminant locus our algorithm arrives at an estimate of illumination.

Remarkably, (yet by definition) our algorithm can estimate the colour of the light even when there is just a single surface in a scene (the lowest colour diversity possible). Moreover, and more importantly, experiments on real images demonstrate that estimation accuracy can be very good.

1. Introduction

Colour constancy algorithms attempt to decouple the colour of light from the colour of surfaces in images. By definition then colour constancy algorithms cannot work if the same image is formed from different illuminant–surface pairs. A red wall viewed under white light is the same as a white wall viewed under red light. However, this apparently sound theoretical argument turns out to be difficult to demonstrate in practice.

For example, Gilchrist and Ramachandran[1] have shown that observers can distinguish between white rooms illuminated by red light and the converse. This is only possible if the physical stimulus reaching the eye differs for the two viewing conditions. Indeed, it does; it is argued that the effect of interreflected light provides sufficient extra information to discriminate between the two scenes.

While interreflection has been shown to be theoretically useful,[2] it is hard to find and use in real images. Fortunately, many physical effects manifest themselves in images and each helps us distinguish between apparently indistinguishable scenes (e.g. a red room under white light and vice versa). In the context of this paper we show how the physics of specular highlights coupled with the statistics of illumination can be used to estimate the prevailing scene illumination.

Many authors, notably Shafer[3] and Tominaga and Wandell[4,5] have proposed a dichromatic model of reflectance. Under this model the light reflected from a surface is comprised of two parts: body and interface reflection. The body part models conventional matte surfaces: light enters the surface, is modulated by the surface reflectance function and then exits. The interface reflectance models highlights. Here, the incident light does not enter the surface but is rather reflected in a mirror like way. As such the interface reflection has the same spectral characteristics as the illumination. The shape of the surface dictates the relative proportion of body and interface reflection that is scattered in different directions.

In terms of a camera image, each RGB for a surface is also comprised of two parts: RGB-body and RGB-interface. Together RGB-body and -interface span what is called the dichromatic plane and, because light is additive, all RGBs for a surface must fall on this plane. All of this in turn implies that the chromaticities of a surface fall along a line in chromaticity space. Somewhere on this line is the chromaticity for RGB-interface i.e. the chromaticity of the illuminant.

Because the dichromatic chromaticity line contains the chromaticity of the illuminant it follows that the illuminant might be estimated by intersecting two chromaticity lines (for two surfaces). Indeed, this algorithm is well known and has been proposed by several authors (including Tominaga

° {graham, gerald}@sys.uea.ac.uk

and Wandell,[4] Lee[7] and Tong and Funt[8]). Unfortunately, these algorithms do not work well in practice. There are two reasons for this failure. First, the intersection calculation is highly sensitive to image noise. In particular, if two dichromatic lines have similar orientation then small orientation changes lead to large shifts in the intersection point and so to large estimate uncertainty. The second problem is that one needs to find two specular surfaces in images. While this is easy in toy' images, perhaps containing highly coloured saturated spheres, it has proven hard to achieve given images of the natural world.

To get around these problems we will use a single surface and so a single dichromatic line. But, in doing so the problem of illuminant estimation apparently becomes ill-posed. In order to progress further we need a second constraint. In previous work many authors[9,10] have argued that the colours of illumination is limited. A saturated purple light while physically possible is practically speaking impossible. So, allowing it to be a possibility makes no sense. In the context of this paper the range of illuminants is modelled by the Planckian locus. As readers will know this is simply a line in colour space which curves from yellow indoor lights through to whitish outdoor illumination and thence to blue sky. Significantly we check the veracity of the locus constraint by plotting the chromaticities of 172 measured lights including daylights and fluorescents. These are all seen to cluster tightly about the Planckian locus.

If the illuminant chromaticity lies somewhere on the dichromatic line of a surface and it also lies somewhere on the Planckian locus then it follows that it can be found by intersecting the dichromatic line with the Planckian locus. Indeed, this simple intersection is at the heart of our new colour constancy algorithm.

By definition our algorithm can solve for the illuminant (and so colour constancy) given the image of one surface. Moreover, so long as the orientation of the dichromatic line is far from that of the illuminant locus then the intersection should be quite stable. As we shall see, this implies that colours unlikely to be illuminants, e.g. greens and purples, can be used to estimate the illuminant. We point out that green is a dominant colour in the natural world and so our algorithm can potentially be applied to many real world images.

To test our algorithm we take images of green leaves under three lights: D65, TL84, and A (blue, white and yellow lights). In all cases our algorithm returns an excellent illuminant estimate.

The rest of the paper is organised as follows: Section 2 provides a brief review of colour image formation, colour constancy and the dichromatic reflection model. Section 3 describes the new algorithm in detail. Section 4 gives some experimental results.

2. Background

An image taken with a linear device such as a digital colour camera is composed of sensor responses that can be described by:

$$\underline{p} = \int_{\omega} C(\lambda)\underline{R}(\lambda)d\lambda \tag{1}$$

where λ is wavelength, $\underline{p}$ is a 3-vector of sensor responses (RGB pixel values), C is the colour signal (the light reflected from an object), $\underline{R}$ is the 3-vector of sensitivity functions of the device. Integration is performed over the visible spectrum ω.

The colour signal $C(\lambda)$ itself depends both on the surface reflectance $S(\lambda)$ and the spectral power distribution $E(\lambda)$. For pure Lambertian (matte) surfaces $C(\lambda)$ is proportional to the product $S(\lambda)E(\lambda)$ and its magnitude depends on the angle(s) between the surface normal and the light direction(s). The brightness of Lambertian surfaces is independent of the viewing direction.

In real life, however, most objects are non-Lambertian, and so have some glossy or highlight component. The combination of matte reflectance together with a geometry dependent highlight component is modelled by the dichromatic reflectance model.[3,4,5,6]

The dichromatic reflection model for inhomogeneous objects states that the colour signal is composed of two additive components, one being associated with the interface reflectance and the other describing the body (or Lambertian) reflectance part.[3] Both of these components can further be decomposed into a term describing the spectral power distribution of the reflectance and a scale factor depending on the geometry. This can be expressed as:

$$C(\theta,\lambda) = m_I(\theta)C_I(\lambda) + m_B(\theta)C_B(\lambda) \tag{2}$$

where $C_I(\lambda)$ and $C_B(\lambda)$ are the spectral power distributions of the interface and the body reflectance respectively, and m_I and m_B are the corresponding weight factors depending on the geometry θ which includes the incident angle of the light, the viewing angle and the phase angle.

It can be seen that the colour signal can be expressed as the weighted sum of the two reflectance components and the colour signals for an object are thus restricted to a plane. In fact, as colour signals cannot be negative and no more light can be reflected from an object than is cast on them (the dichromatic model is restricted to non-fluorescent objects), all colour signals of an object lie on a parallelogram.[3]

Making the roles of light and surface explicit, equation (2) can be further expanded to:

$$C(\theta,\lambda) = m_I(\theta)S_I(\lambda)E(\lambda) + m_B(\theta)S_B(\lambda)E(\lambda) \tag{3}$$

As for many materials, the index of refraction does not change significantly over the visible spectrum it can be assumed to be constant. $S_I(\lambda)$ is thus a constant and equation (3) becomes:

$$C(\theta,\lambda) = m_{I'}(\theta)E(\lambda) + m_B(\theta)S_B(\lambda)E(\lambda) \tag{4}$$

where $m_{I'}$ describes both the geometry depending weighting factor and the constant reflectance of the interface term.

By substituting equation (4) into equation (1) we get the device's sensor responses for dichromatic reflectances:

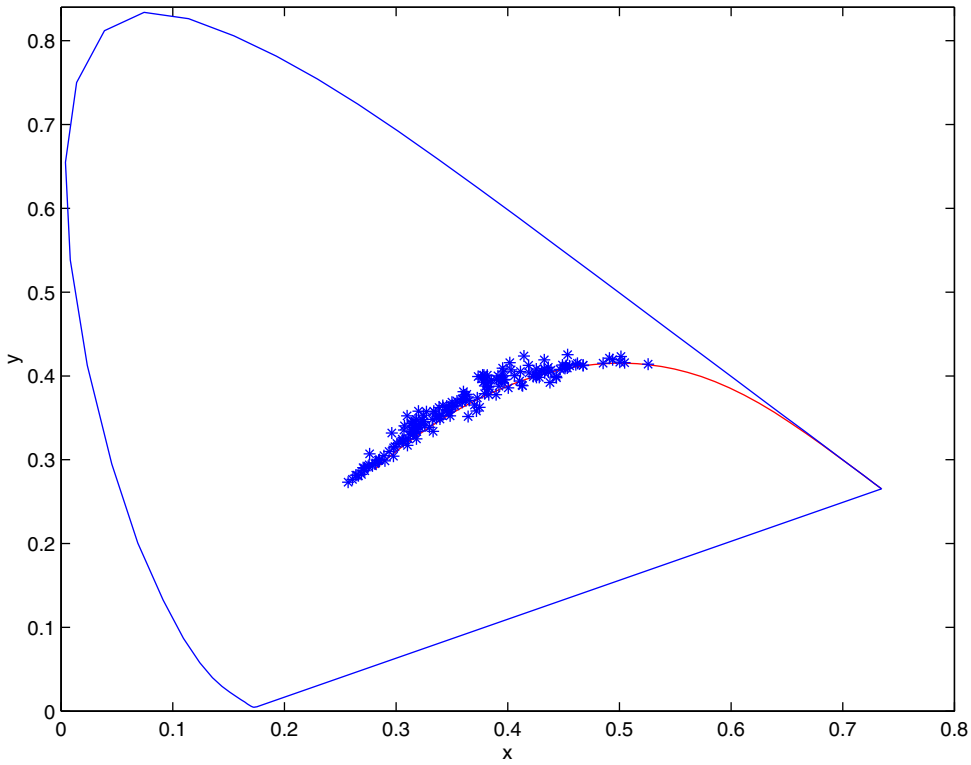

Figure 1. Distribution of 172 illuminants and Planckian locus in xy chromaticity plot

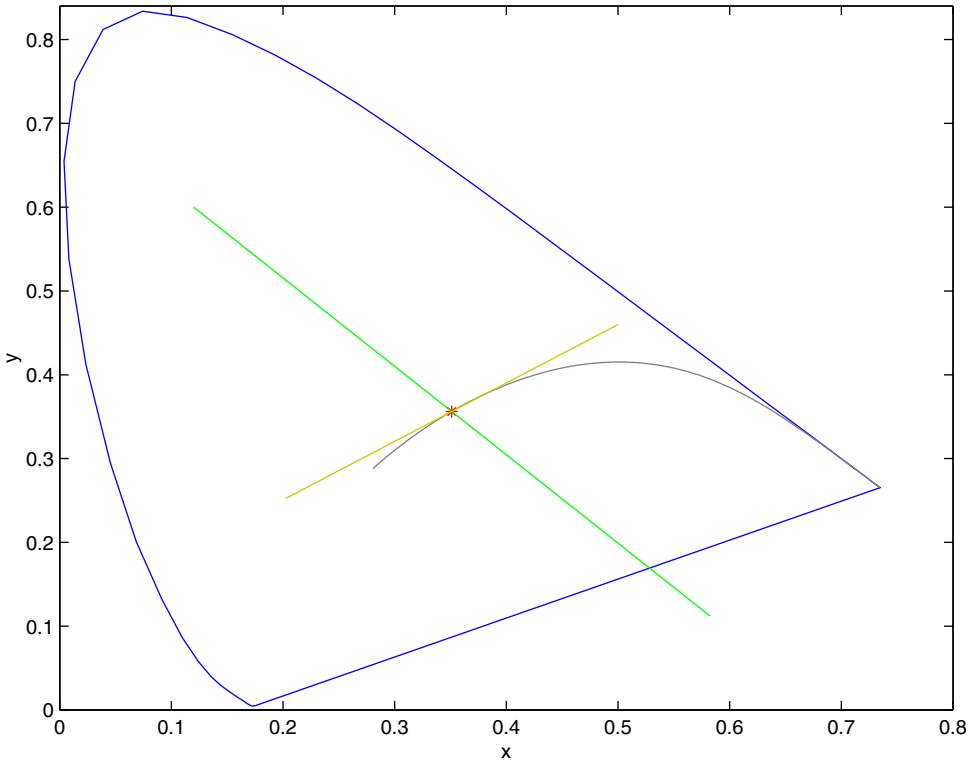

Figure 2. Intersection of dichromatic plane with Planckian locus for a green and a yellow surface

$$\underline{p} = \int_{\omega} m_{I_2}(\theta)E(\lambda)\underline{R}(\lambda)d\lambda + \int_{\omega} m_B(\theta)S_B(\lambda)E(\lambda)\underline{R}(\lambda)d\lambda \qquad (5)$$

which we rewrite as:

$$\begin{pmatrix} R \\ G \\ B \end{pmatrix} = m_{I_2}(\theta)\begin{pmatrix} R \\ G \\ B \end{pmatrix}_I + m_B(\theta)\begin{pmatrix} R \\ G \\ B \end{pmatrix}_B \qquad (6a)$$

where R, G, and B are the red, green, and blue pixel value outputs of the digital camera. Because the RGB of the interface reflectance is equal to the RGB of the illuminant *E* we rewrite (6a) making this observation explicit:

$$\begin{pmatrix} R \\ G \\ B \end{pmatrix} = m_{I_2}(\theta)\begin{pmatrix} R \\ G \\ B \end{pmatrix}_E + m_B(\theta)\begin{pmatrix} R \\ G \\ B \end{pmatrix}_B \qquad (6b)$$

Though the dichromatic reflection model makes several assumptions and simplifications it accurately describes many materials, including plastics, paints, ceramics, leaves, and others.[6]

Dichromatic Colour Constancy

As can be seen from equation (6) the RGBs for a surface lie on a two-dimensional plane one component of which is the RGB of the illuminant. If we consider two objects within the same scene (and assume that the illumination is constant across the scene) then we end up with two RGB planes. Both planes however contain the same illuminant RGB. This means that their intersection must be the illuminant itself. Indeed this is the essence of dichromatic colour constancy.[4]

However, though theoretically sound, dichromatic colour constancy algorithms only perform well under idealised conditions, for real images the estimate of the illuminant RGB turns out not to be that accurate. The reason for this that is in the presence of a small amount of image noise, the intersection of two dichromatic planes can change quite drastically. Hence dichromatic colour constancy tends to work well for highly saturated surfaces taken under laboratory conditions but much less well for real images (say of typical outdoor natural scenes). In fact, the authors know of no dichromatic algorithm which works well for such images.

3. Single-Surface Dichromatic Colour Constancy

The novel colour constancy algorithm proposed in this paper is also based on the fact that many objects exhibit highlights and that the light coming from these objects can be described by the dichromatic reflectance model. However, in contrast to the dichromatic algorithms described above and all other colour constancy algorithms (save that of Yuille[11] which works only under the severest of constraints), it can estimate the illuminant even when there is just a single surface in the scene. Moreover, rather than being an interesting curiosity this single surface constancy behaviour actually represents a significant improvement over previous algorithms. Perhaps, colour constancy is sometimes best calculated using the information from only one surface.

We make use of the fact that the RGBs from an object fall on a plane that contains the illuminant RGB. However, in order to be able to extract this illuminant vector, a further constraint has to be used.

If we look at the distribution of typical light sources then we find that they occupy a highly restricted region of colour space. We took 172 measured light sources, including common daylights and fluorescents, and 100 measurements of illumination reported in [10], and plotted them, in Figure 1, on the xy chromaticity diagram. It is clear that the illuminant chromaticities fall on a long thin "band'' in chromaticity space. As might be expected, this band is very close the Planckian locus of black body radiators, even though it is technically possible to manufacture light sources that do not possess this characteristic. Indeed, it is so close we propose that the Planckian locus can be used as an illuminant constraint.

Single-surface dichromatic colour constancy proceeds in two simple steps. First, the dichromatic plane for a single surface is calculated and this plane is projected to a line in chromaticity space. In the second step the dichromatic line is intersected with the Planckian locus. The intersection point defines the chromaticity of the illuminant.

In proposing this algorithm we are well aware that it will not be a panacea. Rather, it will work well in some circumstances and not others. Indeed, it is well known that all dichromatic colour constancy algorithms tend to work poorly when they rely on measurements made from a white surface. The reason for this is easy to understand. For a white surface the body and interface colours are the same: there is no dichromatic line. Rather, there is a dichromatic point! Moreover, as discussed earlier traditional dichromatic algorithms, which work by intersecting dichromatic planes, work poorly when the planes have similar orientations.

Figure 3. Plant captured under D65, Illuminant A, TL84 (from left to right).

It is interesting then to consider our new algorithm and in particular the conditions under which it will be expected to work. By looking at the xy chromaticity plot in Figure 1 we can qualitatively predict which surface colours will lead to good estimates. For green objects the colour signal plane

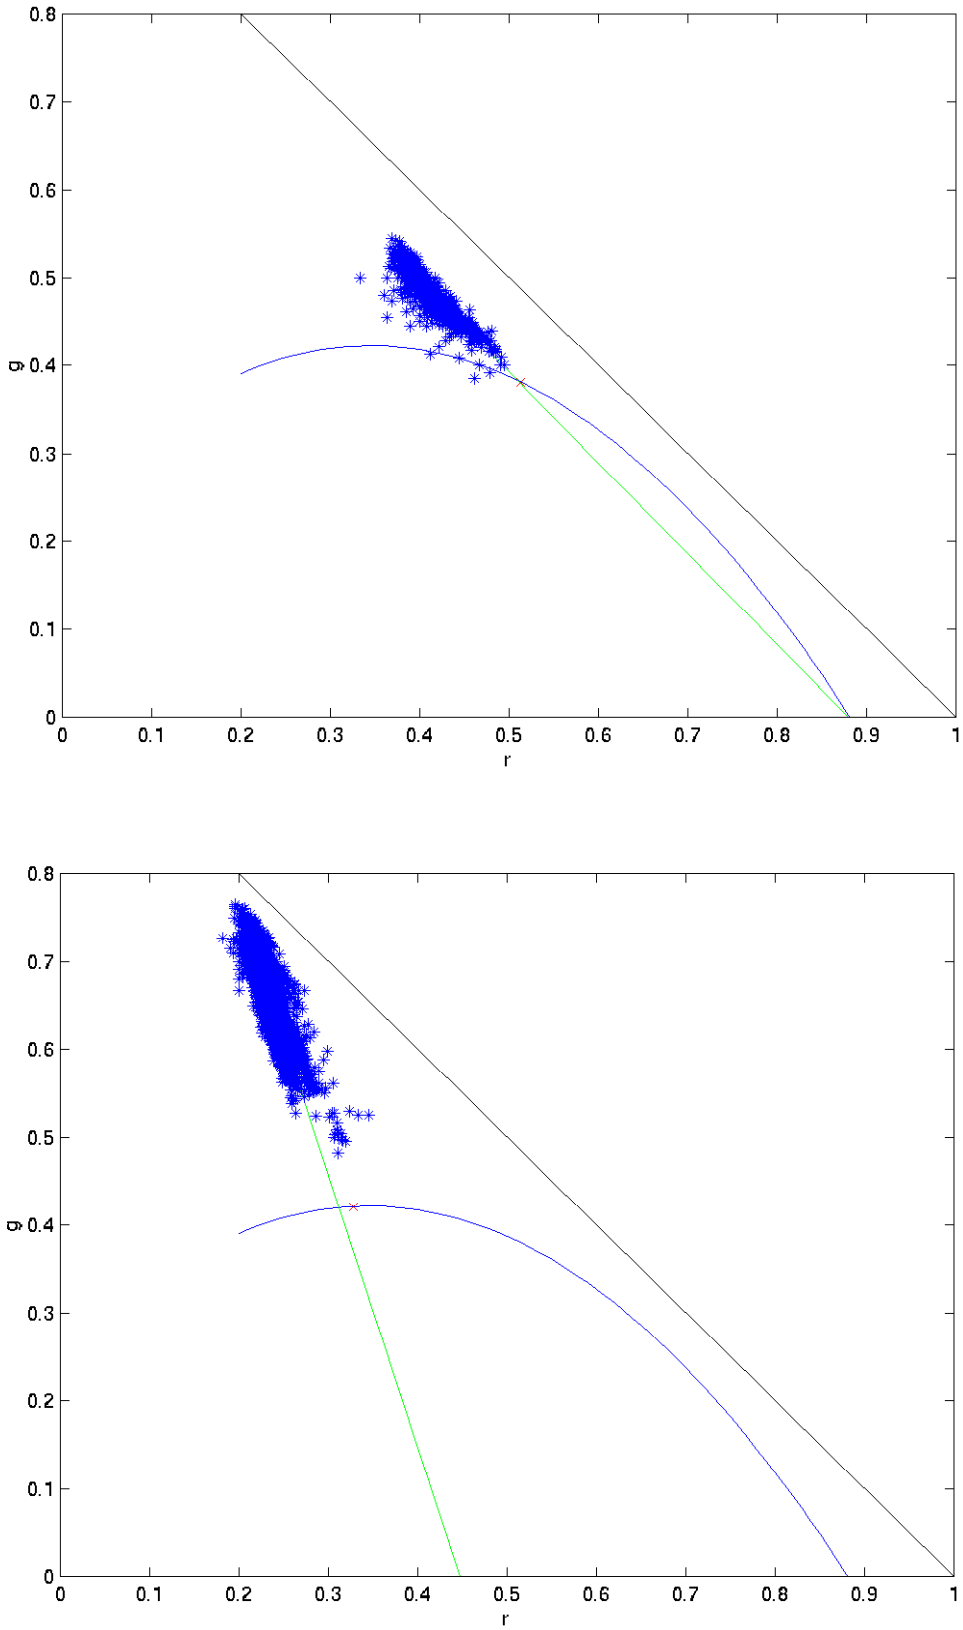

Figure 4. Illuminant estimation plot: top for plant under Illuminant A, bottom under TL84

and the Planckian locus are approximately orthogonal to each other, hence the intersection should give a very reliable estimate of the illuminant. A small amount of noise will not effect the intersection. Good performance is also predicted for purples and magentas. However, when colours are in the yellow and blue illuminant regions the orientation of the plane will be similar to that of the locus itself (see Fig. 2), furthermore it may intersect twice. It is also clear that the results will be more accurate if we have objects with saturated colours as the colour signal plane will be better defined.

This observation, that the solution for green surfaces will be an accurate one, is actually very welcome because if we look around we will notice that "the world is green", ie. many objects in nature are green: grass is green, bushes and forests are predominately green, etc. Consequently, we can expect good results for the new algorithm when applied to many natural images.

4. Experimental Results

In order to prove the validity of our algorithm and to evaluate its performance, a set of experiments were carried out. An Agfa Studiocam was used to capture a green plant in a lightning cabinet with a black background under three different illuminants: a daylight D65 simulator, TL84 (fluorescent light) and Illuminant A (tungsten). The three resulting images are shown as Fig. 3, the different appearance due to different illumination is evident from these images. It is also evident that there are significant highlights.

The best fitting RGB dichromatic planes for the images were computed using singular value decomposition.[4] The computed planes were then projected to lines on the rg chromaticity diagram. Since we do not know the spectral sensitivities of our camera the Planckian locus was approximated by a cubic spline fitted to the three chromaticities of the illuminants (D65, TL84 and A) in rg chromaticity space. Operating in this way also had the advantage that the correct answer was exactly represented on our illuminant locus. A similar cubic spline fitting of the xy chromaticities of our illuminants (measured using a spectro-radiometer) returned a locus which was very similar to the actual Planckian locus.

The intersection of the dichromatic plane with the approximated Planckian locus gives the estimate of the illuminant. The process is shown in Figure 4 (for both the image taken under Illuminant A and TL84) where the distribution of the colour signals (blue asterisks) is plotted together with the resulting dichromatic plane (green line) and the Planckian locus (blue line) in an rg chromaticity diagram; the location of the real illuminant is shown as a red cross. It can be seen that the dichromatic plane intersects with the locus giving an illuminant estimation and that this estimation is very close to the actual illuminant.

To further demonstrate the performance of the algorithm, a Macbeth Checker Chart[12] was captured under the same three illuminants as the plant (the images are shown in Figure 5). The D65 image was chosen to be the correct answer i.e. to represent the correct colours that should be returned by a colour constancy algorithm. We set out to correct the TL84 and Tungsten Macbeth checkers using the illuminant estimates calculated from the plant images. Correcting the plant images themselves turned out not to be interesting since correction accuracy is hard to assess given images of one colour.

(a) Macbeth Checker Chart under D65

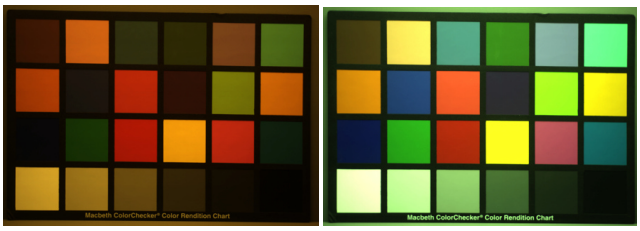

(b) under Illuminant A *(c) under TL84.*

Figure 5.

An estimate of the illuminant chromaticity was transformed into scaling factors for mapping the RGBs to the canonical light. If (r,g,b) denotes the chromaticity estimated illuminant and (r',g',b') the chromaticity of D65 (b=1-r-g and b'=1-r'-g') then the scaling factors are proportional to: (r'/r, g'/g, b'/b). That is the estimated white point is mapped to the white point of D65. Because only the chromaticity of the light is estimated (it is not possible to calculate absolute brightness[13]) the average image pixel value for the transformed images is set equal to the average for the D65 checker. Note that this is a single correction factor based on the global average of all the red, all the green and all the blue pixel values. It controls overall brightness but leaves image colours unchanged.

This correction procedure enables us to explore our algorithms performance on a whole gamut of colours (18 chromatic and 6 neutral colours) and so we can more easily visually assess the algorithm's performance.

The results of correcting the Checker images based on estimates of the illuminant calculated from the plant images are shown Figure 6. Clearly the colours shown in Figure 6 are very close to those shown in Figure 5a (the actual colours under D65). It is evident that the results we get prove to be very good and all the colours are corrected quite accurately. For comparison we have also corrected the images by an illuminant estimation based on running the grey-world algorithm[14] on the plant images. The results of

this are shown in Figure 7. As one would expect the grey world algorithm works very poorly indeed. The average of the plant images regardless of the illuminant is basically green! Also, a correction based on Land's colour constancy algorithm,[15] here referred to as Max. RGB and its results shown in Figure 8, yields significantly inferior results.

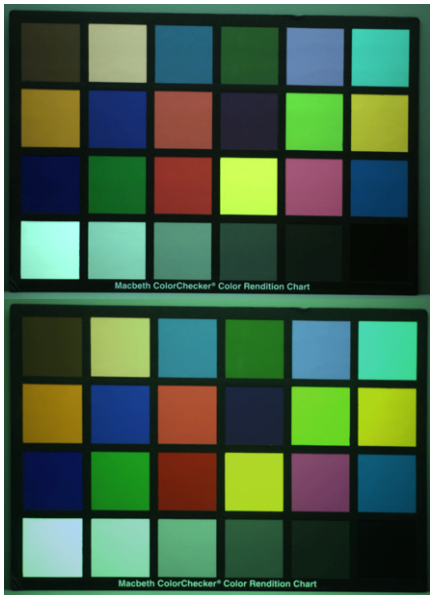

Figure 6. Macbeth Checker Chart under Illuminant A and TL84 corrected for D65 using the novel algorithm.

Figure 7. Macbeth Checker Chart under Illuminant A and TL84 corrected for D65 using grey-world algorithm.

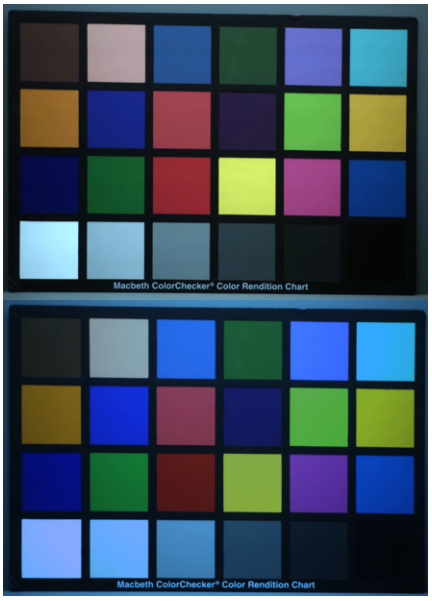

Figure 8. Macbeth Checker Chart under Illuminant A and TL84 corrected for D65 using Max. RGB algorithm.

The visual results are also confirmed by a numerical comparison given in Table 1 where the average error of the Macbeth colour patches and its standard deviation in rg chromaticity units is given for the three algorithms.

Table 1. Mean errors and standard deviations (in brackets) of the 24 Macbeth colours corrected based the mentioned algorithms

	TL84 image	Illum. A image
Single Surface CC	0.0416 (0.0190)	0.0413 (0.0311)
Grey-world	0.1808 (0.0460)	0.0987 (0.0298)
Max. RGB	0.0974 (0.0224)	0.0651 (0.0263)

5. Conclusion

In this paper we have developed a new algorithm for colour constancy which has a number of novel features. First, the algorithm works even when an image only contains a single surface (the most extreme test for all constancy algorithms). Second, the algorithm is refreshingly simple. The set of all possible lights, modelled here as the Planckian locus, is intersected with the set of all image chromaticities. By virtue of the dichromatic reflectance model the latter is a line in chromaticity space which must contain the illuminant chromaticity. By intersecting this line with the Planckian Locus we estimate illumination. That is all there is to our algorithm. Third, and perhaps most novel overall, our algorithm is shown to work well on real images.

Acknowledgement

The authors wish to thank Hewlett-Packard Incorporated for supporting this work. This work was conducted while the

authors were with The Colour and Imaging Institute, University of Derby, Derby, UK.

References

1. Gilchrist A. L., Ramachandran V., *Investigative Ophthalmology and Visual Science*, **33**, 756 (1992).
2. Funt B. V., Drew M. S., Ho J., *Int. Journal Comp. Vision*, **6**, 5-24 (1991).
3. Shafer S. A., *Color Research and Application*, **10**, 210-218 (1985).
4. Tominaga S., Wandell B. A., *J. Opt. Soc. Am. A*, **6**, 576-584 (1989).
5. Tominaga S., Wandell B. A., *J. Opt. Soc. Am. A*, **7**, 312-317 (1990).
6. Tominaga S., *IEEE Trans. Pattern Analysis and Machine Intelligence*, **13**, 658-670 (1991).
7. Lee H.-C., *J. Opt. Soc. Am. A*, **3**, 1694-1699 (1986).
8. Tong F., Funt B.V., *Proc.: Vision Interface 1988*, 98-103 (1988).
9. Finlayson G., *IEEE Transactions on PAMI*, **18**, 1034-1038 (1996).
10. Barnard K., Finlayson G., Funt B., *Computer Vision and Image Understanding*, **65**, 311-321 (1997).
11. Yuille A., *Biological Cybernetics*, **56**, 195-201 (1987).
12. McCamy C.S., Marcus H., Davidson J.G., *J. App. Photog. Eng.*, 95-99 (1976).
13. Maloney L. T., Wandell B. A., *J. Opt. Soc. Am. A*, **3**, 29-33 (1986).
14. Buchsbaum G., *Journal of the Franklin Institute*, **310**, 1-26 (1980).
15. Land E. H., *Scientific American*, 108-129 (1977).