Constrained Dichromatic Colour Constancy

Graham D. Finlayson and Gerald Schaefer

School of Information Systems
University of East Anglia
Norwich NR4 7TJ
United Kingdom
Telephone: +44-1603-592446
Fax: +44-1603-593344
{graham,gerald}@sys.uea.ac.uk

Abstract. Statistics-based colour constancy algorithms work well as long as there are many colours in a scene, they fail however when the encountering scenes comprise few surfaces. In contrast, physics-based algorithms, based on an understanding of physical processes such as highlights and interreflections, are theoretically able to solve for colour constancy even when there are as few as two surfaces in a scene. Unfortunately, physics-based theories rarely work outside the lab. In this paper we show that a combination of physical and statistical knowledge leads to a surprisingly simple and powerful colour constancy algorithm, one that also works well for images of natural scenes.
From a physical standpoint we observe that given the dichromatic model of image formation the colour signals coming from a single uniformly-coloured surface are mapped to a line in chromaticity space. One component of the line is defined by the colour of the illuminant (i.e. specular highlights) and the other is due to its matte, or Lambertian, reflectance. We then make the statistical observation that the chromaticities of common light sources all follow closely the Planckian locus of black-body radiators. It follows that by intersecting the dichromatic line with the Planckian locus we can estimate the chromaticity of the illumination. *We can solve for colour constancy even when there is a single surface in the scene.* When there are many surfaces in a scene the individual estimates from each surface are averaged together to improve accuracy.
In a set of experiments on real images we show our approach delivers very good colour constancy. Moreover, performance is significantly better than previous dichromatic algorithms.

1 Introduction

The sensor responses of a device such as a digital camera depend both on the surfaces in a scene and on the prevailing illumination conditions. Hence, a single surface viewed under two different illuminations will yield two different sets of sensor responses. For humans however, the perceived colour of an object is more or less independent of the illuminant; a white paper appears white both outdoors under bluish daylight and indoors under yellow tungsten light, though

D. Vernon (Ed.): ECCV 2000, LNCS 1842, pp. 342–358, 2000.

the responses of the eyes' colour receptors, the long-, medium-, and short-wave sensitive cones, will be quite different for the two cases. This ability is called colour constancy. Researchers in computer vision have long sought algorithms to make colour cameras equally colour constant.

Perhaps the most studied physics-based colour constancy algorithms, i.e. algorithms which are based on an understanding of how physical processes manifest themselves in images, and the ones which show the most (though still limited) functionality, are based on the dichromatic reflectance model for inhomogeneous dielectrics (proposed by Shafer [11], Tominaga and Wandell [14, 15], and others). Inhomogeneous materials are composed of more than one material with different refractive indices, usually there exist a vehicle dielectric material and embedded pigment particles. Examples of inhomogeneous dielectrics include paints, plastics, and paper. Under the dichromatic model, the light reflected from a surface comprises two physically different types of reflection, interface or surface reflection and body or sub-surface reflection. The body part models conventional matte surfaces, light enters the surface, is scattered and absorbed by the internal pigments, some of the scattered light is then re-emitted randomly, thus giving the body reflection Lambertian character. Interface reflection which models highlights, usually has the same spectral power distribution as the illuminant. Because light is additive the colour signals from inhomogeneous dielectrics will then fall on what is called a dichromatic plane spanned by the reflectance vectors of the body and the interface part respectively.

As the specular reflectance represents essentially the illuminant reflectance, this illuminant vector is contained in the dichromatic plane of an object. The same would obviously be true for a second object. Thus, a simple method for achieving colour constancy is to find the intersection of the two dichromatic planes. Indeed, this algorithm is well known and has been proposed by several authors [2,8,14,16]. When there are more than two dichromatic planes, the best common intersection can be found [14,16]. In a variation on the same theme Lee [8] projects the dichromatic planes into chromaticity space and then intersects the resulting dichromatic lines. In the case of more than two surfaces a voting technique based on the Hough transform is used.

Unfortunately dichromatic colour constancy algorithms have not been shown to work reliably on natural images. The reasons for this are twofold. First, an image must be segmented into regions corresponding to specular objects before such algorithms can be employed. We are not too critical about this problem since segmentation is, in general, a very hard open research problem. However, the second and more serious problem (and the one we address in this paper) is that the dichromatic computation is not robust. When the interface and body reflectance RGBs are close together (the case for most surfaces) the dichromatic plane can only be approximately estimated. Moreover, this uncertainty is magnified when two planes are intersected. This problem is particularly serious when two surfaces have similar colours. In this case the dichromatic planes have similar orientations and the recovery error for illuminant estimation is very high.

In this paper we add statistical knowledge in a new dichromatic colour constancy algorithm. Here we model the range of possible illuminants by the Planckian locus of black-body radiators. The Planckian locus is a line in colour space which curves from yellow indoor lights to whitish outdoor illumination and thence to blue sky light. Significantly experiments show that this locus accounts for most natural and man made light sources. By intersecting a dichromatic line with the Planckian locus we recover the chromaticity of the illuminant. By definition our algorithm can solve for the illuminant given the image of one surface, the lowest colour diversity possible. Moreover, so long as the orientation of the dichromatic line is far from that of the illuminant locus the intersection should be quite stable. As we shall see, this implies that colours unlikely to be illuminants like greens and purples, will provide an accurate estimate of the illumination, whereas the intersection for colours whose dichromatic lines have similar orientations to the Planckian locus is more sensitive to noise. However, even hard cases lead to relatively small errors in recovery.

Experiments establish the following results: As predicted by our error model, estimation accuracy does depend on surface colour: green and purple surfaces work best. However, even for challenging surfaces (yellow and blues) the method still gives reasonable results. Experiments on real images of a green plant (easy case) and a caucasian face (hard case) demonstrate that we can recover very accurate estimates for real scenes. Indeed, recovery is good enough to support pleasing image reproduction (i.e. removing a colour cast due to illumination). Experiments also demonstrate that an average illuminant estimate calculated by averaging the estimates made for individual surfaces leads in general to very good recovery. On average, scenes with as few as 6 surfaces lead to excellent recovery. In contrast, traditional dichromatic algorithms based on finding the best common intersection of many planes perform much more poorly. Even when more than 20 surfaces are present, recovery performance is still not very good (not good enough to support image reproduction).

The rest of the paper is organised as follows. Section 2 provides a brief review of colour image formation, the dichromatic reflection model, and the statistical distribution of likely illuminants. Section 3 describes the new algorithm in detail. Section 4 gives experimental while Section 5 concludes the paper.

2 Background

2.1 Image Formation

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), and $\underline{R}$ 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 on both the surface reflectance $S(\lambda)$ and the spectral power distribution $E(\lambda)$ of the illumination. 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.

2.2 Dichromatic Reflection Model

In the real world, 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 modeled by the dichromatic reflectance model [11,14,15,6,13].

The dichromatic reflection model for inhomogeneous dielectric 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 [11]. 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.

Equation (2) shows that the colour signal can be expressed as the weighted sum of the two reflectance components. Thus the colour signals for an object are restricted to a plane.

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

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

Since 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 $m_{I'}$ now describes both the geometry dependent weighting factor and the constant reflectance of the interface term.

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

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

which we rewrite as

$$\begin{pmatrix} R \\ G \\ B \end{pmatrix} = m_{I'}(\theta) \begin{pmatrix} R \\ G \\ B \end{pmatrix}_I + m_B(\theta) \begin{pmatrix} R \\ G \\ B \end{pmatrix}_B \tag{6}$$

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 (6) making this observation explicit:

$$\begin{pmatrix} R \\ G \\ B \end{pmatrix} = m_{I'}(\theta) \begin{pmatrix} R \\ G \\ B \end{pmatrix}_E + m_B(\theta) \begin{pmatrix} R \\ G \\ B \end{pmatrix}_B \tag{7}$$

Usually chromaticities are written as the two dimensional coordinates (r, g) since $(b = 1 - r - g)$. Clearly, given (7) we can write:

$$\begin{pmatrix} r \\ g \end{pmatrix} = m_{I'}(\theta) \begin{pmatrix} r \\ g \end{pmatrix}_E + m_B(\theta) \begin{pmatrix} r \\ g \end{pmatrix}_B \tag{8}$$

That is, each surface spans a dichromatic line in chromaticity space.

2.3 Dichromatic Colour Constancy

Equation (7) shows that 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 implies that their intersection must be the illuminant itself. Indeed, this is the essence of dichromatic colour constancy [2,8,14,16].

Notice, however, that the plane intersection is unique up to an unknown scaling. We can recover the chromaticity of the illuminant but not its magnitude. This result can be obtained directly by intersecting the two dichromatic chromaticity lines, defined in Equation (8), associated with each surface. Indeed, many dichromatic algorithms and most colour constancy algorithms generally solve for illuminant colour in chromaticity space (no colour constancy algorithm to date can reliably recover light brightness; indeed, most make no attempt to do so[1].)

Though theoretically sound, dichromatic colour constancy algorithms only perform well under idealised conditions. For real images the estimate of the illuminant turns out not to be that accurate. The reason for this is that in the presence of a small amount of image noise the intersection of two dichromatic lines can change quite drastically, depending on the orientations of the dichromatic lines. This is illustrated in Figure 1 where a 'good' intersection for surfaces

[1] Because $E(\lambda)S(\lambda) = \frac{E(\lambda)}{k} kS(\lambda)$ it is in fact impossible to distinguish between a bright light illuminating a dim surface and the converse. So, the magnitude of $E(\lambda)$ is not usually recoverable.

having orientations that are nearly orthogonal to each other is shown as well as an inaccurate illuminant estimation due to a large shift in the intersection point caused by noise for lines with similar orientations. 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 natural scenes.

2.4 Distribution of Common Illuminants

In practice the number and range of light sources is limited [3,1]. Although it is physically possible to manufacture lets say a purple light, it practically does not exist in nature. As a consequence if we want to solve for colour constancy we would do well not to consider these essentially impossible solutions. Rather we can restrict our search to a range of likely illuminants based on our statistical knowledge about them. Indeed, this way of constraining the possible estimates produces the best colour constancy algorithms to date [3,4]. We also point out that for implausible illuminants, like purple lights, human observers do not have good colour constancy.

If we look at the distribution of typical light sources more closely, then we find that they occupy a highly restricted region of colour space. To illustrate this, we took 172 measured light sources, including common daylights and fluorescents, and 100 measurements of illumination reported in [1], and plotted them, in Figure 2, on the xy chromaticity diagram[2]. It is clear that the illuminant chromaticities fall on a long thin 'band' in chromaticity space.

Also displayed in Figure 2 is the Planckian locus of black-body radiators defined by Planck's formula:

$$M_e = \frac{c_1 \lambda^{-5}}{e^{\frac{c_2}{\lambda T}} - 1} \tag{9}$$

where M_e is the spectral concentration of radiant exitance, in watts per square meter per wavelength interval, as a function of wavelength λ, and temperature T in kelvins. c_1 and c_2 are constants and equal to 3.74183×10^{16} Wm^2 and 1.4388×10^{-2} mK respectively.

Planck's black-body formula accurately models the light emitted from metals, such as Tungsten, heated to high temperature. Importantly, the formula also predicts the general shape (though not the detail) of daylight illuminations.

3 Constrained Dichromatic Colour Constancy

The colour constancy algorithm proposed in this paper is again based on the fact that many objects exhibit highlights and that the colour signals coming from

[2] The xy chromaticity diagram is like the rg diagram but x and y are red and green responses of the standard human observer used in colour measurement.

those objects can be described by the dichromatic reflection model. However, in contrast to the dichromatic algorithms described above and all other colour constancy algorithms (save that of Yuille [18] 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.

Constrained dichromatic colour constancy proceeds in two simple steps. First, the dichromatic line for a single surface is calculated. This can be done e.g. by singular value decomposition [14] or by robust line fitting techniques. In the second step the dichromatic line is intersected with the Planckian locus: the intersection then defines the illuminant estimate.

As the Planckian locus (which, can accurately be approximated by either a 2^{nd} degree polynomial or by the daylight locus [17]) is not a straight line, there arise three possibilities in terms of its intersection with a dichromatic line, all of which can be solved for analytically. First, and in the usual case, there will exist one intersection point which then defines the illuminant estimate. However, if the orientation of the dichromatic line is similar to that of the illuminant locus, then we might end up with two intersections (similar to Yuille's algorithm [18]) which represents the second possibility. In this case, several cases need to be considered in order to arrive at a unique answer. If we have prior knowledge about the scene or its domain we might be able to easily discard one of the solutions, especially if they are far apart from each other. For example, for face images, a statistical model of skin colour distributions could be used to identify the correct answer. If we have no means of extracting the right intersection, one could consider the mean of both intersections which will still give a good estimate when the two intersections are relatively close to each other. Alternatively we could look at the distribution of colour signals of the surface, if those cross the Planckian locus from one side then we can choose the intersection on the opposite site, as a dichromatic line will only comprise colour signal between the body and the interface reflectance. Finally, the third possibility is that the dichromatic line does not intersect at all with the Planckian locus. This means that the orientation of the dichromatic line was changed due to image noise. However, we can still solve for the point which is closest to the illuminant locus, and it is clear that there does exist a unique point in such a case.

As we see, in each case we are able to find a unique illuminant estimate and so constrained dichromatic colour constancy based on only a single surface is indeed possible.

In proposing this algorithm we are well aware that its performance will be surface dependent. By looking at the chromaticity plot in Figure 3 we can qualitatively predict which surface colours will lead to good estimates. For green colours the dichromatic line and the Planckian locus are approximately orthogonal to each other, hence the intersection should give a very reliable estimate of the illuminant where a small amount of noise will not affect the intersection. The same will be true for surface colours such as magentas and purples. However,

when colours are in the yellow or blue regions the orientation of the line is similar to that of the locus itself, and hence the intersection is very sensitive to noise. In order to provide some quantitative measurement of the expected accuracy of our algorithm we have developed error models which we introduce in the next section.

3.1 Error Surfaces for Constrained Dichromatic Colour Constancy

We define a dichromatic line in a linear RGB colour space based on a camera with sRGB [7] sensitivities by its two end points, one representing the pure body reflection and the other being the specular part based on a certain illuminant. Gaussian noise of variance 0.005 and mean 0 (chromaticities must fall in the interval [0,1]) is added to both points before they are projected into chromaticity space where their line is intersected with the Planckian locus. For realism we restrict to colour temperatures between 2000 and 10000 kelvin. Within this range we find all typical man-made and natural illuminants. As the temperature increases from 2000K to 10000K so illuminants progress from orange to yellow to white to blue. The intersection point calculated is converted to the corresponding black-body temperature. The distance between the estimated and the actual illuminant, the difference in temperature, then defines the error for that specific body reflectance colour. In order to get a statistical quantity, this experiment was repeated 100 times for each body colour and then averaged to yield a single predicted error. The whole procedure has to be carried out for each point in chromaticity space to produce a complete error surface.

Figure 4 shows such an error surface generated based on a 6500K bluish light source. (Further models for other lights are given elsewhere [5].) As expected, our previous observations are verified here. Greens, purples, and magentas give good illuminant estimates, while for colours close to the Planckian locus such as yellows and blues the error is highest (look at Figure 1 to see where particular surface colours are mapped to in the chromaticity diagram). The error surface also demonstrates that in general we will get better results from colours with higher saturation, which is also what we expect for dichromatic algorithms. Notice that performance for white surfaces is also very good. This is to be expected since if one is on the locus at the correct estimate then all lines must go through the correct answer. Good estimation for white is important since, statistically, achromatic colours are more likely than saturated colours. In contrast the conventional dichromatic algorithm fails for the achromatic case: the dichromatic plane collapses to a line and so no intersection can be found.

3.2 Integration of Multiple Surfaces

Even though our algorithm is designed by definition for single surfaces, averaging estimates from multiple surfaces is expected to improve algorithm accuracy. Notice, however, that we cannot average in chromaticity space as the Planckian locus is non linear (and so a pointwise average might fall off the locus). Rather each intersection is a point coded by its temperature and the average temperature

is computed. Given the temperature and Planck's formula it is easy to generate the corresponding chromaticity.

4 Experimental Results

In our first experiment we measured, using a spectraradiometer the light reflected from the 24 reflectances on a Macbeth colour chart (a standard reference chart with 24 surfaces [9]) viewed under 6500K (bluish light) and at two orientations. The resulting spectra were projected, using Equation (1), onto sRGB camera curves to generate 48 RGBs. Calculating RGBs in this way (rather than using a camera directly) ensures that Equation (1) holds exactly: we can consider algorithm's performance independent of confounding features (e.g. that a camera is only approximately linear). Each pair of RGBs (one per surface) defines a dichromatic line. Intersecting each dichromatic line returns an illuminant estimate. Taking the 24 estimates of the illuminant we can plot the recovery error as a function of chromaticity as done in Figure 5 and compare these errors with those predicted by our error model (Figure 4). The resemblance of the two error distributions is evident.

From the distribution of errors, a green surface should lead to excellent recovery and a pink surface, close to the Planckian locus, should be more difficult to deal with. We tested our algorithm using real camera images of a green plant viewed under a 6500K bluish daylight simulator, under fluorescent TL84 (≈5000K, whitish), and under tungsten light (≈2800K, yellowish). The best fitting dichromatic RGB planes were projected to lines on the rg chromaticity space where they were intersected with the Planckian locus. The plant images together with the resulting intersections are shown in Figures 6, 7, and 8. It can be seen that in each case the intersection point is indeed very close to the actual illuminant. A pink face image provides a harder challenge for our algorithm since skin colour lies close to the Planckian locus (it is desaturated pink). However, illuminant estimation from faces is a particularly valuable task because skin colour changes dramatically with changing illumination and so complicates face recognition [12] and tracking [10]. In Figure 9 a face viewed under 4200K is shown together with a plot showing the dichromatic intersection that yields the estimated illuminant. Again, our algorithm manages to provide an estimate close to the actual illuminant, the error in terms of correlated colour temperature difference is remarkably only 200 K.

As we have outlined in Section 3.2, combining cues from multiple surfaces will improve the accuracy of the illuminant estimate. To verify that and to determine a minimum number of surfaces which will lead to sufficient colour constancy was the task of our final experiment. We took images of the Macbeth colour checker at two orientations under 3 lights (Tungsten (2800K, yellow) , TL84 (5000K, whitish) and D65 (6500K, bluish)). We then randomly selected between 1 and 24 patches from the checker chart and ran our algorithm for each of the 3 lights. In order to get meaningful statistical values we repeat this procedure many times. The result of this experiment is shown in Figure 10 which demonstrates that the

curve representing the average error (expressed in K) drops significantly already for a combination of only a few surfaces. In Figure 10 we have also drawn a line at about 500 kelvin. Experiments have shown that a recovery error of 500 K suffices as an acceptibility threshold for the case of image reproduction (for more details see [5]). So, visually acceptable estimation is achieved with as few as 6 surfaces.

Finally, we wanted to compare this performance to that of traditional colour constancy algorithms. We implemented Lee's method for intersecting multiple dichromatic planes in chromaticity space [8]. Figure 11 summarizes estimation error as a function of the number of surfaces. Notice that the conventional dichromatic algorithm fails to reach the acceptability threshold. Performance is comparatively much worse than for our new algorithm.

5 Conclusion

In this paper we have developed a new algorithm for colour constancy which is (rather remarkably) able to provide an illuminant estimate for a scene containing a single surface. Our algorithm combines a physics-based model of image formation, the dichromatic reflection model, with a constraint on the illuminants that models all possible light sources as lying on the Planckian locus. The dichromatic model predicts that the colour signals from a single inhomogeneous dielectric all fall on a line in chromaticity space. Statistically we observe that almost all natural and man made illuminants fall close to the Planckian locus (a curved line from blue to yellow in chromaticity space). The intersection of the dichromatic line with the Planckian locus gives our illuminant estimate. An error model (validated by experiment) shows that some surfaces lead to better recovery than others. Surfaces far from and orthogonal to the Planckian locus (e.g. Greens and Purples) work best. However, even 'hard' surfaces (e.g. pink faces) lead to reasonable estimation. If many surfaces are present in a scene then the average estimate can be used.

Experiments on real images of a green plant (easy case) and a caucasian face (hard case) demonstrate that we can recover very accurate estimates for real scenes. Recovery is good enough to support pleasing image reproduction (i.e. removing a colour cast due to illumination). Experiments also demonstrate that an average illuminant estimate calculated by averaging the estimates made for individual surfaces leads in general to very good recovery. On average, scenes with as few as 6 surfaces lead to excellent results. In contrast, traditional dichromatic algorithms based on finding the best common intersection of many planes perform much more poorly. Even when more than 20 surfaces are present, recovery performance is still not very good (not good enough to support image reproduction).

Acknowledgements

The authors wish to thank Hewlett-Packard Incorporated for supporting this work.

References

1. K. Barnard, G.D. Finlayson, and B.V. Funt. Color constancy for scenes with varying illumination. *Computer Vision and Image Understanding*, 65(2):311–321, February 1997.
2. M. D'Zmura and P. Lennie. Mechanisms of color constancy. *Journal of the Optical Society of America A*, 3(10):1662–1672, 1986.
3. G.D. Finlayson. Color in perspective. *IEEE Trans. PAMI*, 18(10):1034–1038, October 1996.
4. G.D. Finlayson, P.M. Hubel, and S. Hordley. Color by correlation. In *The 5^{th} Color Imaging Conference*, pages 6–11, Scottsdale, Arizona, 1997.
5. G.D. Finlayson and G. Schaefer. Constrained dichromatic colour constancy. Technical report, School of Information Systems, University of East Anglia, Norwich, UK, 1999.
6. G. Healey. Using color for geometry-insensitive segmentation. *Journal of the Optical Society of America A*, 6(6):920–937, 1989.
7. IEC. Default RGB colour space - sRGB. *Standards document*, 1998.
8. H.-C. Lee. Method for computing the scene-illuminant from specular highlights. *Journal of the Optical Society of America A*, 3(10):1694–1699, 1986.
9. C.S. McCamy, H. Marcus, and J.G. Davidson. A color-rendition chart. *J. App. Photog. Eng.*, pages 95–99, 1976.
10. Y. Raja, S. McKenna, and S. Gong. Colour model selection and adaptation in dynamic scenes. In *Proceedings European Conference Computer Vision*, Freiburg, Germany, 1998.
11. S.A. Shafer. Using color to separate reflection components. *Color Res. App.*, 10(4):210–218, 1985.
12. M. Störring, H.J. Andersen, and E. Granum. Skin colour detection under changing lighting conditions. In *Proceedings 7th Symposium on Intelligent Robotics Systems*, pages 187–195, Coimbra, Portugal, June 1999.
13. S. Tominaga. Surface identification using the dichromatic reflection model. *IEEE Trans. PAMI*, 13(7):658–670, July 1991.
14. S. Tominaga and B.A. Wandell. Standard surface-reflectance model and illuminant estimation. *Journal of the Optical Society of America A*, 6(4):576–584, April 1989.
15. S. Tominaga and B.A. Wandell. Component estimation of surface spectral reflectance. *Journal of the Optical Society of America A*, 7(2):312–317, 1990.
16. F. Tong and B.V. Funt. Specularity removal for shape from shading. In *Proceedings: Vision Interface 1988*, pages 98–103, Edmonton, Alberta, Canada, 1988.
17. G. Wyszecki and W.S. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulas*. Wiley, New York, 2nd edition, 1982.
18. A. Yuille. A method for computing spectral reflectance. *Biological Cybernetics*, 56:195–201, 1987.

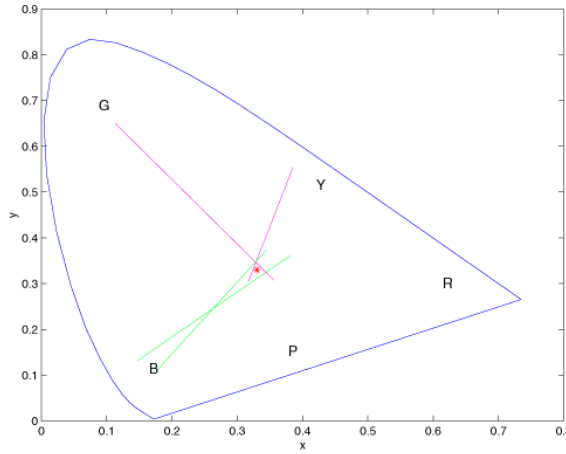

Fig. 1. Inaccuracy due to image noise in the illuminant estimation by dichromatic colour constancy. Only when the orientations of the dichromatic planes (here projected to dichromatic lines in xy chromaticity space) is far from each other (purple lines) the solution won't be affected too much by noise. If however their orientations are similar, only a small amount of noise can lead to a big shift in the intersection point (green lines). The red asterisk represents the chromaticity of the actual illuminant. Also shown is the location of certain colours (Green, Yellow, Red, Purple, and Blue) in the chromaticity diagram.

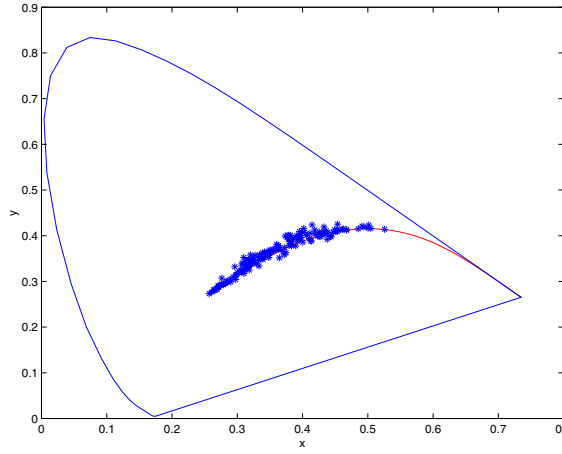

Fig. 2. Distribution of 172 measured illuminants and the Planckian locus (red line) plotted in xy chromaticity space. It can be seen that the illuminants are clustered tightly around the Planckian locus.

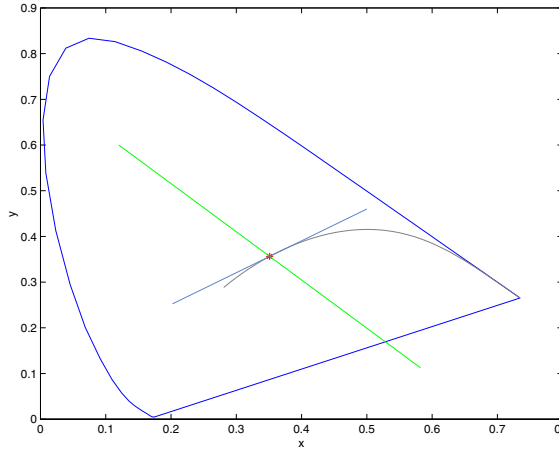

Fig. 3. Intersection of dichromatic lines with the Planckian locus for a green (or purple) object (green line) and for a blue (or yellow) object (blue line). While the orientations for green/purple objects are approximately orthogonal to that of the Planckian locus, for blue/yellow objects the dichromatic lines have similar orientations to the locus.

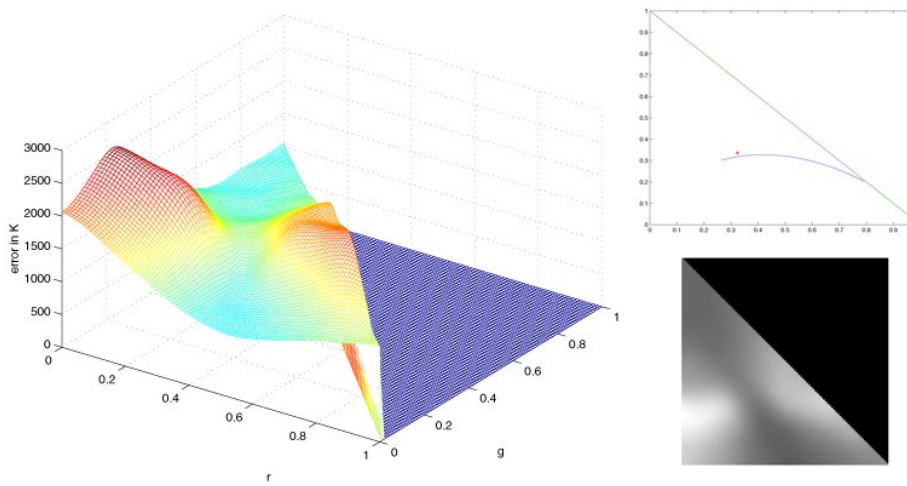

Fig. 4. Error surface generated for D65 illuminant and displayed as 3D mesh (left) and 2D intensity (bottom right) image where white corresponds to the highest and black to the lowest error. The diagram at the top right shows the location of the Planckian locus (blue line) and the illuminant (red asterisk).

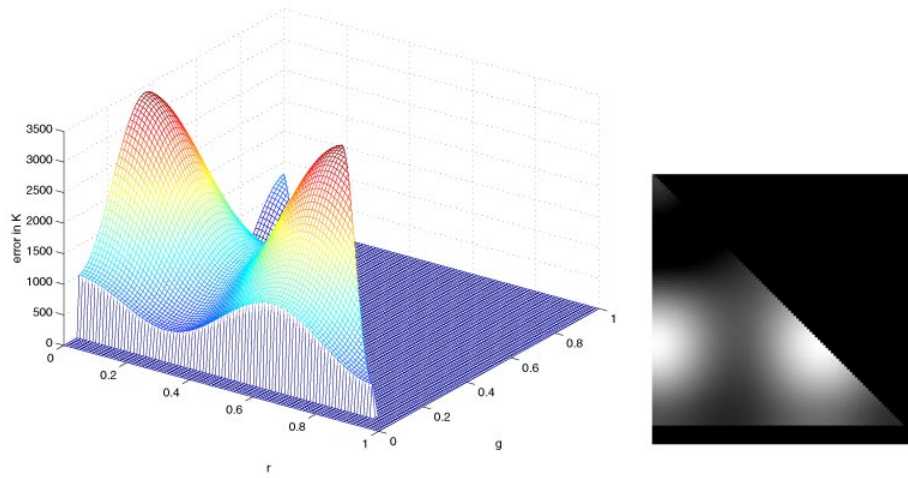

Fig. 5. Error surface generated obtained from spectral measurements of the Macbeth Checker Chart (left as 3D mesh, right as intensity image). Values outside the range of the Macbeth chart were set to 0 (black).

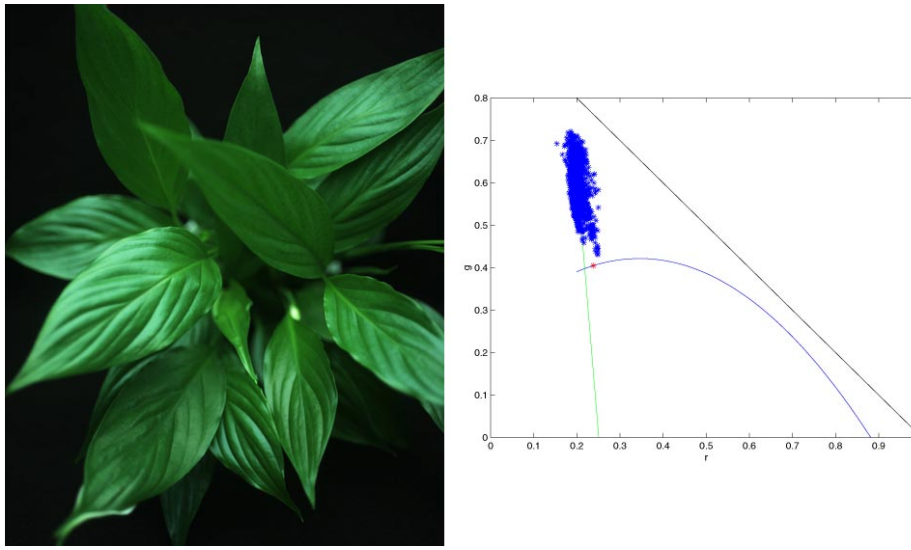

Fig. 6. Image of a green plant captured under D65 (left) and result of the intersection giving the illuminant estimate (right). The real illuminant is plotted as the red asterisk. The blue asterisks show the distribution of the colour signals.

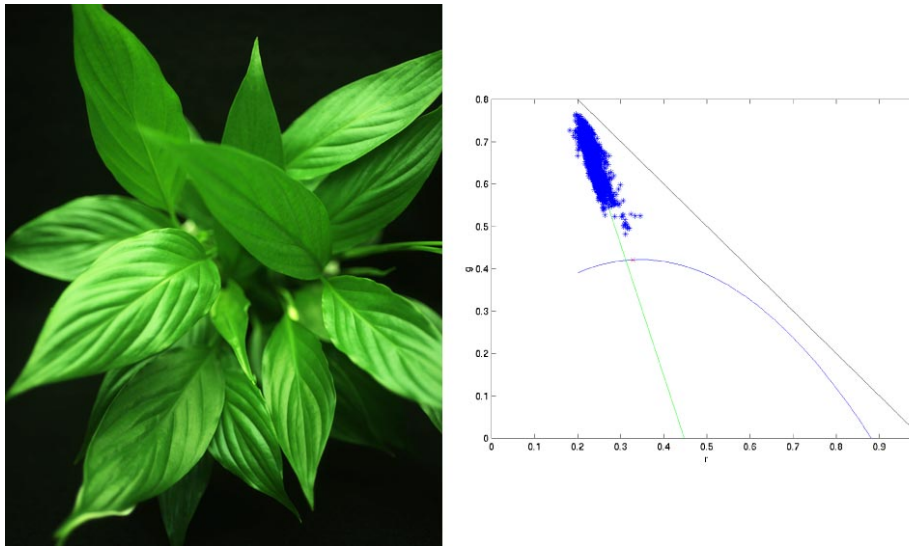

Fig. 7. Image of a green plant captured under TL84 (left) and result of the intersection giving the illuminant estimate (right). The real illuminant is plotted as the red asterisk. The blue asterisks show the distribution of the colour signals.

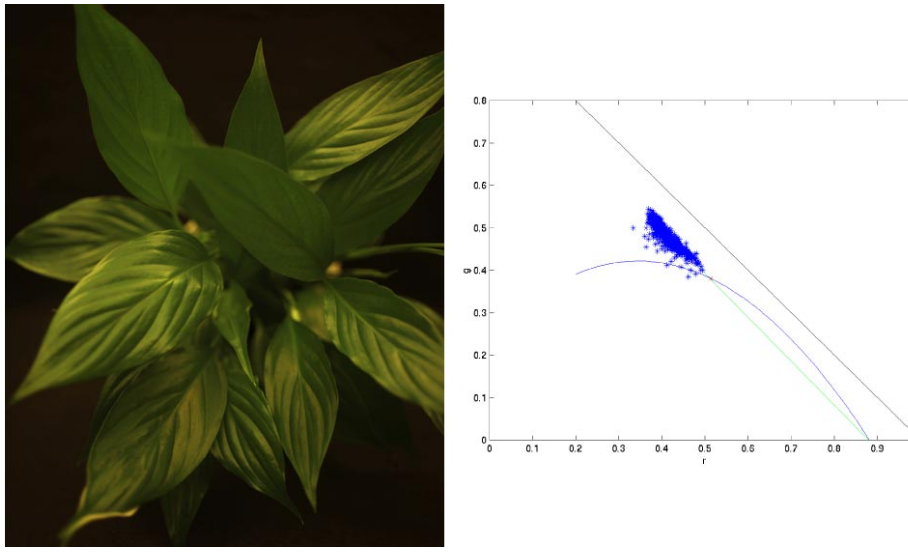

Fig. 8. Image of a green plant captured under Illuminant A (left) and result of the intersection giving the illuminant estimate (right). The real illuminant is plotted as the red asterisk. The blue asterisks show the distribution of the colour signals.

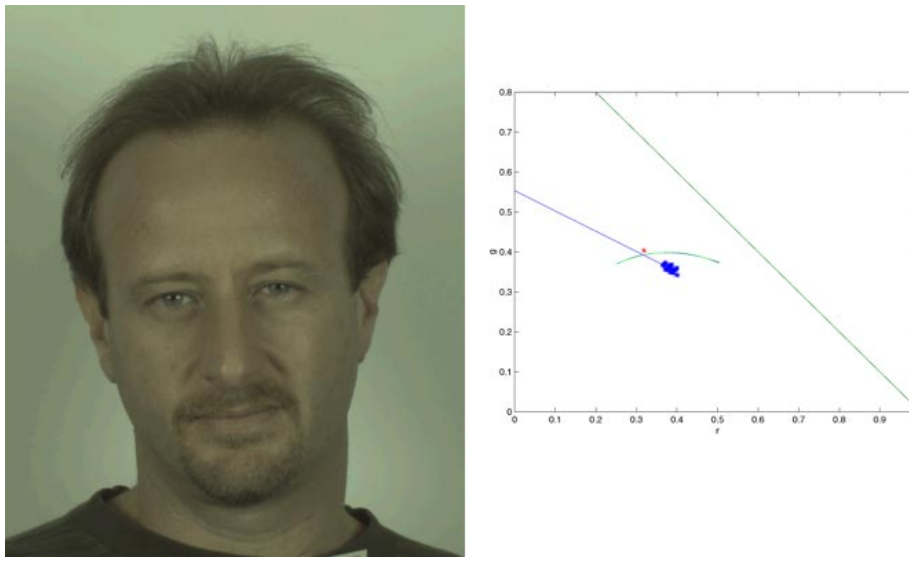

Fig. 9. Image of a face captured under light with correlated colour temperature of 4200 K (left) and result of the intersection giving the illuminant estimate (right). The real illuminant is plotted as the red asterisk. The blue asterisks show the distribution of the colour signals.

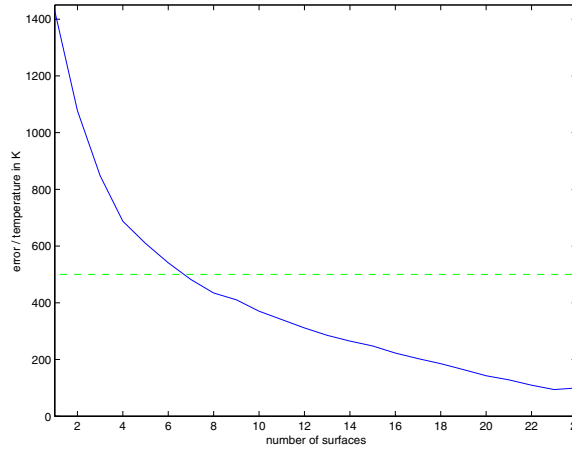

Fig. 10. Performance of finding a solution by combining multiple surfaces. The average estimation error expressed in kelvins is plotted as a function of the number of surfaces used.The horizontal line represents the acceptibility threshold at 500 K. It can be seen that this threshold is reached with as few as about six surfaces.

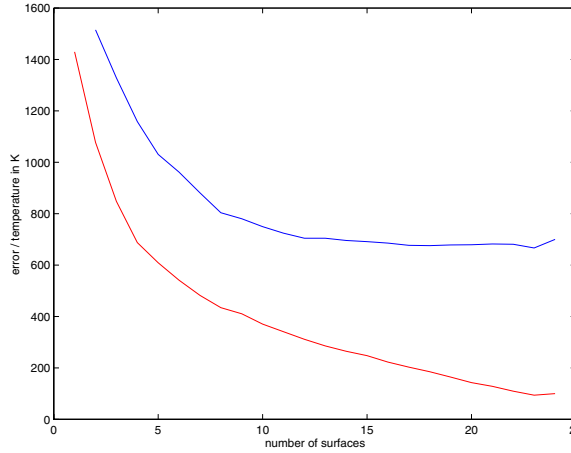

Fig. 11. Performance of traditional dichromatic colour constancy compared to our new approach. The average estimation error expressed in kelvins as a function of the number of surfaces used. The blue line corresponds to traditional dichromatic colour constancy, the red line to our new constrained dichromatic colour constancy.