Revisiting Surface Colour Estimation under Varying Illumination

Graham D. Finlayson and Stuart E. Lynch, University of East Anglia, Norwich, NR4 7TJ

Abstract

There are now several works reported in the literature which attempt to estimate surface colour when the same surface is viewed under two or more lights. There are many practical situations where such information is available including at shadow edges or in surveillance applications where the same scene is viewed over time. Crucially, because typical lights are highly constrained, they fall on or close to the Planckian locus, varying illumination algorithms for surface estimation can, plausibly, estimate surface colour even for scenes with little colour diversity.

One of the first varying illumination algorithms made the empirical observation that the mappings, 2x2 diagonal matrices, taking spectral band ratio chromaticities (e.g. r/b and g/b) for surfaces viewed under a range of typical lights to corresponding values under a reference canonical light (e.g. D65) tended to lie on a 2D line. It follows that applying this 'linear set' of maps to the chromaticity of an arbitrary surface under unknown light results in a line along which the D65 counterpart should lie. Viewing the same surface under a second light results in a second constraint line. The intersection of the two lines results in an estimate of the surface chromaticity under D65.

While this method can work well, Kawakami et al showed that serious estimation errors can result in the presence of even small amounts of image noise. While the noise tended to make only small changes in the slope and intercept of the constraint lines the intersection point could move a significant distance; indeed, the shifted intersection might correspond to a highly improbable (physically impossible) light. To solve this 'intersetion stability' problem they proposed limiting the set of maps not only to lie on a line but on a line segment (e.g. only allow illuminants that are physically plausible and likely). This observation, which necessitated dealing with the problem non-intersecting line segments, formed the foundation of a new algorithm which was shown to deliver a step change in surface colour estimation. In this paper we extend Kawakami's work in two ways. First, we deal with the non-intersecting line problem using a 'total least-squares' approach (as oppose to assuming one or other of the line segments is in error in Kawakami's work). Second, we optimise over the position and length of the line segment map-set used. Experiments demonstrate that our new method delivers significantly improved surface colour estimation. Compared with the Kawakami method we deliver over 50% improved surface colour estimates. We also show that the Kawakami method can be improved by optimising the line segment map set but even in this case our new method still provides about a 25% decrease in estimation error.

Introduction

Colour constancy is the ability to see surfaces independent of the illumination that is cast on them. Humans have reasonably good colour constancy (arguably, better than the current best algorithms). For example, a white t-shirt looks white under most illuminations. Most theories of colour constancy, e.g. [3, 5, 9], attempt to relate an easily derivable summary statistic (mean, max etc) to the colour of the light. These methods tend to work well when there is large surface colour diversity in a scene.

An alternate strategy, exploited by Kawakami et al and Finlayson et al [2, 6, 7, 8] is to assume that we have access to the colour of a surface under two or more unknown illuminants. The key observation exploited in these works is that illumination colour is limited. To a first approximation typical lights tend to lie on or near the Planckian Locus which is a curved line in chromaticity space. This line effectively provides a constraint for solving the two light constancy problem.

Consider the case where we are imaging a greenish surface viewed under two lights. Further, let us suppose the goal of colour constancy is to map the image colour to its corresponding appearance under a known canonical light. Because we know illuminants lie near the Planckian locus (and or are bluish, whitish and yellowish), we can solve for the set of maps (diagonal matrices) that might map a surface colour to its canonical appearance. In [2] this set could be parametised by a straight line in a 2d mapping space. This mapping set applied to the camera response resulted in a linear set (a *surface constraint* line) of possible chromaticities, any one of which could have been the true surface colour under the canonical light. In the case where we had a second light, we ended up with a second set of possible canonical chromaticites. Intersecting the two chromaticities resulted in a single estimate of the surface colour.

The method worked best given two lights which were very far apart in colour. Indeed, when lights were similar in colour, the two constraint lines tended to be very similar. And, in this case small perturbations in the input camera RGBs (e.g. due to noise) would result in small changes in the lines but, significantly, very large changes in the intersection of the lines. This intersection instability was examined in detail by Kawakami et al. [6]. There they proposed that the instability was in part due to modelling illumination maps by a line of infinite length. This clearly was an unnecessarily weak assumption. Planckian illuminants are bounded by temperature (say 2500K to 20000K) and so the Planckian locus is a curved segement in chromticity space. This implies the map set taking image colours to canonical counterparts is also represented as a line segment. Kawakami et al's 'intersecting line segment' algorithm delivered a step change in the accuracy of surface colour estimation.

However, the Kawakami approach raised two interesting technical questions that, we believe, were only partially addressed. First, representing the map sets as a line segment can (and often does) lead to the case where the surface constraint lines do not intersect. In the Kamakami method this problem is solved by, effectively, assuming one of the line segments is in error, and then moving the line to enforce an intersection. We revisited this problem and propose a method which finds the best intersection assuming error in both constraint lines. Effectively, our method is a ‘total least-squares’ solution to the problem. Second, the question of ‘what line segment’ is not addressed in Kawakami. In this paper we find that there is a small benefit in using a slightly shorter segment than that derived from the endpoints of physically realisable lines. Not only are these lights ‘unlikely’, the fact we exclude them means that our estimates of surface colour under highly chromatic (blue or reddish) lights will retain a small element of the light colour.

Background

Let us assume the RGB response in a camera can be written as:

$$p_k = \int_{\omega} E(\lambda)S(\lambda)R_k(\lambda)d\lambda \qquad (1)$$

Where p denotes e response, $E(\lambda)$ is the illumination function, $S(\lambda)$ is the surface reflectance and $R_k(\lambda)$ is the spectral sensitivity of the devices kth sensor. This is integrated across the visible spectrum ω. If we assume our device sensor is narrow-band then we can rewrite equation (1) as the following:

$$p_k^{j,i} = S_k^j E_k^i \qquad (2)$$

Where $S_k^j = S^j(\lambda)R_k(\lambda)$ and $E_k^i = E^i(\lambda)R_k(\lambda)$. If we take a single surface S_k^j and view it under an arbitary illuminant E_k^i and a canonical illuminant E_k^c (where $i \neq c$) then (2):

$$p_k^{j,c} = \frac{E_k^c}{E_k^i} p_k^{j,i} \qquad (3)$$

This equation can be usefully rewritten as:

$$p_{j,c} = \mathscr{D}^{i,c} p^{j,i} \qquad (4)$$

where $\mathscr{D}^{i,c}$ denotes the 3x3 diagonal matrix mapping responses under light i to the canonical light c.

Given a camera response $[r,g,b]$ let us define the chromaticity function c as:

$$([c_1 c_2]) = \left[\frac{r}{b}, \frac{g}{b}\right] \qquad (5)$$

This is a slightly non standard chromaticity space (usually we divide by the sum of r, g and b). This space is adopted because if a diagonal matrix relates image colours across illuminants then this remains true for chromaticities:

$$c_{j,c} = D^{i,c} c^{j,i} \qquad (6)$$

Of course, a 2-d diagonal matrix is parametrised by two numbers (the diagonal elements). Thus we can plot maps on a graph. In previous work, Finlayson et al found, empirically, that the set of maps taking chromaticities under unknown light to canonical counterparts tended to lie on a line. That is, the curved line of Planckian lights was mapped to a straight line using this idea of ’maps to the canonical light’. The set of mappings from unknown illumination to a known canonical light is represented using the following equation:

$$\frac{1}{E_2} = m\frac{1}{E_1} + k \qquad (7)$$

where E_1 and E_2 denote the chromaticity coordinates of the ith light E^i.

Let us take a chromaticity c of an unknown surface under two unknown light $e = 1$ and $e = 2$. According to our linear model of illuminant maps (7) we can, after some algebraic manipulation write down two line equations which define the location of the chromaticities under the canonical light:

$$c_2^c = \left(m\frac{c_2^e}{c_1^e}\right)(1/E_2) + c_2^e k \qquad e = (1,2) \qquad (8)$$

The intersection of these two *surface constraint* lines is the chromaticty under the canonical light.

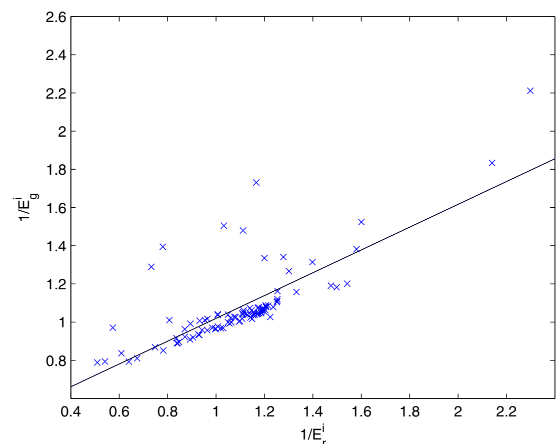

Figure 1. *Typical daylights plotted in inverse chromaticity space, with best fit line.*

For images where there is no noise problem or for illumant colours that are far apart, this constraint line intersection approach delivered good constancy performance. However, Kawakami et al [6, 7] showed that serious estimation errors can result in the presence of even modest amounts of image noise. Part of the problem is that the mapping set of (7) is a line of infinite extent. Physically this does not make sense. The Planckian locus is bounded in chromaticity space and so, the corresponding map set should be bounded. Kawakami et al argued (convincingly) that rather than modelling the map set as a line, we should model it as a line segment. The Kawakami algorithm has 2 modes of operation. In mode 1 the two constraint line segments (bounding the estimate colour of the surface under the canonical light) intersect. In this case the Kawakami algorithm returns the same result as the Finlayson et al algorithm. In the second mode, the line segments do not intersect. Here, based on an argument related to expected sensor noise, one or other of the line segments is assumed to be in error. Then this segment is moved toward the other segment to force an intersection. This process is illustrated in Figure 2.

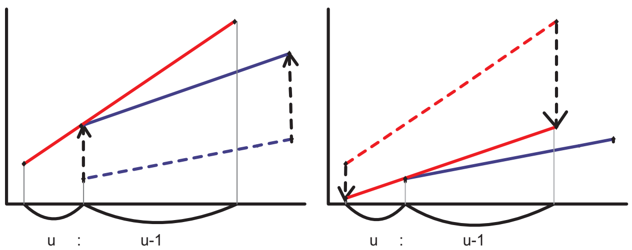

Figure 2. *Left: Case of adjusting first line segment Right: Case of adjusting second line segment. The ratio u does not change in both cases.*

Of course the reader might wonder what happens if shifting the line segment does not, in itself, solve the intersection problem (e.g. the ranges of the x-coordinates of the segments do not overlap). While this problem is acknowledged in the Kawakami work, no solution is presented.

Improved constraint line intersection

Here we wish to revisit how we find the point which best approximates the intersection (of non-intersecting line segments). In Figure 3 we show two non - intersecting line segments. We ask which point (anywhere on the Cartesian plane) is simultaneously closest to both segments. Let us denote the line segments as l_1 and l_2. Abstractly, we might write our objective function as:

$$\min_{p} \operatorname{dist}(p, l_1) + \operatorname{dist}(p, l_2)$$

where $dist(p,l)$ is the length between point p and line l. Clearly, the distance to the line segments is defined by a perpendicular projection. Such projections for the endpoints of each line to the other line segment are shown in Figure 3. Let us now consider where the actual intersection point might be. Could it be, for example, between the endpoints for both lines? This is impossible. The easy case is shown in Figure 4 where the ranges of the x coordinates of the line segments do not overlap. Here the best intersection is the middle of the two closest endpoints. In the second case we assume the range of the x-coordinates do not overlap. Here as a function of the x coordinate and barring parallel lines the two segments will clearly converge in one direction (get closer) and diverge in the other. They converge in the left direction of figure 3. Clearly, as we move an intersection point in the converging direction the length to both lines (their respective perpendicular projections) will decrease. Of course we can move the point only so far until the closest point on either line segment is the endpoint of one of the lines. Thus, we find the intersection point by calculating the closest point to each endpoint in the other line segment. We find the projection point with the minimum distance. The intersection point is simply the average of the two points. All computation is in closed-form.

One advantage of our intersection problem is that unlike the Kawakami method, we always have a well defined intersection point.

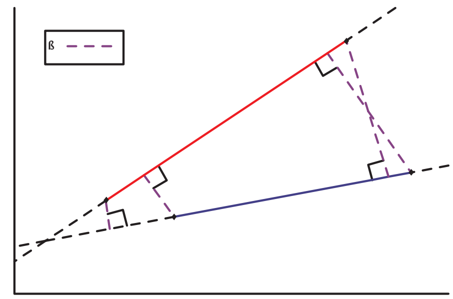

Figure 3. *Demonstrates the β line segments defined on the two device response line segments (shown in red and blue).*

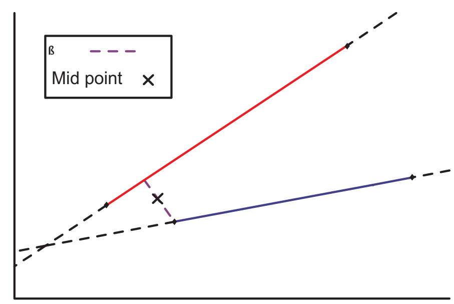

Figure 4. *Shows the accepted β line segment with mid point.*

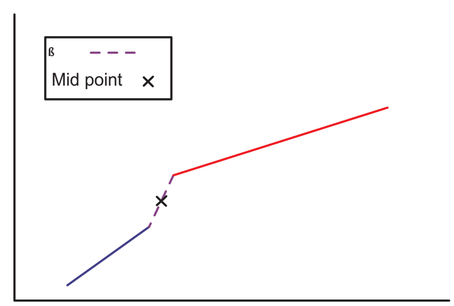

Figure 5. *Demonstrates the case where no acceptable perpendicular intersections are found. Therefore the shortest line between endpoints is used.*

Of course, the intersection point we calculate will depend on the line segment used to model the illuminant maps. Indeed, if we made the segment longer and longer we would eventually observe the same behaviour as the original Finlayson et al algorithm: as the segment grows in length the non-intersection problem would become less frequent. In the opposite direction if we assumed the illuminant map to be a single map e.g. the identity transform then the resulting constancy algorithm resorts to averaging the two input chromaticities (a sort of grey-world approach).

Thus, our second innovation is to consider algorithm performance as a function of line segment length. We predict using a slightly shorter line compared with the physical limits implied by Planckian illumination will result in improved performance.

The most saturated lights in our model appear less frequently than those closer to white. By shortening the line segment representing the set of light maps we will perform slightly less well for strongly chromatic light but better on average.

We also predict that adding the length of the line segment into the mix will also provide a means to improve the Kawakami algorithm.

Results

In this section we compare the results of our algorithm against that proposed by Kawakami et al. We conducted two sets of experiments, one on synthetically generated data, and the other on a set of real image data. The error measure used is the angular error between the RGB of the estimated surface colour and the actual RGB (for a reference light) as it is an error measure independent of intensity. Since we do no expect to recover intensity this is appropriate. Figure 6 shows a visual representation of what effect varying levels of angular error can have on a colour. This figure shows that a recovery angular error of 30 degrees would be the same as suggesting a surface was orange when it was actually green. An error of up to 5 degrees should give an answer in the acceptable colour region.

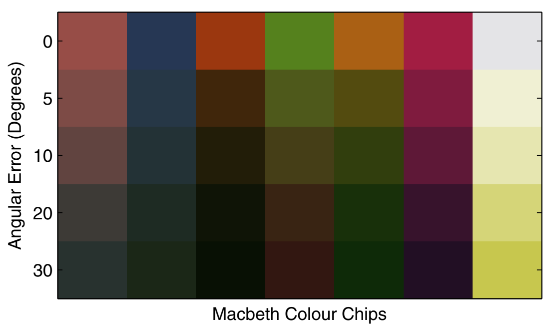

Figure 6. *Examples of the effect varying levels of angular error can have.*

Synthetic Data Experiments

For our experiments we generated a set of 99 test images using the 24 chips of the Macbeth Colour checker [10] and 99 recorded daylights [11] and SONY 930-DXC camera sensitivities [1]. A uniform white light is used as our canonical illuminant. We also introduced normally distributed random noise with a tolerance of 1% SNR into the images. For a single surface colour the input to all algorithms is a pair of RGBs: the same surface viewed under a random pair of lights (with noise). Because the method recovers surface colour chromaticity we make a 3D RGB vector by mapping $[c_1\ c_2]$ to $[c_1\ c_2\ 1]$. To measure how well we have recovered the surface colour (the correct answer defined as the RGB vector under the canonical illuminant) we calculated the angle between the recovered estimate and the actual RGB.

We computed all the results of each surface under any pair the lights. The angular error was calculated between the estimated surface and the actual surface. The median value was then calculated for each surface over all the lights. As a comparison we look at the median recovery error where one of the Macbeth colour checker responses is randomly selected. Our 'guessing' estimate is calculated by taking each one of the 24 Macbeth chips, and working out the median angular error of assuming each one of the other surfaces as the correct answer. Initially we compared our result (predicated on the best line segment and our new intersection method) against the default Kawakami method(figure 7).

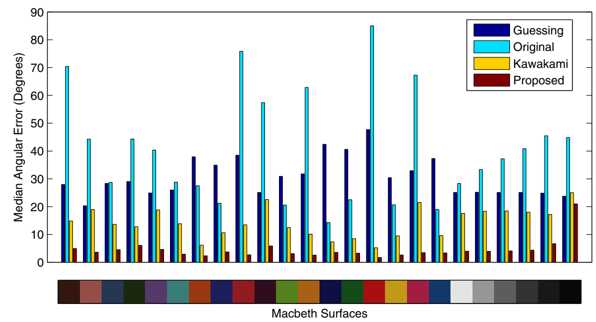

Figure 7. *Results using Kawakami default method and our proposed method, contrasted with the original Finlayson method and 'guessing'. The data contains 1% noise.*

The reader is no doubt surprised at how poor the original Finlayson et al method performs. This is in large part due to the experiment reported here including the cases which Finlayson et al indicated would result in their algorithm working poory: the cases where the illuminant change is small. When two illuminants are close together very small changes in the orientation of the line e.g. induced by the presences of 1% noise can result in large changes in the intersection point. It is known that the Finlayson et al method presented previously only works for illuminants which are far apart (e.g. sun and shadow) and our results shown in 7 and 8 tell the same story.

In the presence of large amount of noise, it turns out that that simply guessing the correct surface gives a better result (though this ranking flips as noise is reduced). Our method shows a significant improvement on guessing in every case, and significantly better than the Kawakami improved algorithm. Also the median error for most surfaces is less than 5 degrees: our recovery returns chromaticities that are visually similar.

As a second test, we also optimised the Kawakami method over all choices of illuminant line segment. Here, performance improves and is similar to that delivered by our algorithm. Though, our median angular error remains about 25% lower than for Kawakami.

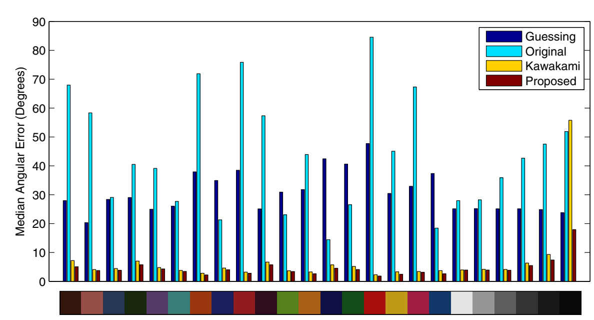

Figure 8. *Results using optimal line segments with 1% noise.*

Real Image Data

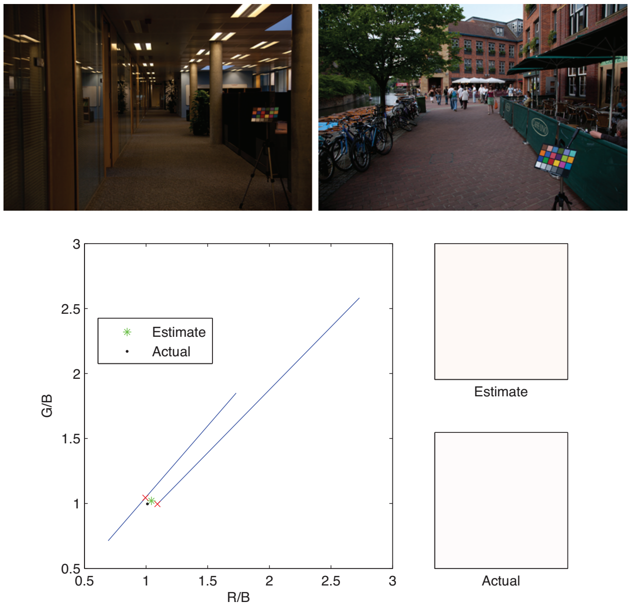

Figure 9. Images from the Microsoft dataset with different illumination. Left: Indoor Right: Outdoor

To assess our method on real image data we used the Microsoft Cambridge dataset used in [4]. The 428 images used were captured using a Canon 1D in RAW format. The Cambridge data set contains 568 images stored as RAW or rendered jpegs. For our purposes we wish to have linear camera responses and so considered only the linear images. Each image contains a Macbeth colour checker and so, after extracting this, we can rerun our previous experiment. However, only 428 of the RAW images supplied have a checker that is sufficiently large to be extractable from the image. So, our test will use these images. Taking all pairs of checkers we have '428 choose 2' = 'illuminant pairs to test' In the synthetic test, D65 was chosen as the reference illuminant. Here we choose an arbitrary image from the test set. Now we extract the 24 RGBs from the checker. The corresponding chromaticities for these responses are, for the purposes of this experiment, the correct answer.

In a preprocessing step we calculate the 2D diagonal map taking the remaining 427 colour checker's to the reference counterpart in a least squares sense. We then optimise for the best line segment that represents this set of maps.

We now randomly select a pair of images (for two unknown lights) and extract the two colour checkers (shown in Figure 9). For each surface in the checker we have two chromaticities. One for each light. We map each by the line segment of illuminant maps. This results in two sets of possible locations where the surface chromaticity should lie for the canonical light. The best intersection of the two segments is our estimate of the surface colour.

The angular error was calculated between the the RGB value of the chips in the image and the estimated values from the algorithm. The median was taken for each chip of the Macbeth Colour Checker. The results are shown in Figure 10. A distinct improvement in our method over Kawakami's is displayed in real image data.

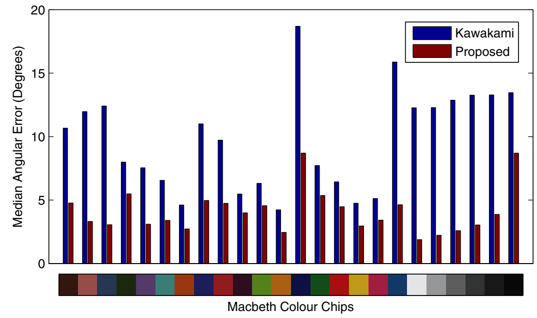

Figure 10. Comparison of results between Kawakami and Proposed method on real images.

The median of the medians is 3.64 (i.e. less than 5). This is remarkable. Given real data (and synthetic data) assuming we can identify the same surface under two lights (e.g. at a shadow edge or over time in a surveillance situation), we can with tolerable accuracy estimate the surface colour.

Conclusions

We have proposed a method which extends previous work to stabilize surface estimation under varying illumination. We developed a method that is robust and able to produce a well defined solution. Our experimental evaluation delivers favourable results when compared against antecedent algorithms.

Acknowledgements

The authors gratefully acknowledge the support of EPSRC grant EP/H022236/1.

References

[1] K. Barnard, L. Martin, B. Funt, and A. Coath. A data set for color research. *Color Research & Application*, 27(3):147–151, 2002.

[2] GD Finlayson, BV Funt, and K. Barnard. Color constancy under varying illumination. In *Proceedings of the Fifth International Conference on Computer Vision*, pages 720–725. IEEE Computer Society, 1995.

[3] G.D. Finlayson and E. Trezzi. Shades of gray and colour constancy. In *Proc. of the Twelfth Color Imaging Conference*, pages 37–41, 2004.

[4] P.V. Gehler, C. Rother, A. Blake, T. Minka, and T. Sharp. Bayesian color constancy revisited. In *IEEE Conference on Computer Vision and Pattern Recognition, 2008. CVPR 2008*, pages 1–8, 2008.

[5] R. Gershon, A.D. Jepson, and J.K. Tsotsos. From [R, G, B] to surface reflectance: computing color constant descriptors in images. *Perception*, 17:755–758, 1988.

[6] R. Kawakami and K. Ikeuchi. Stabilizing Illumination Chromaticity Estimation using the Illumination Line Segment. In *Procs. of the Tenth IAPR Conf. on Machine Vision Applications*, pages 437–440. Citeseer, 2007.

[7] R. Kawakami and K. Ikeuchi. Color estimation from a single surface color. 2009.

[8] R. Kawakami, R.T. Tan, and K. Ikeuchi. Consistent surface color for texturing large objects in outdoor scenes. *Digitally Archiving Cultural Objects*, pages 279–295.

[9] L.T. Maloney and B.A. Wandell. Color constancy: a method for recovering surface spectral reflectance. *Readings in computer vision: issues, problems, principles, and paradigms*, page 293, 1987.

[10] C.S. McCamy, H. Marcus, and JG Davidson. A color-rendition chart. *J. App. Photog. Eng*, 2(3):95–99, 1976.

[11] P. Pant, P. Koirala, M. Hauta-Kasari, and J. Parkkinen. Estimating Color Signal at Different Correlated Color Temperature of Daylight. In *Advanced Concepts for Intelligent Vision Systems*, pages 587–597. Springer, 2009.

Author Biography

Graham D. Finlayson obtained his BSc in Computer Science from the University of Strathclyde (Glasgow, Scotland) in 1989. He then pursued his graduate education at Simon Fraser University (Vancouver, Canada) where he was awarded his MSc and PhD degrees in 1992 and 1995 respectively. From August 1995 until September 1997, Dr. Finlayson was a Lecturer in Computer Science at the University of York (York, UK) and from October 1997 until August 1999 he was a Reader in Colour Imaging at the Colour & Imaging institute, University of Derby (Derby, UK). In September 1999 he was appointed a Professor in the School of Computing Sciences, University of East Anglia (Norwich, UK).

Stuart Lynch received his BSc in Computer Science in 2009 from the University of East Anglia, Norwich (UK). Since then he has been a PhD student at the University of East Anglia, Norwich (UK) under the supervision of Prof. Graham D. Finlayson.