Surface reflectance recovery under point light illumination

Robert B. Fisher Aristides P. Gionis

Machine Vision Unit, Department of Artificial Intelligence
University of Edinburgh
Email: {rbf,aristide}@aifh.ed.ac.uk

Abstract

In this paper, a novel algorithm for colour recovery is presented. It assumes that the 3-D geometry of the scene is known. The spectral power distribution of a point illumination source, and the response function of the sensor are calibrated jointly. This algorithm has been used for the colour recovery part of an integrated system, developed in our laboratory, for environmental modelling. The geometry of the scene is recovered using a laser stripe range-finder and this information is exploited by the colour recovery algorithm. A point light source, attached to the whole system, has been used for the illumination of the scene, in order to confine undesirable side-effects of the ambient light. The joint spectral power distribution of this point light source and the response function of the camera are obtained with off-line calibration.

1 Introduction

The objective of this paper is to present a new algorithm for obtaining surface colour information in a controled enviroment. The geometry of the surfaces depicted in the input intensity images is assumed to be known and the scene is illuminated by a point light source whose spectral power distribution has been jointly calibrated with the response function of the camera. The motivation for this work is the ability of humans to perceive the colour of a surface as an invariant characteristic of the surface. This ability is known as "colour constancy".

What is invariant for each surface, and depends only on its micro-structure, is the reflectance of the surface (or colour descriptors), which is the percentage of the incident light that be reflected from the surface as a function of wavelength. Therefore, the colour spectrum measured by a visual system (human eye or camera) is the product of its reflectance spectrum by the spectrum of the incident light.

The first attempt at explaining the human colour constancy mechanism was made by [8] and [11], who found that the vision system does not achieve perfect colour constancy but it emphasize the changes in colour under illuminant and reflectance changes. However, this mechanism is not yet well explained, but there is evidence that several simultaneous mechanisms contribute to it [1]. More recent computational approaches, having as input the RGB responses of a colour camera,

attempt to separate the reflectance spectrum of a surface from the illumination spectrum. However, without restrictions and limitations on the type of illuminance and the surfaces that will be encountered, colour constancy is not, in general, possible [15]. The different algorithms that have been developed usually impose some of these restrictions and/or require additional information to eliminate them.

Colour descriptors are powerful invariants that could be used for many applications: an autonomous lawn-mower that continues its work as clouds appear, an industrial system for object classification based on their colour, a colour-indexing matching system [16], or even a virtual reality system that renders an indoor environment with different illumination from the one that was present at the scene acquisition time [5].

The main points of the classical approach to the estimation of the surface reflectance (suggested by [15]) are that 1) the ambient light is used as illumination source and 2) a finite-dimensional linear model is used to represent both surface reflectance and illumination source. The approach presented in this paper has the main differences that 1) a point light source is used for the illumination of the scene instead of ambient light and 2) the combined illuminant and sensor system is calibrated on a pixel by pixel basis. This has been chosen in order to make neglectable side effects of the ambient light, such as shadows, mutual reflections, illuminant variations, and sensor sensitivity variations

2 Previous work on colour constancy

One of the first computational theories for colour constancy was the retinex theory [13], the main assumption of which is that it applies to the planar Mondrian world. The Mondrian world consists of patches of differently coloured paper pasted on a planar background. The illumination across the Mondrian world is constrained to vary smoothly. The main idea for the separation of the surface reflectance from the illumination is that sharp changes in intensity of the images correspond to edges between the patches, i.e. surface reflectance changes, while smoother changes of the colour signal correspond to illumination changes. By ignoring the small intensity changes of the image, and accepting only the large ones, the surface reflectance can be reconstructed.

More recent approaches to the problem make use of a finite-dimensional linear model ([15], [2], [14], [7], [9], [6]), in which the surface reflectance and the illumination are expressed as a weighted sum of a fixed set of basis functions. If we denote the surface spectral reflectance at the point x with $R(\lambda)$, we can express it as

$$R(\lambda) = \sum_{j=1}^{n} \rho_j R_j(\lambda)$$

where the $R_j(\lambda)$ basis reflectances must be considered fixed and known, the weights ρ_j are unknown, and the number of them, n, is referred to as the number of degrees of freedom in the model. Common values for the parameters of the model are $n = 3$ (corresponding to red, green, and blue response channels), and $\lambda = 400, \ldots 770$nm, (corresponding to the visible spectrum). In a similar way, the ambient light $I(\lambda)$

can be expressed by a linear model as

$$I(\lambda) = \sum_{i=1}^{m} \epsilon_i I_i(\lambda)$$

with fixed, known basis lights $I_i(\lambda)$, and weights ϵ_i.

The set of the basis functions for the reflectance and the illumination should be carefully chosen so as many materials and different illumination conditions can be expressed by them. Studies were conducted by Cohen [3] for the reflectance function and by Judd *et al.* [12] for the illumination function, and the resulting basis functions account for 99% of the overall variance.

Many algorithms have been developed for the estimation of the reflectance weights ρ_j, and illumination weights ϵ_i. Maloney and Wandell [14] showed how to estimate them and recover surface reflectance under a number of different conditions: 1) the illumination must be constant over a segment of the image, 2) sufficient colour information is available over the image, and 3) the number of sensor channels n should be greater than m (the number of basis functions in the illumination model). This last limitation is important because it implies that we have to use either $n = 4$ response channels for obtaining colour information in the image (while humans appear to need only $m = 2$ spectral sources), or $m = 2$ basis functions for the scene reflectance (which are too few to represent the usually richer ambient light).

Ho *et al.* [9] showed how one can avoid this "fourth" response channel by taking additional information by a finer sampling of the colour signal spectrum. A different method followed by Funt *et al.* [6] uses the mutual reflection of surfaces near the concave boundaries between these surfaces, as well as measurements from points that are not affected by mutual reflection in order to get the extra information needed for reconstructing the surface reflectance function.

One of the common assumptions of the methods mentioned above is that the incident illumination remains constant across the scene. Recently, Finlayson *et al.* [4] showed how to achieve colour constancy using two RGB measurements of the same surface under two spectrally distinct incident illuminations.

A different approach was followed by Hurlbert [10], who demonstrated how to synthesize the operator that transforms the reflectance to the image irradiance by associative learning from a set of examples. The algorithm has good performance in the restricted Mondrian world for which it was designed.

3 Description of the problem and its assumptions

As we said before, the algorithm presented in this paper has been used as the reflectance recovery subsystem of a hand-held environment modelling system [5]. The main purposes of this system is the 3-D geometry and reflectance recovery of the objects in the scene. The system consists of a camera, a laser striper, and a magnetic position sensor (Flock of Birds) that one can move rigidly on an arbitrary trajectory over the scene. As far as the reflectance recovery is concerned,

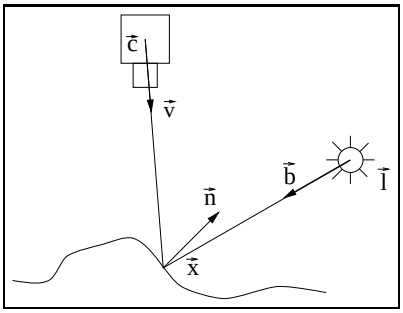

Figure 1: The point $\vec{x}(i,j)$ with surface reflectance $R_{\lambda,\vec{x}}$, and surface orientation $\vec{n}_{\vec{x}}(i,j)$ is projected to the pixel (i,j) of a pinhole camera at $\vec{c}$. The scene is illuminated by a point light source at $\vec{l}$ whose power distribution depends on the direction $\vec{b}$.

the arbitrary motion of the hand-held unit over the scene under the ambient illumination introduce many side-effects like shadows and mutual reflections. In order to eliminate these side-effects, a point light source has been attached on the hand-held unit and dominates the illumination conditions.

The main drawback using point light source is that the incident illumination power per area unit varies over the scene. More specifically:

- The incident illumination power per area unit at each scene point is proportional to the square of the distance between the surface point and the light source. Fortunately, the geometry of the scene is known and the distance between the points of the scene and the light source can be easily recovered.
- The incident illumination power per area unit at each surface point depends on the orientation of the surface at this point.
- The illumination power of the point light source is not uniformly distributed over the space. It would be uniform if an ideal point light was used, instead of a commercial halogen light bulb. To solve this problem, we illuminate a white planar surface and we build a map of the illumination power of the point light source as a function of space direction. Strictly speaking, this map contains information of the spatial distribution power of the light source, as it has been filtered by the camera's response function.

As we do not want the camera to dynamically adjust its response function, the automatic gain control (AGC) is disabled.

All 3-D vectors and points are defined in terms of a global coordinate system. Let us consider an intensity image; each pixel (i,j) of the image corresponds to a point of the real world $\vec{x}(i,j)$, with normal $\vec{n}_{\vec{x}}(i,j)$ (see Figure 1). The line between the optical centre of the camera $\vec{c}$, and the pixel (i,j) defines the viewing direction $\vec{v}(i,j)$ towards this point. The vector from the point light source $\vec{l}$ to the point $\vec{x}(i,j)$ defines the light direction $\vec{b}$, and the light intensity at this direction $\vec{b}$

at unit distance is $d_\lambda(\vec{b})$. The measured image intensity at each pixel is $m_\lambda(i,j)$, and the sensor's response function is $s_\lambda(i,j)$. Finally, the incident illumination at this point is $I_{\lambda,\vec{x}}$, and the surface reflectance $R_{\lambda,\vec{x}}$. The surface is assumed to be Lambertian.

Two assumptions concerning the camera gain mechanism are made: First, we assume that the camera gain is described by a gamma function law $g(x) = x^\gamma$, and second that the response function of the camera is uniform across all the pixels, i.e. $s_\lambda(i,j) = s_\lambda$, although this is generally not true.

4 Underlying theory

The incident irradiance on a surface point $\vec{x}(i,j)$ is :

$$I_{\lambda,\vec{x}} = d_\lambda(\vec{b}) \frac{1}{\|\vec{x}(i,j) - \vec{l}\|^2} |\vec{n}_{\vec{x}}(i,j) \cdot \vec{b}|$$

while, from the assumption that our world consists of Lambertian surfaces we have that the irradiance at the surface is (see Section 1):

$$E_{\lambda,\vec{x}} = R_{\lambda,\vec{x}} I_{\lambda,\vec{x}}$$

and the measured colour is:

$$m_\lambda(i,j) = s_\lambda(i,j) g(R_{\lambda,\vec{x}} I_{\lambda,\vec{x}})$$

Using the gamma function law for the gain function of the camera, and solving the above equations for $R_{\lambda,\vec{x}}$, we get:

$$R_{\lambda,\vec{x}} = \frac{(m_\lambda(i,j))^{\frac{1}{\gamma}}}{(s_\lambda(i,j))^{\frac{1}{\gamma}} d_\lambda(\vec{b})} \frac{\|\vec{x}(i,j) - \vec{l}\|^2}{|\vec{n}_{\vec{x}}(i,j) \cdot \vec{b}|}$$

and using the assumption that the response of the camera is uniform, we finally have:

$$R_{\lambda,\vec{x}} = \frac{(m_\lambda(i,j))^{\frac{1}{\gamma}}}{(s_\lambda)^{\frac{1}{\gamma}} d_\lambda(\vec{b})} \frac{\|\vec{x}(i,j) - \vec{l}\|^2}{|\vec{n}_{\vec{x}}(i,j) \cdot \vec{b}|} \quad (1)$$

Equation (1) shows how we can estimate the surface reflectance as a function of the measured colour responses $m_\lambda(i,j)$, the scene geometry ($\vec{x}(i,j)$, $\vec{n}_{\vec{x}}(i,j)$), the geometry of the hand-held arrangement ($\vec{c}$, $\vec{l}$), the direction and the spectral power distribution of the point light source ($\vec{b}$, $d_\lambda(\vec{b})$), and the characteristics of the camera ($\vec{v}(i,j)$, γ, s_λ).

5 Surface reflectance recovery

In this section, we discuss the estimation of all the parameters contained in the right-hand side of the Equation (1).

The 3-D coordinates of the scene points $\vec{x}(i,j)$, and the orientation vectors $\vec{n}_{\vec{x}}(i,j)$ of the surfaces at these points are estimated by the laser stripe triangulation range-finder, that is part of the hand-held scanner.

The vectors $\vec{v}(i,j)$ are estimated using a simple technique for camera calibration: Two calibration planes with many marked features of known coordinates in the world coordinate system are used. For each of these planes, a sample of pixels that correspond to known 3-D coordinates can be found (i.e. the ones that correspond to the marked features). The data sample of each plane is fitted to all pixels using a fourth degree polynomial approximation. Using two calibration planes, we obtain two points on the optical ray of each pixel, and therefore the vector $\vec{v}(i,j)$ itself.

We compute the optical centre $\vec{c}$ of the camera as the intersection point of the vectors $\vec{v}(i,j)$ (in a least squares sense). The position of the point light source $\vec{l}$ is physically measured.

The gamma factor γ of the gain function of the sensor is estimated using images of a white plane at different known distances of the camera. Measuring the average intensity drop-off over the distance we can solve for finding the γ factor that minimises the error between the actual and the estimated intensity values.

The light direction $\vec{b}$ can be easily estimated from:

$$\vec{b} = \frac{\vec{x}(i,j) - \vec{l}}{\|\vec{x}(i,j) - \vec{l}\|}$$

The spectral power distribution $d_\lambda(\vec{b})$ of the point light source as a function of the light direction $\vec{b}$ is estimated as it is filtered by the camera's response function, i.e. the compound term $(s_\lambda)^{\frac{1}{\gamma}} d_\lambda(\vec{b})$ is computed as it appears in the Equation (1). This is done with off-line calibration: an image of a white Lambertian plane at known distance and orientation to the sensor is taken. If we denote by $M_\lambda(i,j)$ the measured colour signal of the camera to this white plane, we know that, by definition, $R_{\lambda,\vec{x}} = 1$, and therefore:

$$1 = \frac{(M_\lambda(i,j))^{\frac{1}{\gamma}}}{(s_\lambda)^{\frac{1}{\gamma}} d_\lambda(\vec{b})} \cdot \frac{\|\vec{x}(i,j) - \vec{l}\|^2}{|\vec{n}_{\vec{x}}(i,j) \cdot \vec{b}|}$$

Solving, we obtain:

$$(s_\lambda)^{\frac{1}{\gamma}} d_\lambda(\vec{b}) = (M_\lambda(i,j))^{\frac{1}{\gamma}} \frac{\|\vec{x}(i,j) - \vec{l}\|^2}{|\vec{n}_{\vec{x}}(i,j) \cdot \vec{b}|}$$

The direction vector $\vec{b}$ is expressed in terms of its spherical coordinates (ϕ, θ) ($|\vec{b}| = 1$), and a map is built that approximates in a locally linear way the function $(s_\lambda)^{\frac{1}{\gamma}} d_\lambda(\vec{b})$ upon the (ϕ, θ) indices.

Setting the camera's measured colour signal received from a Lambertian white plane equal to 1 implies that we choose the surface reflectance of this plane to represent the reflectance of what we interpret as white. However, this is not an actual problem, since in any case the absolute brightness of the scene's surfaces cannot be recovered: colour constancy is always recovered up to a scale factor.

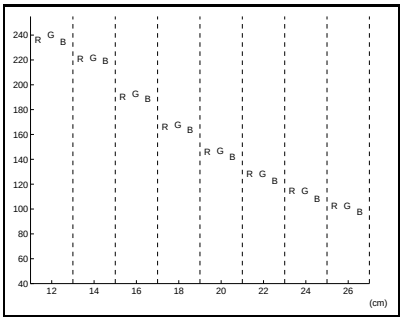

(a) Measured intensity of images of a surface at different distances

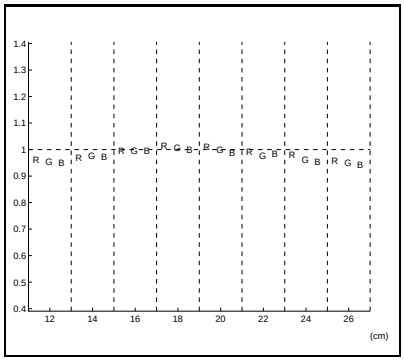

(b) Estimated reflectance of images of a surface at different distances

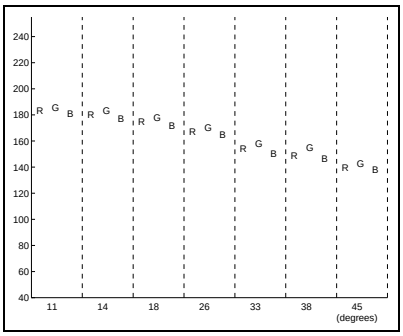

(c) Measured intensity of images of a surface at different orientations

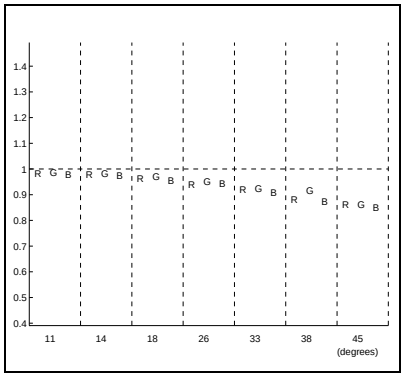

(d) Estimated reflectance of images of a surface at different orientations

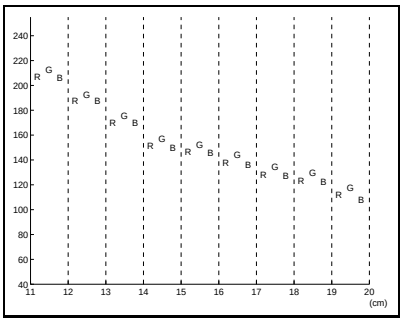

(e) Measured intensity of images of a specular image at different distances

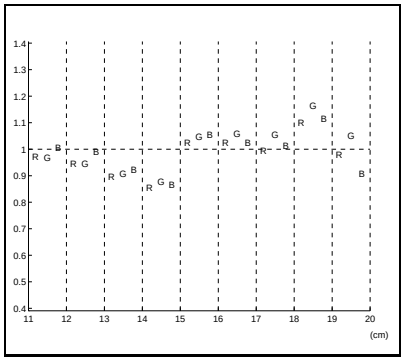

(f) Estimated reflectance of images of a specular image at different distances

Figure 2: Figure (a) shows the RGB intensity of a white surface as measured at the raw images over images at different distances. Figure (b) shows the estimated reflectance of the same images. The same experiment is presented for images at different orientations in figures (c) and (d), and for a specular surface in Figures (e) and (f). The range of distances is limited between 12 and 20 cm because outside this range of distances the pixel values of the intensity images are near 0 or near 255 and they can not be used reliable.

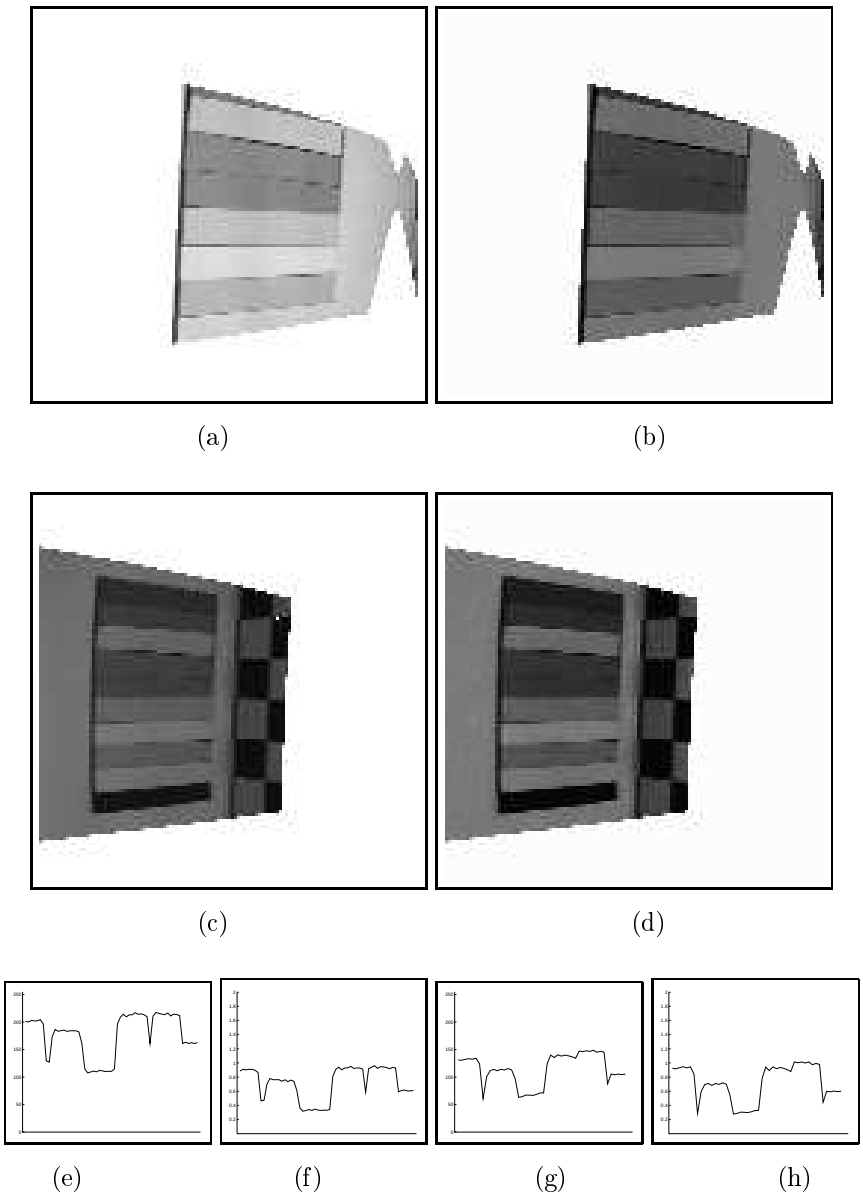

Figure 3: Images (a) and (c) shows the raw images of the same coloured surface at two different positions. Images (b) and (d) show the recovered reflectance of the surface in each case. The quality of the estimated reflectance is counted in terms of the invariance of the RGB values of the images (b) and (d). Figures (e)-(h) correspond to images (a)-(d) respectively, and show the measured intensity and the estimated reflectance across a vertical line cutting the same coloured stripes of the tested board. These values correspond only to the red channel of the camera, but similar results appear for the other two channels.

6 Experimental results

In order to test the validity of the algorithm independent of the quality of the scene geometry obtained by the range-finder, we estimate the surface reflectance of planar surfaces for which we can find analytically the position and the orientation of their points. Dealing only with planar surfaces does not affect the quality of the reported results, since all the calculations are local.

A number of images of some planar surfaces is taken at different distances and orientations. Changes of the illumination falling on the images caused differences in the measured intensity of regions that correspond to the same surface. Since we did not have any way to precisely measure the real reflectance of the tested surfaces, the quality of the results is judged by comparing the surface reflectance values estimated from each image. In the ideal case they should be the same, expressing the invariant characteristic of the surface reflectance.

Figure 2(a) shows the RGB signal intensity as measured in the raw images of a planar white surface at different distances and the same orientation. The range of distances is 12-20cm. Outside this distance range the images became too bright or too dark and the reflectance recovery algorithm is not effective. In this figure, we can see that the measured RGB signal intensity drops as the distance increases. The RGB reflectance estimated by the algorithm is shown in Figure 2(b).

The same experiment is repeated for images the surface at different orientations and the same distance. The results are shown in Figures 2(c) and 2(d). This time, the recovery is not quite as successful and this is probably happens because of the imperfect Lambertian nature of the used surface. Other materials of different surface type and colour were tested and similar results were found, with poorer results correspond to the less Lambertian nature of surfaces. An example is presented in the Figures 2(e) and 2(f), where the reflectance of a quite specular surface is recovered in an obviously unreasonable way.

Finally, Figure 3 gives a visual percept of the algorithm in action. Images (a) and (c) show two different images of a surface consisting of many colour patches. Images (b) and (d) present the estimated reflectance of the surface in each case. The invariance of the reflectance for each colour patch is apparent, even given the gray-level print out of the images.

7 Discussion conclusion

We have presented a new approach to estimate the reflectance of Lambertian surfaces. A light has been used for the illumination of the scene and the theoretical analysis shows how the reflectance of the objects in the scene can be recovered as a function of the 3-D scene geometry, the direction and the power of the incident light to the objects, and the gain mechanism of the camera. It turns out that good results can be obtained with careful calibration of the hand-held unit geometry and the spectral power distribution of the point light source. The estimated reflectance is quite sensitive to the γ factor of the gain function of the camera, and better estimation of the γ value can further improve the results.

Acknowledgements Thanks to A.W. Fitzgibbon and M. Wright for their helpful ideas and suggestions during the implementation. Also to M. Pilu for his comments on the manuscript. This work was supported by EPSRC grant GR/H/86905. The second author is funded by a Bakalas Foundation scholarship.

References

[1] K.T. Blackwell and G. Buchsbaum. Quantitative studies of colour constancy. *J. Opt. Soc. America A*, 5:1772–1780, 1988.

[2] M.H. Brill and G. West. Contributions to the theory of invariance of color under condition of varying illumination. *J. Math. Biol.*, 11:337–350, 1981.

[3] J. Cohen. Dependency of the spectral reflectance curves of the munsell colour chips. *Psychonomic Sci.*, 1:369–370, 1964.

[4] G.D. Finlayson, B.V. Funt, and K. Barnard. Color constancy under varying illumination. In *5th ICCV*, pages 720–725. IEEE Computer Society Press, June 1995.

[5] R.B. Fisher, A. Fitzgibbon, A. Gionis, M. Wright, and D. Eggert. A hand-held optical surface scanner for environmental modeling and virtual reality. In *Proc. Virtual Reality World '96*, February 1996.

[6] B.V. Funt, M.S. Drew, and J. Ho. Colour constancy from mutual reflection. *IJCV*, 6(1):5–24, 1991.

[7] R. Gerson. *The use of colour in computational vision.* PhD thesis, University of Toronto, 1988.

[8] H. Helson. Fundamental problems in color vision. i. the principles governing changes in hue saturation and lightness of non-selective samples in chromatic illumination. *J. Exp. Psychol.*, 23(439), 1938.

[9] J. Ho, B.V. Funt, and M.S. Drew. Separating a colour signal into illumination and surface reflectance components: Theory and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(10):966–977, 1990.

[10] A.C. Hurlbert. Neural network approaches to color vision. In Harry Wechsler, editor, *Neural Networks for Perception*, chapter II.15, pages 265–284. Academin Press, 1992.

[11] D.B. Judd. Hue saturation and lightness of surface colors with chromatic illumination. *J. Opt. Soc. America*, 30(2), 1940.

[12] D.B. Judd, D.C. MacAdam, and G. Wyszecki. Spectral distribution of daylight as a function of correlated color temperature. *J. Opt. Soc. America*, 54:1031–1040, 1964.

[13] E.H. Land and J.J. McCann. Lightness and retinex theory. *J. Opt. Soc. America*, 61:1–11, 1971.

[14] L.T. Maloney and B.A. Wandell. Color constancy: A method for recovering surface spectral reflectance. *J. Opt. Soc. America A*, 3:29–33, 1986.

[15] P. Sällström. Colour and physics: some remarks concerning the physical aspects of human colour vision. Technical Report 73-09, Institute of Physics Rep., University of Stockholm, Stockholm, 1973.

[16] Michael J. Swain. Color indexing. Technical report, Dep. of Computer Science, University of Rochester, Rochester, NY, 1990.