

Dichromatic Illumination Estimation via Hough Transforms in 3D

Lilong Shi and Brian Funt; Simon Fraser University; Vancouver, BC., Canada

Abstract

A new illumination-estimation method is proposed based on the dichromatic reflection model combined with Hough transform processing. Other researchers have shown that using the dichromatic reflection model under the assumption of neutral interface reflection, the color of the illuminating light can be estimated by intersecting the dichromatic planes created by two or more differently coloured regions. Our proposed method employs two Hough transforms in sequence in RGB space. The first Hough Transform creates a dichromatic plane histogram representing the number of pixels belonging to dichromatic planes created by differently coloured scene regions. The second Hough Transform creates an illumination axis histogram representing the total number of pixels satisfying the dichromatic model for each posited illumination axis. This method overcomes limitations of previous approaches that include requirements such as: that the number of distinct surfaces be known in advance, that the image be pre-segmented into regions of uniform colour, and that the image contain distinct specularities. Many of these methods rely on the assumption that there are sufficiently large, connected regions of a single, highly specular material in the scene. Comparing the performance of the proposed approach with previous non-training methods on a set of real images, the proposed method yields better results while requiring no prior knowledge of the image content.

Introduction

In this paper, we will focus our attention on the problem of estimating the color of the light by exploiting the principles of color image formation laid down by the dichromatic reflection model [1]. It states that, in RGB space, the colors reflected by an inhomogeneous dielectric material lie on a plane that is spanned by two characteristic colors; namely, the color of the specular component reflected from the air-surface interface, and the color reflected from the body of the material. If neutral interface reflection is assumed [2], then the chromaticity of the specular reflection is the same as that of the illuminating light. As a result, the color of the illuminant can be estimated by intersecting the planes that the set of RGBs from two or more different materials describe.

Based on the dichromatic reflection model [1], Lee [3] introduced a method for computing the scene illuminant chromaticity by intersecting lines in chromaticity space. Although Lee's method performs sufficiently well on synthetic images of spheres, its application to real-world scenes is sensitive to noise and inhomogeneities such as textured surfaces. Another approach using dichromatic regions of different colored surfaces, is called color line search [4]. It involves automatic detection of specular regions, a Hough transform step, and consistency check step. However, this approach requires correct detection of regions of interest, and can fail when specular highlights are incorrectly identified or absent from the scene. The method proposed by Tan et. al.[5] describes an inverse-intensity chromaticity space in which the

correlation between illumination chromaticity and image chromaticity can be analyzed. Once again, this method relies on correctly identifying the highlight regions, and does not perform any better than competing methods. The proposal to solve for the intersection of the dichromatic planes directly as described by Toro et. al.[6] assumes that in any patch of the given image, a fixed number of different materials coexist. The illumination colour can be calculated by solving a set of simultaneous linear equations using a Veronese projection of multilinear constraints. However, this approach assumes that the number of different surfaces in an image is already known. It also does not yield any better results than previous methods when applied to real images. The method proposed by Schaefer in [7] achieves competitive results, but the approach requires the illumination to be from a set of known light sources.

In this paper we proposed a robust method for determining the illumination axis. The method detects dichromatic planes while placing few restrictions on the image content, such as the number of surfaces, the surface colours, or the identification of specular regions. The approach involves two Hough Transforms in sequence that result in a histogram representing the likelihood that a candidate intersection line is the image illumination axis. The final illumination estimate is determined by intelligently choosing from amongst the most likely candidates.

Dichromatic Reflection Model

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 reflection and the other describing the body reflection part [1]. This can be expressed as

$$C(\theta, \lambda) = m_I(\theta)C_I(\lambda) + m_B(\theta)C_B(\lambda) \quad (1)$$

where $C_I(\lambda)$ and $C_B(\lambda)$ are the spectral power distributions of the interface and the body reflection 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.

Suppose R , G , and B are the red, green, and blue pixel value outputs of a digital camera, then each color vector $(R, G, B)^T$ is determined by a linear combination of a surface reflection component $(R_i, G_i, B_i)^T$ and a body reflection $(R_b, G_b, B_b)^T$ component. Equation 2 shows that the colour signal can be expressed as the weighted sum of these two reflectance components. Thus the colour signals for an object are restricted to a plane.

$$\begin{pmatrix} R \\ G \\ B \end{pmatrix} = w_i \begin{pmatrix} R_i \\ G_i \\ B_i \end{pmatrix} + w_b \begin{pmatrix} R_b \\ G_b \\ B_b \end{pmatrix} \quad (2)$$

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. Although theoretically sound, dichromatic colour constancy algorithms do not always perform well on real images. For example, image noise may cause the intersection of two dichromatic lines planes to change quite drastically. In addition, textured and non-uniform surfaces may mean that the distribution of colours does not lie on distinguishable dichromatic planes.

Method

We use a Hough Transform [8] for dichromatic plane detection. In the 3D case, a plane is parameterized as:

$$\rho = (x \cos(\phi) + y \sin(\phi)) \cos(\theta) + z \sin(\theta) \quad (3)$$

where ρ is the distance between a plane and the origin, ϕ is angle relative to the z axis, θ is angle relative to the y axis. In the discrete case, the parameter space (ρ, ϕ, θ) is quantized into bins, so the Hough Transform is represented as a three-dimensional histogram.

According to the dichromatic model, all dichromatic planes should pass through the origin. This implies that the “distance” ρ in Equation 3 is zero, so the RGBs reflected from a dichromatic surface satisfy the following parametric plane equation

$$R \cos(\phi) \cos(\theta) + G \sin(\phi) \cos(\theta) + B \sin(\theta) = 0 \quad (4)$$

All pixels from the same surface belong to a single plane defined by the two angles ϕ and θ . Hence a 2D Hough Transform can be used to create a *dichromatic plane histogram* \mathbf{H}_1 . Each bin of the histogram represents the number of pixels belonging to a distinct dichromatic plane specified by the pair of angles (ϕ, θ) satisfying Equation 4. A high value in the histogram implies the existence of this dichromatic plane in the image, while a lower value implies its absence. Two examples of *dichromatic plane histograms* are shown in Figure 1(b) and (e).

Since the illumination axis is the intersection of all dichromatic planes [6], it must be perpendicular to the normal of each dichromatic plane. Therefore, the axis perpendicular to the normals of the largest number of dichromatic planes is a good candidate for the illumination axis. To determine it, we use a second Hough Transform to create an *illumination histogram* \mathbf{H}_2 based on the data from \mathbf{H}_1 . To use the data from \mathbf{H}_1 , we first calculate the normals of the dichromatic plane in the dichromatic plane set. The normal of a dichromatic plane described by (ϕ, θ) is $\mathbf{n} = (u, v, w)$ where

$$\begin{aligned} u &= \cos \theta \cos \phi \\ v &= \cos \theta \sin \phi \\ w &= \sin \theta \end{aligned} \quad (5)$$

When an illumination axis is represented in polar form by the two angles α and β , it is perpendicular to the normal \mathbf{n} of a dichromatic plane if and only if it satisfies the following equation

$$u \cdot \cos(\beta) \cos(\alpha) + v \cdot \sin(\beta) \cos(\alpha) + w \cdot \sin(\alpha) = 0 \quad (6)$$

Based on Equation 6, a second 2D Hough Transform parameterized by (α, β) is used to create an *illumination histogram*. The count for a bin in the *illumination histogram* is calculated in the following manner. When the normal of dichromatic plane (ϕ, θ) is perpendicular to illumination axis (α, β) , the count from the corresponding bin of \mathbf{H}_1 is added to that of the corresponding bin of \mathbf{H}_2 . The bin count of a bin \mathbf{b} in the resulting histogram indicates the number of image pixels that conform to the dichromatic model under the illumination that \mathbf{b} represents in that contributing pixels all come from a collection of dichromatic planes that share a common intersection, and a common intersection represents a shared illumination. Therefore, a high bin count implies a high probability that the bin corresponds to the true scene illumination. Figures 1(c) and (f) provide two examples of *illumination histograms* for the same object under two different illuminations.

In principle, the correct illumination can be determined by searching for the global maximum in the *illumination histogram*. However, due to noise and the non-dichromatic properties of some surfaces that may be present in the image, the global maximum of the *illumination histogram* does not always correctly correspond to the true scene illumination. Although the global maximum may not always indicate the correct illumination, generally one of the local maxima will. Hence the problem becomes how to select between the local maxima. Our strategy is to select the local maximum inside a bounding disc centered at the illumination as estimated by another illumination-estimation method. In particular, in the experiments reported here we use the Shades of Gray (SoG) method [9]. The disc radius is based on the average and standard deviation of the error of this method.

In summary, the complete estimation consists of the following steps:

1. Normalize the image \mathbf{I} (scale intensities, remove dark pixels, etc)
2. Transform 3D pixels in \mathbf{I} into the dichromatic Hough space \mathbf{H}_1 using Equation 5
3. Transform \mathbf{H}_1 into illumination Hough Space \mathbf{H}_2 using Equation 6 and 7
4. Estimate image illumination L by SoG
5. Find the nearest local maximum in \mathbf{H}_2 inside a bounding disc centered at L
6. Convert polar coordinate representation of this local maximum into chromaticity coordinates

In summary, the two Hough Transforms can be thought of as two voting procedures. First, each pixel votes for the candidate dichromatic planes that pass through it. Second, each dichromatic plane in turn casts a weighted vote (weighted by the number of pixels on that plane) for each candidate illumination axis that passes through it. Finally, the illumination axes that receive the highest votes are considered likely candidates for the true illumination.




Figure 1. Two images of the same object under different illuminants. (a) and (b) are the original images; (b) and (c) are the dichromatic plane histograms of the images in (a) and (b) respectively, after the first Hough Transform, with ϕ and θ ranging from 0 to 179; (c) and (d) are the illumination histograms of the images in (a) and (b) respectively, after the second Hough Transform, with α and β range from 0 to 89. The arrows in the two figures indicate the locations of the true illuminants. The correspondence between the true illumination and the histogram peaks is evident.

Methods	Training Required	Median Angular Error
Grey World	no	7.0
Max-RGB	no	6.5
Multilinear Constraint	no	5.8
Shades of Grey (n=6)	no	3.7
Grey Edge	no	3.2
2nd order Grey Edge	no	2.7
GSI	no	3.5
Color by Correlation	yes	3.2
Neural Networks	yes	7.8
TPS	yes	0.6
2D SVR	yes	4.7
3D SVR	yes	2.2
3D Hough Transform	no	1.7

Table 1. Comparison of performance of the proposed method with that of other non-training methods (Grey World, Max RGB, Multilinear Constraint, SoG, GSI, Grey Edge, 2nd order Grey Edge) and training methods (Color by Correlation, TPS, SVR 2D and SVR 3D, Neural Networks) measured in terms of median angular errors based on the SFU image dataset of 321 images. The entries for GW, Max RGB, SoG, GE, 2nd GE, CbyC, and NN are reproduced from Table II, page 2211 of [13].

Experimental Results

The proposed method was tested on the Simon Fraser University colour image database [10], which contains 321 images of 32 scenes under 11 different illuminants. In our experiments, an image is resized to 200x200 and normalized such that the range of intensity in any image is [0, 255], and then the first Hough transform is applied to all pixels (excluding pixels over 250 or under 10). The space of planes is defined by angles (ϕ, θ) , whose values are integers in [0 to 179]. The result of the first Hough transform, \mathbf{H}_1 , is therefore a 180x180 2D histogram as shown in Figure 1b and 2e. The

illumination axis space is defined by angles α and β with integer values in [0, 89]. Hence, the illumination histogram \mathbf{H}_2 calculated by a Hough transform of \mathbf{H}_1 is a 90x90 2D histogram (Figure 1c and 1f).

The performance is evaluated in terms of the angular difference in degrees between the RGB of the estimated and actual illumination. In Table 1, our approach shows good performance when compared to competing illumination-estimation methods [6,11-16].

Conclusion

We have presented an illumination-estimation method that uses the constraints provided by the dichromatic model in a new and quite robust way. The method is based on two Hough transform voting procedures. First, each image pixel votes for every dichromatic plane it could fall on. This results in a 2D histogram representing the likelihood of each plane. Second, each dichromatic plane votes for each candidate illumination axis that could pass through that plane. Finally, an illumination axis is chosen from among those receiving the highest number of votes based on the resulting illumination being close to that of the SoG illumination estimate.

In general, the total complexity of our method is $O(NM+MK)$, where N is the number of pixels in an image; M is the cardinality of the candidate dichromatic plane set; K is the cardinality of the candidate illumination axis set. In our experiment, the dichromatic planes and illumination axes were searched exhaustively with $M = 180 \times 180$ and $K = 90 \times 90$.

In conclusion, we proposed a robust method that creates a 2D illumination axis histogram that represents the likelihood of the possible illuminations. Our approach makes no assumption about the number of surfaces or the surface colours, yet performs well in comparison to the other methods tested.

References

- [1] S.A. Shafer, "Using color to separate reflection components", *Color Res. App.*, 10, 4 (1985).
- [2] H.-C. Lee, E. J. Breneman, and C. P. Schulte. "Modeling light reflection for computer color vision", *IEEE Trans. Pattern Anal. Machine Intell.*, 12, 4, 402–409, (1990).
- [3] H.-C. Lee. Illuminant color from shading. In *Proc. SPIE: Perceiving, Measuring, and Using Color*, 1250, 236–244, (1990).
- [4] T. M. Lehmann and C. Palm, "Color Line Search for Illuminant Estimation in Real-World Scenes", *J. Opt. Soc. Am. A*, 18, 2679–2691 (2001).
- [5] R. T. Tan, K. Nishino, and K. Ikeuchi, "Color constancy through inverse-intensity chromaticity space," *J. Opt. Soc. Am. A* 21, 321–334 (2004)
- [6] J. Toro and B. Funt, "A Multilinear Constraint on Dichromatic Planes for Illumination Estimation", *IEEE Transactions on Image Processing*, 16, 1, 92–97 (2007)
- [7] G. Schaefer. "Robust dichromatic colour constancy", In *Int. Conf. Image Analysis Recogn.*, 2, 257–264 (2004)
- [8] K. Okada, S. Kagami, M. Inaba, H. Inoue, Plane segment finder: algorithm, implementation and applications, In *Proc. of Int. Conf. on Robotics and Automation*, 2120–2125, (2001).
- [9] G. D. Finlayson and E. Trezzi, Shades of Gray and Colour Constancy, in *Proc. IS&T/SID Twelfth Color Imaging Conference: Color Science, Systems and Applications*, Society for Imaging Science and Technology, Scottsdale, 37–41, (2004).
- [10] K. Barnard, L. Martin, B. Funt, and A. Coath, "A data set for colour research," *Color Research and Application*, 27, 3, 147–151 (2002).
- [11] K. Barnard, L. Martin, A. Coath, and B. Funt, "A comparison of computational color constancy algorithms-part ii: Experiments with image data," *IEEE Trans. on Image Processing*, 11, 9, 985–996 (2002).
- [12] S. Hordley and G. Finlayson, "Re-evaluation of color constancy algorithm performance," *The Journal of the Optical Society of America A.*, 23, 5, 1008–1020 (2006).
- [13] J. van de Weijer, Th. Gevers, A. Gijsenij, "Edge-Based Color Constancy", *IEEE Trans. Image Processing*, 16, 9, 2207–2214 (2007).
- [14] W. Xiong, B. Funt, L. Shi, S. Kim, B. Kang, and S.D. Lee, Automatic White Balancing via Gray Surface Identification , *Proc. of the Fifteenth IS&T Color Imaging Conference*, Albuquerque, 143–146, (2007).
- [15] W. Xiong, , L. Shi, B. FuntS. Kim, B. Kang, and S.D. Lee, "Illumination Estimation via Thin-Plate Spline Interpolation," *Proc. of the Fifteenth IS&T Color Imaging Conference*, Albuquerque, 25–29, (2007).
- [16] B.V. Funt, and W.H. Xiong. Estimating Illumination Chromaticity via Support Vector Regression, in *Proceedings of the IS&T/SID Twelfth Color Imaging Conference: Color Science, Systems and Applications*, 47–52, (2004).

Author Biography

Lilong Shi received the B.Sc. degree and the M.Sc. degree in computer science from the Simon Fraser University, Vancouver, Canada, in 2003 and 2005. He is currently a PhD candidate in Simon Fraser University, Vancouver, Canada. His research interests are in illumination estimation, color constancy, and image processing.

Brian Funt received the Ph.D. degree in computer science from the University of British Columbia Vancouver, BC, Canada, in 1976. He is a Professor with the School of Computing Science, Simon Fraser University, Vancouver, where he has been since 1980. He has published widely in color, color constancy, retinex, color calibration, and illumination estimation.