Illumination Estimation from Dichromatic Planes

Javier Toro and Djemel Ziou, Departement d'Informatique, Universite de Sherbrooke, Sherbrooke, QC, Canada
Brian Funt, School of Computing Science, Simon Fraser University, Burnaby, BC, Canada

Abstract

Adopting the dichromatic reflection model under the assumption of neutral interface reflection, the color of the illuminating light can be estimated by intersecting the planes that the color response of two or more different materials describe. From the color response of any given region, most approaches estimate a single plane on the assumption that only a single material is imaged. This assumption, however, is often violated in cluttered scenes. In this paper, rather than a single planar model, several coexisting planes are used to explain the observed color response. In estimating the illuminant, a set of candidate lights is assessed for goodness of fit given the assumed number of coexisting planes. The candidate light giving the minimum error fit is then chosen as representative of the scene illuminant. The performance of the proposed approach is explored on real images.

Introduction

A number of methods for color constancy have been proposed. Major approaches have been classified either as statistics-or physics-based techniques.[4] In any case, the goal is to discount the color of the illuminating light in order to obtain an invariant description of the color of observed surfaces. In this paper, we will focus our attention on the problem of estimating the color of the light exploiting the principles of color image formation laid down by the dichromatic reflection model.[12] An advantage of approaches exploiting the dichromatic model is that very few different surfaces are required in order to reach a conclusion about the color of the light. These approaches follow from the dichromatic hypothesis that the color response of an inhomogeneous dielectric material lies on a plane in RGB space,[9] and that for a material with neutral interface reflection,[8] this plane contains the color of the light. As a result, the color of the illuminant can be estimated by intersecting the planes that the color response of two or more different materials describe.

There have been several previous approaches to utilizing dichromatic planes for illuminant estimation,[3,7,10,11,13] However, an unresolved issue is how to estimate reliably the plane where the colors of each observed material rest. In estimating a plane, we must necessarily pool information from a region of the given image. The problem is that it is not known *a priori* what the extent of a material-coherent region is, and this makes it difficult to prevent incompatible information (i.e., the information provided by colors displayed by different materials) from being combined. In some approaches, it is assumed that the extent of each region in the scene is known[13] or that in any given small region of the image only one material is imaged.[3] Other approaches attempt to estimate the extent of these coherent regions prior to the calculation of the planes (or in some cases, lines in chromaticity space).[7,10] Recently, a mechanism aimed at deciding when an estimated plane is likely to produce a reliable light estimate has been incorporated[11] and leads to a noticeable improvement. In general, these methods rely on the assumption that there are sufficiently large connected regions of a single material in the scene. This assumption, however, may be violated in highly cluttered scenes as pointed out by Finlayson and Schaefer in Ref. [3].

In this paper, we assume that in any patch of the given image, a fixed number of different materials do coexist. As such, in order to account for the multiplicity of materials, several planes (a number equal to that of assumed coexisting materials) are used to describe the observed colors, whatever the properties of the materials may be. In this description all planes intersect at the same line, which is defined by the color of the light. The color of the illuminant is then estimated by searching for the light that, under the assumed number of planar models, provides the best description for the colors displayed by the patch. The best description is obtained if there is an arrangement of the planes such that every color from the patch lies on at least one plane.

A brief description of the dichromatic reflection model is given in the next section. Using this model, a general constraint on the observed colors under any given illuminant is formulated in the section entitled *A Global Constraint*. In the section on *Solving for the Illuminant*, it is shown how this constraint can be used to estimate the color of the light. A quantitative evaluation of the proposed method and some final remarks are given in the *Experimental Results* and the *Concluding Remarks*.

Inhomogeneous Dielectric Materials

Under the dichromatic reflection model.[12] the color of an inhomogeneous material is written as a linear combination of two independent colors: the colors of the diffuse and the specular reflections of the material. If the spectral properties of the light falling on the material remain unchanged across the scene, the colors displayed by the material can be written as

$$\bar{c}(\bar{x}) = m_s(\bar{x})\bar{c}_s + m_b(\bar{x})\bar{c}_b, \tag{1}$$

where $\bar{x}$ stands for pixel position, $\bar{c}_s$ and $\bar{c}_b$ are the colors of the specular and the diffuse reflections, and $m_s(\bar{x})$ and $m_b(\bar{x})$ are scalar functions that capture the dependencies of the model on the geometry of the viewing situation. The color response $\bar{c}(\bar{x})$ thus lies on a plane in RGB space (referred to as the dichromatic plane of the material) described by vectors $\bar{c}_s$ and $\bar{c}_b$. Moreover, if the specular reflection of the observed material is assumed to be neutral,[8] the chromaticity of color $\bar{c}_s$ is the same as that of the illuminating light.

When several different materials are observed, and provided that material changes are abrupt, the color vector $\bar{c}_b$ alone becomes a piecewise constant function of position. Then instead of a single plane, the color response $\bar{c}(\bar{x})$ will reside in a collection of planes all intersecting at the same line give by vector $\bar{c}_s$. The method proposed here rests on these assumptions.

A Global Constraint

Equation (1) indicates how the colors of an inhomogeneous dielectric material distribute in color space. When several materials are observed, the color response lies on planes all intersecting at the same line defined by the color of the light. The light can therefore be viewed as the axis on which the planes pivot. Accordingly, any of the dichromatic planes can be written in terms of a single parameter (an angle) that describes the orientation of the plane. Indeed, a plane turning around $\bar{c}_s$ can be specified by its normal vector (which we denote as $\bar{u}(\gamma,\bar{c}_s)$), orthogonal to $\bar{c}_s$, given as

$$\bar{u}(\gamma,\bar{c}_s) = R(\bar{c}_s)[\cos(\gamma)\sin(\gamma)0]^T, \tag{2}$$

where γ is the orientation-related parameter of the plane and $R(\bar{c}_s)$ is a rotation matrix that depends on the color of the light. Varying γ the vector on the right-hand side of Equation (2) defines a family of planes all turning on the b-axis. Rotation matrix $R(\bar{c}_s)$ changes the axis of rotation to that specified by vector $\bar{c}_s$, as shown in Figure 1.

The desired coordinate transformation rotation matrix, $R(\bar{c}_s)$, must bring the b-axis, represented by vector $[001]^T$, in correspondence with the axis specified by the unit vector $\hat{c}_s = \bar{c}_s/|\bar{c}_s|$. Such a transformation can be written as $R(\hat{c}_s) = [\hat{b}_1 \hat{b}_2 \hat{c}_s]$, , where $\hat{b}_1$ and $\hat{b}_2$ are chosen as normal vectors orthogonal to one another and to $\hat{c}_s$. Clearly, the set $\{\hat{b}_1 \hat{b}_2 \hat{c}_s\}$ forms an orthonormal basis of R^3. A real square matrix whose columns are an orthonormal basis of its definition space is orthogonal. An orthogonal matrix defines a linear transformation which preserves the inner product and therefore is a rotation, a reflection or a combination of the two. In particular, for matrix $R(\bar{c}_s)$ to describe a proper rotation, the basis $\{\hat{b}_1, \hat{b}_2, \hat{c}_s\}$ must be positively oriented. We note, however, that this last requirement is of no major consequence for our purposes and can therefore be relaxed. An improper rotation does not change the axis on which the dichromatic planes pivot, only their direction of spin. To actually produce $R(\bar{c}_s)$ we must procure $\hat{b}_1$ and $\hat{b}_2$. There are several different ways in which these vectors can be obtained. For instance, we can arbitrarily fix two of the entries of $\hat{b}_1$ and obtain the third from the orthogonality of this vector with $\hat{c}_s$. Vector $\hat{b}_1$ and $\hat{b}_2$ can then be obtained as the cross product of $\hat{c}_s$. Alternatively, matrix $R(\bar{c}_s)$ can be obtained as a composition of two elementary rotations (an elementary rotation being a rotation about one of the coordinate axes). If we write the color of the light as $\bar{c}_s = [r\cos(\theta)\sin(\phi)\ r\sin(\theta)\sin(\phi)\ r\cos(\phi)]^T$, the rotation matrix can then be written as $R(\bar{c}_s) = R_b(\theta)R_g(\phi)$, which explicitly shows the dependence of $R(\bar{c}_s)$ on the angles θ and ϕ of vector $\bar{c}_s$. Matrices $R_b(\theta)$ and $R_g(\phi)$ are rotations about the b-and g-axes in an anti-clockwise direction by angles θ and ϕ, respectively, when looking towards the origin:

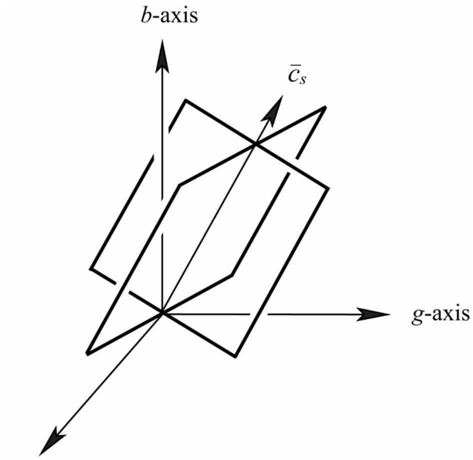

Figure 1. Dichromatic planes around color vector $\bar{c}_s$. The γ-dependent vector on the right-hand side of Equation (2) defines a plane tangential to the b-axis. cs)changes the tangential axis to that specified by vector Rotation matrix R($\bar{c}_s$). Two different planes, which correspond to two different values of γ, are shown in the figure. The actual value of parameter γ for a given plane depends on the specific derivation of matrix R($\bar{c}_s$), that derivation being carried out either using the properties of orthogonal matrices or the elementary rotations method.

$$R_b(\theta) = \begin{bmatrix} \cos(\theta) & -\sin(\theta) & 0 \\ \sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 1 \end{bmatrix}, \tag{3}$$

$$R_g(\phi) = \begin{bmatrix} \cos(\phi) & 0 & \sin(\phi) \\ 0 & 1 & 0 \\ -\sin(\phi) & 0 & \cos(\phi) \end{bmatrix}. \tag{4}$$

To derive the global constraint used here from the above restriction on the dichromatic planes, first suppose that the color response $\bar{c}(\bar{x})$ of a single material is observed. Since all colors must fall on a plane, it is not difficult to see that there is an angle γ such that the following expression is satisfied:

$$\bar{c}(\bar{x})^T \bar{u}(\gamma_i,\bar{c}_s) = 0\,. \tag{5}$$

This is the constraint equation for a single material in the scene. If several materials are observed, there is a constraint of the same form as that of Equation (5) on the colors of each of the observed materials, which will generally be satisfied by a different γ. Placing all these constraints together, if n different materials are observed, the totality of observed colors must therefore satisfy

$$\prod_{i=1}^{n} \bar{c}(\bar{x})^T \bar{u}(\gamma_i,\bar{c}_s) = 0, \tag{6}$$

for some $\gamma_1,\ldots,\gamma_n$. This is the constraint equation of the observed scene.

Solving for the Illuminant

In principle, vector $\bar{c}_s$ is not known. It is the quantity we actually want to solve for. Nevertheless, Equation (6) gives an avenue to the estimation of this vector. Intuitively, we see that if the observed color response is well behaved—where good behavior is understood as having a color response spanning actual planes and not lying on degenerate entities such as points or lines passing through the origin—then there is a unique (discarding the amplitude) light-related vector such that Equation (6) is satisfied. The equation can be used as a constraint on the chromaticity of the light. In the least-square sense, the estimation of the illuminant's chromaticity can be carried out as:

$$\bar{c}_{s_e} = \arg\min_{\bar{c}_s} \min_{\gamma_1,\ldots,\gamma_n} \sum_{\bar{x}} \left(\prod_{i=1}^{n} \bar{c}(\bar{x})^T \bar{u}(\gamma_i, \bar{c}_s) \right)^2, \tag{7}$$

with the additional constraint that $\bar{c}_s$ be a unit vector. This expression seeks the vector $\bar{c}_s$ that under the assumed number of planes, n, gives the best explanation to the observed colors. Note that in this expression the best explanation is taken as that configuration of planes that best fits the data, where goodness of fit is measured as the total deviation of the response values from the fitted planes (summed square of residuals).

In practice, to solve for the color of the illuminant, we assume that there is a fixed set of candidate lights to choose from. Given this set, we evaluate how well each light "fits" the observed color response. The candidate light giving the minimum error fit is then chosen as the illuminating light. More precisely, suppose that there are κ candidate lights, $\bar{c}_{s_1},\ldots,\bar{c}_{s_\kappa}$. For each $\bar{c}_s \in \{\bar{c}_{s_1},\ldots,\bar{c}_{s_\kappa}\}$, the fit is assessed by measuring the total residual error

$$\sum_{\bar{x}} \left(\prod_{i=1}^{n} \bar{c}(\bar{x})^T \bar{u}(\gamma_i, \bar{c}_s) \right)^2$$

resulting after all the different planar models have been fitted. The planes are fitted by using an iterative scheme in which the functional of Equation (7) is minimized with respect to a single γ at the time. That is, the planes are fitted by iteratively minimizing over the angles γ_k, $k = 1,\ldots,n$, the expression:

$$\min_{\gamma_k} \sum_{\bar{x}} w(\bar{x};k) \left(\bar{c}(\bar{x})^T \bar{u}(\gamma_k, \bar{c}_s) \right)^2, \tag{8}$$

where $w(\bar{x};k)$ is a weighting function that depends on previous γs estimates:

$$w(\bar{x};k) = \prod_{i \in \{1,\ldots,n\}/\{k\}} \left(\bar{c}(\bar{x})^T \bar{u}(\gamma_k, \bar{c}_s) \right)^2. \tag{9}$$

For a γ with no previous estimate, the corresponding dot product in Equation (9) is taken equal to one. The minimization of Equation (8) is iterated for all γs until their change is small. Observe that the expression of Equation (8) is almost the same as that of Equation (7) with the difference being that in (8) the only free parameter is γ_k.

We note that Equation (8) can be solved analytically. Substituting (2) into (8) and after some manipulation, it is not difficult to verify that the equation can be written as:

$$\min_{\gamma_k} \sum_{\bar{x}} \left(d_1(\bar{x}) \cos(\gamma_k) + d_2(\bar{x}) \sin(\gamma_k) \right)^2, \tag{10}$$

where $d_1(\bar{x})$ and $d_2(\bar{x})$ depend on $w(\bar{x};k)$ and $R(\bar{c}_s)$. More precisely, if the rotation matrix is written as $R(\bar{c}_s) = [\bar{r}_1 \bar{r}_2 \bar{r}_3]$, then $d_i(\bar{x}) = \sqrt{w(\bar{x};k)}\ \bar{r}_i^T \bar{c}(\bar{x})$, $i = 1, 2$. Differentiating the summation term of Equation (10) with respect to γ_k, setting the result equal to zero, and solving for γ_k yields:

$$\gamma_k = \frac{1}{2} \arctan \left(\frac{-2 \sum_{\bar{x}} d_1(\bar{x}) d_2(\bar{x})}{\sum_{\bar{x}} d_2(\bar{x})^2 - d_1(\bar{x})^2} \right). \tag{11}$$

Note that an additional check is needed in order to verify that the calculated γ_k corresponds to a minimum. If it is not, then the sought value is orthogonal to the found orientation.

So far, nothing has been said about the choice of the number of planar models to use in estimating the light. It is easy to see that in ideal conditions (i.e., when the observed colors strictly lie on dichromatic planes), the constraint of Equation (6) still holds even if the number of models used exceeds the actual number of different materials in the image. We therefore can formulate the constraint using a reasonably large number of models without worrying about the actual diversity of materials in the scene.

Experimental Results

To test the performance of the proposed approach we used the database produced at the Computational Vision Lab, Simon Fraser University.[2] This database contains images of 32 scenes taken under 11 different illuminants. The database is broken down into four subsets: *mondrian*, which is a group of images showing minimal specularities; *specular*, comprising images with non-negligible specularities; *metallic*, containing images with metallic specularities; and *fluorescent*, where the imaged scenes contain at least one fluorescent surface. A detailed description of this database can be found in Ref. [2].

In principle, the proposed scheme could be applied to the image as a whole. Here, we proceed differently. The image is first partitioned into blocks and the algorithm is applied to each of these blocks individually. For each block, the algorithm provides the fitting error of each candidate light (i.e., one of the 11 different illuminants used to construct the database). An overall assessment of how well the observed colors are described under a specific light is obtained by averaging the fitting errors across all image blocks under that specific illuminant. The candidate light that exhibits the best average fit is chosen as the estimate of the light. The reason for using such a procedure is that experimentally it provided better performance given a fixed number of planar models. This may be explained by the reduced complexity of the color response when our study is confined to a smaller subset of the total observations. In particular, the results reported here were obtained by partitioning the image into 200 × 200 non-overlapping blocks (each block covering an area about a sixth of the total image), and using four planar models to describe the colors of each patch.

Following the recommendation of Hordley and Finlayson[5] on the evaluation of color constancy algorithms, in addition to the mean, Tables 1 and 2 show the performance of the approach using the sample median and the corresponding confidence intervals for two different measures. The median, in particular, has been advocated as being a more appropriate estimate of the central tendency of measured errors. In Table 1, the performance is shown in terms of the Euclidean distance between the estimated and the actual light in the rg-chromaticity space. In Table 2, the performance is measured in terms of the angular difference in degrees between the two lights.

Table 1: Chromaticity error in rg-space. These results were obtained using four planar models to describe the colors in each 200 × 200 block. Confidence intervals (c. i.) were calculated using the routines produced by Kaplan.[6]

	mean	95% c. i.	median	95% c. i.
dataset	**0.056**	**0.050-0.060**	**0.044**	**0.037-0.046**
mondrian	0.043	0.036-0.051	0.027	0.019-0.036
specular	0.051	0.042-0.061	0.044	0.037-0.050
metallic	0.076	0.067-0.087	0.062	0.046-0.085
fluorescent	0.055	0.042-0.070	0.044	0.034-0.058

Table 2: Angular error in degrees. In these tests, four planar models were used to explain the colors of each 200 × 200 block.

	mean	95% c. i.	median	95% c. i.
Dataset	**7.83**	**7.17-8.54**	**6.05**	**4.30-7.64**
mondrian	5.97	4.94-6.99	3.76	2.23-4.30
specular	7.14	5.79-8.56	6.05	4.27-7.78
metallic	10.99	9.57-12.47	8.14	7.90-11.69
fluorescent	7.66	5.88-9.53	6.05	4.27-7.89

We note that in previously proposed techniques based on the dichromatic reflection model experimentation has been carried out either on the mondrian subset or on a database specifically generated to test the proposed approach. In the mondrian subset, strong specularities are not prevalent but nevertheless the specular component is sufficient to cause noticeable dichromatic behavior. Furthermore, since imaged surfaces in this subset are mostly flat or convex, interreflections, which may locally change the color of the illuminating light, are almost absent. This may explain why it is more likely to obtain from this set a more accurate estimate of the light than from the rest of the database, as shown in Tables 1 and 2. Here, we obtained statistics on the performance for the whole database looking at establishing a comparison of the proposed approach with major statistics-based techniques.[1] Clearly, the database contains images that violate the assumptions on color image formation upon which our model rests. Nevertheless, such an assessment may provide an idea of how the proposed method may perform when confronted with natural unconstrained scenes. The mean chromaticity error obtained for the whole database, shown in Table 1, compares well with those of statistics-based techniques, as reported in Ref. [1]. The confidence intervals in Tables 1 and 2 describe how sure the measured statistics are within the given upper and lower bounds. These results are a positive evidence on the potential of our proposed technique for estimating the color of the illumination in uncontrolled scenes.

Concluding Remarks

In this paper we have introduced a multi-linear constraint on the colors displayed by different inhomogeneous dielectric materials under a given illuminating light. We showed that this constraint can be used to estimate the illuminant from the set of observed colors. This is achieved by searching for the light that under the assumed number of planar models provides the best fit to the observed color response, where the goodness of fit is measured as the sum of squares of residuals. The experimental results on the SFU dataset, which includes scenes that do not strictly comply with the dichromatic hypothesis on color formation, show good performance in estimating the color of the illuminant.

References

1. K. Barnard, L. Martin, A. Coath, and B. Funt. A comparison of computational color constancy algorithms – part II: Experiments with images. IEEE Transactions on Image Processing, 11, 9, pp. 985–996 (2002).
2. K. Barnard, L. Martin, B. Funt, and A. Coath. A data set for color research. Color Research and Applications, 27, 3, pp. 148– 152 (2002). Available: http://www.cs.sfu.ca/~colour/data/colourconstancytestimages/index.html.
3. G. Finlayson and G. Schaefer. Convex and non-convex illuminant constraint for dichromatic color constancy. International Conference on Computer Vision and Pattern Recognition, vol. I, pp. 598–604 (2001).
4. G. Finlayson and G. Schaefer. Solving for colour constancy using a constrained dichromatic reflection model. International Journal of Computer Vision, 42, 3, pp. 127–144 (2001).
5. S. Hordley and G. Finlayson. Re-evaluating colour constancy algorithms. International Conference on Pattern Recognition, vol. 1, pp. 76–79 (2004).
6. D. Kaplan. Resampling Stats in Matlab. Available: http://www.macalester.edu/~kaplan/resampling/.
7. H.-C. Lee. Method for computing the scene-illuminant chromaticity from specular highlights. Journal of the Optical Society of America A, 3, 10, pp. 1694–1699 (1986).
8. H.-C. Lee, E. J. Breneman, and C. P. Schulte. Modeling light reflection for computer color vision. IEEE Transactions on Pattern Analysis and Machine Intelligence, 12, 4, pp. 402–409 (1990).
9. C.L. Novak and S. Shafer. Anatomy of a color histogram. International Conference on Computer Vision and Pattern Recognition, pp. 599–605 (1992).
10. V. Risson. Application de la Morphologie Math´ematique `a l'Analyse des Conditions d'Eclairage des Images Couleur. Ph.D. thesis, Ecole des Mines de Paris (2001).
11. G. Schaefer. Robust dichromatic colour constancy. International Conference on Image Analysis and Recognition, vol. 2, pp. 257–264 (2004).
12. S. A. Shafer. Using color to separate reflection components. Tech. Rep. 136, Department of Computer Science, University of Rochester, New York (1984).
13. S. Tominaga and B. A. Wandell. Standard surface-reflectance model and illuminant estimation. Journal of the Optical Society of America A, 6, 4, pp. 576–584 (1989).