Illumination Estimation Using a Multilinear Constraint on Dichromatic Planes

Javier Toro[1] and Brian Funt[2]

LBMU[1], Centre Hospitalier de l'Université de Montréal
Montréal, QC, Canada H2L 2W5

School of Computing Science[2], Simon Fraser University
Burnaby, BC, Canada V5A 1S6

Abstract

A new multilinear constraint on the color of the scene illuminant based on the dichromatic reflection model is proposed. The formulation avoids the problem, common to previous dichromatic methods, of having to first identify pixels corresponding to the same surface material. Once pixels from two or more materials have been identified, their corresponding dichromatic planes can be intersected to yield the illuminant color. However, it is not always easy to determine which pixels from an arbitrary region of an image belong to which dichromatic plane. The image region may cover an area of the scene encompassing several different materials and hence pixels from several different dichromatic planes. The new multilinear constraint accounts for this multiplicity of materials and provides a mechanism for choosing the most plausible illuminant from a finite set of candidate illuminants. The performance of this new method is tested on a database of real images.

1 Introduction

Image colors vary significantly with changes in the color of the light incident upon a scene. Being able to estimate accurately the scene illuminant color from an image is at the heart of the color constancy problem. Once the color of the light is known, it is easy to adjust the image colors accordingly [17]. The two main approaches to illumination estimation can be classified roughly as statistics-based versus physics-based [6]. Statistics-based methods [10] [7] [3] have been quite successful; however, they require a relatively large number of differently colored surfaces to be present. They also make no use of the underlying physics of image formation. Physics-based techniques derive constraints from physical principles so that the illuminant can be obtained as a solution to a set of equations.

Of the physics-based constraints explored so far such as interreflections, shadows, chromatic aberration [8] [4] [9] and specularities [6] [12] [14], specularities have proven the most useful. Specularities constrain the illuminant according to the dichromatic reflection model [15]. It states that the colors reflected by an inhomogeneous dielectric material will be a linear combination of two characteristic colors; namely, the color of

BMVC 2005 doi:10.5244/C.19.54

the specular component reflected from the air-surface interface and the color reflected from the body of the material. If neutral interface reflection is further assumed [13], as it is customary, then the chromaticity of the specular reflection is the same as that of the illuminating light.

Under the dichromatic model, the image colors observed from a surface patch of a single material must lie on a plane in color space. Given two different surface patches illuminated by the same light, the color of the light can be estimated as the intersection of the two dichromatic planes. Although this method has been shown to work [5] [6] [14], it requires *a priori* knowledge as to which image colors emanate from a single material.

The approach proposed here also exploits the dichromatic model constraint, but overcomes the requirement that image colors be pre-grouped in terms of scene materials. This is achieved by describing the image colors in terms of a multilinear model consisting of several planes simultaneously oriented around an axis defined by the color of the illuminant. From a set of candidate colors, the color of the light is estimated by assessing how well each of the candidates explains the observed color response under the assumption that the observed response can be described by a certain fixed number of coexisting linear models. The resulting multilinear constraint is reformulated into a set of simultaneous linear equations using a Veronese projection. The smallest eigenvalue of the resulting matrix provides a quantitative measure of how well a candidate illuminant explains the image data. The candidate exhibiting the smallest eigenvalue is then chosen as representative of the scene illuminant.

2 The Dichromatic Reflection Model

The dichromatic reflection models [15] asserts that the colors displayed by an inhomogeneous dielectric material live in a two-dimensional subspace of the color space. This two-dimensional subspace is ordinarily referred to as the dichromatic plane of the material. Under the assumption of neutral interface reflection [13], this subspace is the span of two characteristic colors: the color of the illuminating light and a color that depends on the reflection properties of the observed material. Note then that if different materials are illuminated by the same light, the colors that these materials display lie on a finite collection of two-dimensional subspaces and that the intersection of this collection is a one-dimensional subspace which is the span of the color of the light.

3 A Constraint on the Observed Colors

The dichromatic reflection model constrains how the observed colors must distribute in color spaces. As mentioned in the previous section, they must organize into a particular collection of two-dimensional subspaces. In our approach, instead of looking at the constraint in color space, a reduced space is used to develop a multilinear constraint on the observed colors. The proposed constraint is then employed for the estimation of the light.

Given the RGB image of a scene that complies with the dichromatic reflection model, suppose that the color of the illuminating light is known. Then, the projection of the observed colors onto the two-dimensional subspace that is orthogonal to the color of the light is a collection of one-dimensional subspaces of the aforementioned two-dimensional subspace. Note that if the number of different materials observed in the scene is equal to n, the number of one-dimensional subspaces is also equal to n. Also note that if the observed colors are projected onto a subspace orthogonal to a light whose chromaticity is different from that of the actual light, then the projected colors do not necessarily reside within a collection of n one-dimensional subspaces. This observation is core to the approach proposed here for the estimation of the light.

Developing on this observation, assume that $\mathbf{m}_1$ and $\mathbf{m}_2$ are some orthonormal basis of the two-dimensional subspace that is orthogonal to the color of the light (which is denote here by vector $\mathbf{w}$). The projection of any color $\mathbf{c}(\mathbf{x})$ into the subspace defined by vectors $\mathbf{m}_1$ and $\mathbf{m}_2$ can then be written as:

$$\mathbf{d}(\mathbf{x};\mathbf{w}) = M(\mathbf{w})\mathbf{c}(\mathbf{x}), \tag{1}$$

where $M(\mathbf{w})$ is a 2×3 matrix derived from the orthonormal basis as:

$$M(\mathbf{w}) = \begin{bmatrix} \mathbf{m}_1^T \\ \mathbf{m}_2^T \end{bmatrix}. \tag{2}$$

Suppose that $\mathbf{w}$ is the color of the actual illuminating light. Then, since each projected color $\mathbf{d}(\mathbf{x};\mathbf{w})$ must lie on one of the n one-dimensional subspaces, there must be a collection of two-dimensional vectors $\{\mathbf{u}_i\}_{i=1}^n$ such such that $\mathbf{d}(\mathbf{x};\mathbf{w})^T\mathbf{u}_i = 0$ for some $i \in \{1,\ldots,n\}$. Therefore any projected color also satisfies:

$$\prod_{i=1}^{n} \mathbf{d}(\mathbf{x};\mathbf{w})^T\mathbf{u}_i = 0. \tag{3}$$

This is the multilinear constraint on the observed colors that we use for the estimation of the light. Since the constraint is developed on a subspace that is orthogonal to the light (vector $\mathbf{w}$), note that the constraint is also satisfied for any other color of the light written as $k\mathbf{w}$, $k \neq 0$. Consequently, from this equation it is only possible to estimate the chromaticity of the illuminating light, not its intensity. Also note that if there are at least two materials whose color responses differ sufficiently that they properly fill their corresponding subspaces (i.e., not a degenerate one-dimensional subspace), then the color of the light satisfying (3) is unique up to a multiplicative factor. Clearly, if the observed colors are sufficiently varied, then the projections of the colors onto any other two-dimensional subspace (orthogonal to any other arbitrary light that is not in the span of the actual illuminant) will not reside within a collection of n one-dimensional subspaces, but within a larger collection of such subspaces.

4 Estimating the Color of the Light

To estimate the color of the illuminating light, we assume that there is a finite set of candidate lights to choose from. We then check how well each of these candidates explains the observed color response.

Consider the image of a scene satisfying the dichromatic reflection model under the assumption of neutral interface reflection. Assume that there are n different materials in the imaged scene. Then, from Equation (3), we can see that a given candidate light is the best explanation of the observed colors if there is a set of n vectors $\{\mathbf{u}_i\}_{i=1}^n$ such that Equation (3) holds. This seems to require that in order to judge whether or not a candidate light is the actual illuminant, we must first estimate the **u**'s. However, as will be shown next, this turns out not to be necessary.

Let us begin by considering how the **u**'s may be estimated. Clearly, estimation of the **u**'s from Equation (3) is a nonlinear problem. Nevertheless, we note that by recasting Equation (3), this nonlinear problem can be partially rendered into a linear one. Multiplying out (3), it is not difficult to verify that the equation can be written as the sum of $n+1$ terms of the form $d_1^{n_1} d_2^{n_2}$ $(n_1 + n_2 = n)$ with associated coefficients depending on $\{\mathbf{u}_i\}_{i=1}^n$. The totality of these terms can be represented using the Veronese map of degree n on two variables ($\mathbf{d} = [d_1\ d_2]^T$), which is written as:

$$\nu_n(\mathbf{d}) = [d_1^n \quad d_1^{n-1} d_2^1 \quad d_1^{n-2} d_2^2 \quad \dots \quad d_2^n]^T. \tag{4}$$

Denote the overall coefficient associated with the term $d_1^{n_1} d_2^{n_2}$ as a_{n_1,n_2}. Then, Equation (3) can be recast as:

$$\nu_n(\mathbf{d}(\mathbf{x};\mathbf{w}))^T \mathbf{a}_n = \sum a_{n_1,n_2} d_1(\mathbf{x};\mathbf{w})^{n_1} d_2(\mathbf{x};\mathbf{w})^{n_2} = 0, \tag{5}$$

which is a linear expression on the coefficients $\mathbf{a}_n$. Note that if we can solve for the coefficients $\mathbf{a}_n$, then from $\mathbf{a}_n$ we can solve for the **u**'s by using a technique for polynomial factorization. The coefficients $\mathbf{a}_n$ can be solved for by constructing a system of linear equations using colors from different image locations $\{\mathbf{x}_i\}_{i=1}^m$, $m \geqslant n+1$. Stacking (5) in matrix form, the resulting system can be written as:

$$\Lambda_n \mathbf{a}_n = \begin{bmatrix} \nu_n(\mathbf{d}(\mathbf{x}_1;\mathbf{w}))^T \\ \nu_n(\mathbf{d}(\mathbf{x}_2;\mathbf{w}))^T \\ \vdots \\ \nu_n(\mathbf{d}(\mathbf{x}_m;\mathbf{w}))^T \end{bmatrix} \mathbf{a}_n = \mathbf{0}. \tag{6}$$

Vidal *et al.* [16] proved that the rank of a matrix with the same structure as that of the $m \times (n+1)$ matrix Λ_n is equal to n if among the considered colors (using the terminology of our particular problem), there are colors from the n different materials observed in the scene (which is ensured if all observed colors are considered). From this result, it immediately follows that a candidate color **w** is the actual color of the light if the smallest eigenvalue of matrix $\Lambda_n^T \Lambda_n$ is equal to zero. Indeed, in the least-square sense, the problem of finding a coefficient $\mathbf{a}_n$ satisfying (6) can be written as $\min_{\mathbf{a}_n} \|\Lambda_n \mathbf{a}_n\|^2 = \min_{\mathbf{a}_n} \mathbf{a}_n^T \Lambda_n^T \Lambda_n \mathbf{a}_n$. Since matrix Λ_n has rank n, the null space of the $(n+1) \times (n+1)$ matrix $\Lambda_n^T \Lambda_n$ has dimension equal to 1.

We see then that it is possible to assess whether or not a candidate light is the actual illuminating light simply by inspecting the smallest eigenvalue of $\Lambda_n^T \Lambda_n$. There is no need to solve for the **u**'s. In practice, due to noise or nonconformance of the observed colors to the assumed color image formation model, this eigenvalue does not vanish. Nevertheless, from the given set of candidate lights, the light whose smallest eigenvalue is minimum across the set of smallest eigenvalues is chosen as the illuminating light.

Vidal *et al.* [16] also showed that the number of different models (the number of different materials in the scene) can be estimated from a rank constraint on matrix Λ_n. Here, we simply assume that a fixed number of materials do coexist. This assumption is made relying on the observation that, in principle, the constraint of Equation (3) remains valid even if a number of different models larger than the actual value is used to formulate the constraint [16].

5 Experimental Results

To assess the performance of the proposed approach, we carried out experiments on the database produced at the Computational Vision Lab, Simon Fraser University [2]. This database is populated with images that do not strictly observe the assumptions on color image formation laid down by the dichromatic reflection model. Nevertheless, by such an assessment we seek to determine how the proposed algorithm performs in unconstrained natural scenes as well as to establish a comparison with the main statistics-based techniques.

The SFU database contains images of 32 scenes under 11 different illuminants. These images are divided into four categories: (1) *mondrian*, a set of images with minimal specularities; (2) *specular*, images with non-negligible dielectric specularities; (3) *metallic*, images with metallic specularities; and (4) *fluorescent*, images of scenes with at least one fluorescent surface.

As a strategy to maintain a moderate computational cost, we split any given image into non-overlapping regions so that we can then use a constraint with a relatively small number of models on each of these regions. We assume that in a smaller region the number of coexisting materials is also likely to be smaller. Next, from each block, we obtain a measure of how well each candidate light explains the observed colors. The average of these measures across all blocks is then used as an overall assessment of each light. In our experiments, images are split into blocks of 200×200 pixels, and it is assumed that in any of these blocks up to 4 different materials coexist. The 11 lights used to construct the database are taken as the candidate lights.

Tables 1 and 2 show the performance of the proposed approach. Table 1 shows the chromaticity error measured as the Euclidean distance between the chromaticity of the actual and the estimated light. In Table 2 the performance is shown as the angular difference in degrees between the estimated and the actual values of the light. In both tables, results are reported by the mean, the median and their corresponding 95% confidence intervals (c.i.). The mean has been used to report the performance of major statistics-based techniques [1]. More recently, the median has been advocated as a more appropriate measure of the central tendency of measured errors [11]. The confidence interval gives some idea about how uncertain we are about the measured statistics.

Overall, the errors reported in Tables 1 and 2 are higher than those produced by the best performing algorithms reported in [1]. Nevertheless, the approach shows a better or comparable performance to that of commonly used algorithms such as the gray world method. Note that there is a noticeable difference between the mean and median values. This is an indication that there are a few images with high errors. We inspected the images with a chromaticity error higher than 0.2. This is an ensemble of 31 images out of the 518 comprising the database. To test the validitiy of the assumption that each image block contains 4 materials, we re-estimated the illuminant on this ensemble assuming 8 models. We then compared the errors between the two set of results. To this end, for a given image of the ensemble, we measured the difference between the errors of the newly and the previously estimated light. If this difference is negative, the newly estimated light is better. Out of the 31 images in the ensemble, 13 had a negative difference, with a mean of -0.096; 7 images, a positive difference, with a mean of 0.035. We also re-estimated the light using a partition of 100×100 non-overlapping blocks, and 4 models. In this test, 10 images had a negative difference, with a mean of -0.172; 2 images, a positive difference, with a mean of 0.032. These tests suggest that the overall accuracy might be improved by a better estimate of the number of materials appearing in each region. Underestimating the number causes the multilinear constraint to be violated; overestimating requires more computation.

	mean	95% c. i.	median	95% c. i.
dataset	**0.070**	**0.063 - 0.076**	**0.050**	**0.044 - 0.058**
mondrian	0.059	0.050 - 0.068	0.037	0.026 - 0.044
specular	0.057	0.046 - 0.069	0.046	0.037 - 0.052
metallic	0.097	0.085 - 0.11	0.092	0.075 - 0.104
fluorescent	0.059	0.045 - 0.074	0.044	0.037 - 0.058

Table 1: Errors in rg-chromaticity. Errors obtained using 4 planar models to explain the colors of 200×200 blocks. As a measure of the tendency of errors, the mean, median and corresponding confidence intervals (c.i.) are given.

	mean	95% c. i.	median	95% c. i.
dataset	**9.64**	**8.79 - 10.50**	**7.77**	**6.34 - 7.91**
mondrian	8.17	6.88 - 9.42	4.30	3.76 - 6.34
specular	7.82	6.37 - 9.30	7.71	4.29 - 7.89
metallic	13.63	12.01 - 15.31	11.92	9.67 - 14.32
fluorescent	7.87	6.13 - 9.84	6.34	5.54 - 7.91

Table 2: Angular errors in degrees. The estimate of the light is obtained by averaging the assessment of each candidate over non-overlapping 200×200 blocks of the images under the assumption that in each block up to 4 different materials coexist.

6 Concluding Remarks

In this paper, an approach for estimating the color of the illuminating light based on the dichromatic reflection model has been proposed. We have shown that through a multilinear constraint on the observed colors it is possible to evaluate how well these colors are explained by a given candidate light. This evaluation is carried out by inspecting the smallest eigenvalue of a matrix that is derived from the observed colors, the color of the candidate light, and an assumption as to the number of different materials expected to be present. Experiments on the SFU database show promising results. The approach eliminates the need for grouping image colors in terms of materials that is usually required for methods based on intersecting dichromatic planes.

References

[1] K. Barnard, L. Martin, A. Coath, and B. Funt. A comparison of computational color constancy algorithms – part II: Experiments with images. *IEEE Trans. Image Processing*, 2002.

[2] K. Barnard, L. Martin, B. Funt, and A. Coath. A data set for color research. *Color Res. Appl.*, 27(3):148–152, 2002. available: `http://www.cs.sfu.ca/~colour/data/colour_constancy_test_images/index.html`.

[3] V. Cardei, B. Funt, and K. Barnard. Estimating the scene illumination chromaticity using a neural network. *J. Opt. Soc. Am. A*, 19(12):2363–2373, 2002.

[4] G. Finlayson and B. Funt. Colour constancy using shadows. *Perception*, 23:89–90, 1994.

[5] G. Finlayson and G. Schaefer. Convex and non-convex illuminant constraint for dichromatic color constancy. In *Int. Conf. Comput. Vision Pattern Recogn.*, volume I, pages 598–604, 2001.

[6] G. Finlayson and G. Schaefer. Solving for colour constancy using a constrained dichromatic reflection model. *Int. J. Comput. Vision*, 42(3):127–144, 2001.

[7] D. Forsyth. A novel algorithm for color constancy. *Int. J. Comput. Vision*, 5(1):5–36, 1990.

[8] B. Funt, M. Drew, and J. Ho. Color constancy from mutual reflection. *Int. J. Comput. Vision*, 6(1):5–24, 1991.

[9] B. Funt and J. Ho. Color from black and white. *Int. J. Comput. Vision*, 3(2):109–118, 1989.

[10] G. Finlayson S. Hordley and P. Hubel. Color by correlation: A simple, unifying framework for color constancy. *IEEE Trans. Pattern Anal. Machine Intell.*, 23(11):1209–1221, 2001.

[11] S. Hordley and G. Finlayson. Re-evaluating colour constancy algorithms. In *Int. Conf. Patt. Recogn.*, volume 1, pages 76–79, 2004.

[12] H.-C. Lee. Illuminant color from shading. In *Proc. SPIE: Perceiving, Measuring, and Using Color*, volume 1250, pages 236–244, 1990.

[13] 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.

[14] G. Schaefer. Robust dichromatic colour constancy. In *Int. Conf. Image Analysis Recogn.*, volume 2, pages 257–264, 2004.

[15] S. Shafer. Using color to separate reflection components. *Color Res. Appl.*, 10(4):210–218, 1985.

[16] R. Vidal, Y. Ma, and S. Sastry. Generalized principal component analysis (GPCA). In *Int. Conf. Comput. Vision Pattern Recogn.*, volume 1, pages 621–628, 2003.

[17] J. von Kries. Chromatic adaptation. *[originally published in Festschrift der Albrecht-Ludwigs-Universitat (Fribourg, Germany, 1902), pp. 145-148.]*. In Sources of Color Vision, L. D. MacAdam, ed. (MIT Press, Cambridge, Mass., 1970), pages 109-126.