 Copyright 2000, IS&T

Illumination Detection in Linear Space

Bo Tao — Streaming21 Inc., Los Gatos, CA
Ingeborg Tastl — Sony USRL, San Jose, CA
Naoya Katoh — Sony PNC, Tokyo, Japan

Abstract

Scene illumination in digital photography is of limited types, including Daylights of various temperatures, Horizon, Incandescent and various Fluorescents. For white balancing purposes in consumer digital photography, it is often sufficient to compute the scene illumination with limited accuracy. Therefore, rather than estimating the tristimulus color value or the illumination spectrum as a continuous (vector) variable, we can detect the illumination via hypothesis testing and selecting from a finite set of candidate illuminations. An illumination detection approach is introduced in this work by incorporating linear decomposition of the spectral surface reflectance.

Illumination estimation, illumination detection, hypothesis testing, linear decomposition, white balancing.

Introduction

Illumination estimation is an important issue in both color science and digital photography applications. A number of illumination estimation techniques have been proposed in the past, while many new ones are being introduced at each conference and by each new camera model, which shows the importance of this research area.

There are basically two groups of techniques. The first group of techniques treats the color or the spectrum of the unknown illumination as a continuous vector variable and uses signal estimation approach.[1] Examples include the widely used gray-world algorithm,[2] estimation using specular reflectance,[3] the neural network approach[4,5] and the linear space approach,[6] among many others.

The other group of techniques is a relatively new development. Rather than estimating a continuous variable, this type of techniques determines the illumination type by selecting from a finite set of candidate illuminations. To differentiate these two groups of techniques, we call this second group **illumination detection**, drawing the analogy between signal estimation and signal detection in communications theory.[1] Illumination detection is a promising approach because in the real world, the scene illumination is only of limited types, such as various daylights, horizon, incandescent and several fluorescents. Furthermore, for many applications including digital photography, it is acceptable to have an estimation of finite accuracy. Color by correlation[7] is an excellent example of the illumination detection approach.

Illumination estimation and detection are closely related to each other. For example, the binary histogram of the colors present in an image is depends on the scene illumination and thus can be used to compute the scene illumination. Both illumination estimation[4] and illumination detection[7] approaches have been introduced based on the same statistic. In practice, one can compute the illumination using either the estimation or the detection approach based on consideration of accuracy and execution speed.

In this work, we will introduce a new illumination detection technique using linear projection of the spectral surface reflectance. It is well known that the reflectance of most surfaces and the spectral power distribution of many illuminations can be approximated very well in a low-dimensional linear space.[8] Compared to other techniques, linear decomposition of reflectance and illumination takes into account the physical properties of the surface and the illumination and simplifies the illumination estimation (or detection) problem by providing a compact representation. A number of techniques have thus been developed using linear decomposition.[6,9–12] However, they all can be seen as taking an estimation approach. In this paper, we will put the linear projection technique into an illumination detection framework.

There are several potential advantages to combine the detection approach with the linear projection technique. The first is that, in the detection framework using hypothesis testing, the spectral power distribution of the candidate illumination is assumed to be known for each hypothesis (or equivalently, for each candidate illumination). Therefore, a lot more information is available to solve for the surface reflectance. Another advantage is that, by knowing the spectral power distribution of the candidate illumination, we can avoid the problem of linearly decomposing the illumination spectrum. As we have shown before,[12] although a 3D linear space is usually good enough for surface reflectance, as much as 6 dimensions are often needed to approximate the illumination spectrum. For a typical RGB image, we only have 3 known color values at each pixel, the same number as the 3 unknown surface reflectance coefficients. Therefore, by avoiding decomposing the illumination spectrum, the ambiguity of solving for an underdetermined problem is avoided.

In the next section, we are going to present our new illumination detection technique by breaking it into several parts. We first discuss projecting the surface reflectance onto a linear space. Then we discuss how linear projection can be used in an illumination detection framework. One issue for linear projection-based techniques is that the computation is often very high. We avoid this pitfall by

Copyright 2000, IS&T

selecting representative pixels from the image. This greatly reduces the computational load. Finally, we give some preliminary experimental results to conclude this paper.

Illumination Detection in Linear Space

Linear Decomposition of Spectral Surface Reflectance

Let us sample each spectrum at a discrete number of wavelengths within the visible range. We index these sample wavelengths using n=1, …, N. Let L represent the spectral power distribution of the illumination, R_k represent the surface reflectance at pixel k=1, …, K and S_j represent sensor sensitivity functions with j=1, …, J where J is the number of sensors. For example, for a 512x512 image, K=512 x 512=262,144, and for an RGB image, J=3. The reflectance at each pixel R_k is constrained to be between 0 and 1 by their own physical properties. The color value of sensor j at pixel k is represented by $f_{k,j}$. The value is given by the sum of the product of L, R_k, and S_j, i.e.

$$f_{k,j} = \eta \sum_{n=1}^{N} L(n) \times R_k(n) \times S_j(n) \tag{1}$$

where η is a scaling factor. For digital photography applications, η is a function of exposure level and other parameters and is often difficult to estimate. The goal of illumination computation is to compute from the pixel color values in the image the sensor response to the illumination itself, i.e.

$$l_j = \eta \sum_{n=1}^{N} L(n) \times S_j(n) \quad . \tag{2}$$

Or even ideally, the goal is to estimate L itself.

As discussed previously, the spectral reflectance function R_k can be well approximated by using a low-dimensional linear model. That is \

$$R_k(n) = \sum_{m=1}^{M} \beta_{k,m} B_m(n)$$

where $\{B_m\}$ is the set of basis functions for reflectance, $\{\beta_{k,m}\}$ are weighting coefficients and M is the dimension of the linear space. The best basis in the least square sense is the set of eigenvectors of the covariance matrix of the reflectance, which has the largest eigenvalues. In our implementation, the basis functions for reflectance are derived from a set of more than 20,500 samples containing Munsell patches as well as measured reflectance spectrum from two databases described in [13] and [14].

Hypothesis Testing Using Linear Decomposition

Hypothesis testing for signal detection is a widely used technique. In our algorithm, we take a straightforward approach by making each hypothesis corresponding to a candidate illumination type. Let D be the number of hypotheses. Thus, we have pre-selected D candidate illuminations, L^a , $d=1, \ldots, D$, and hypothesis H^d is that the unknown scene illumination L is equal to L^a , i.e.

$$H^d: \quad f_{k,j} = \eta \sum_{n=1}^{N} L^d(n) \times R_k(n) \times S_j(n). \tag{3}$$

The essence of signal detection is to compute a cost function C^d for each hypothesis. The one with the lowest cost is declared the winner (other variations are possible though). In the following, we show how to compute a very simple yet powerful cost function.

By using the reflectance basis functions mentioned above, we can rewrite eq. (3) as

$$\begin{aligned} H^d: \quad f_{k,j} &= \eta \sum_{m=1}^{M} \beta_{k,m} \sum_{n=1}^{N} L^d(n) \times B_m(n) \times S_j(n) \\ &= \eta \sum_{m=1}^{M} \beta_{k,m} A^d_{j,m} \end{aligned} \tag{4}$$

with

$$A^d_{j,m} = \sum_{n=1}^{N} L^d(n) \times B_m(n) \times S_j(n).$$

For RGB images, we have J=3; we also use M=3 basis functions for surface reflectance. Henceforth, at pixel k, we have

$$\vec{F_k} = \eta A^d \vec{\beta^d_k} \tag{5}$$

in which

$$\vec{F_k} = \begin{bmatrix} f_{k,1} \\ f_{k,2} \\ f_{k,3} \end{bmatrix} \text{ and } \vec{\beta^d_k} = \begin{bmatrix} \beta^d_{k,1} \\ \beta^d_{k,2} \\ \beta^d_{k,3} \end{bmatrix}$$

are 3x1 vectors, and A^d is a 3x3 matrix with entries $A_{j,m}{}^d$. Thus, we have

$$\vec{\beta^d_k} = \frac{1}{\eta} inv(A^d) \vec{F_k}. \tag{6}$$

If the surface reflectance lies exactly in the 3D linear space spanned by $\{B_1, B_2, B_3\}$, then when hypothesis d is true, we have

$$\vec{\beta^d_k} = \vec{\beta^*_k} \text{ where } \vec{\beta^*_k}$$

denote the actual decomposition coefficients. On the other hand, since most surface reflectance can be very well approximated by this linear space, we know that

$$\vec{\beta^d_k} \approx \vec{\beta^*_k}$$

 Copyright 2000, IS&T

if hypothesis d is true. Therefore, the estimated reflectance for pixel k,

$$R_k^d(n) = \sum_{m=1}^{3} \beta_{k,m}^d B_m(n) \qquad (7)$$

is close to the actual reflectance $R_k^*(n)$. In contrast, when the hypothesis is false, the estimated reflectance will be different from the actual value.

To calculate the cost function for hypothesis testing, we use the constraints associated with the surface reflectance. In theory, many constraints can be used together. However, for fast computation, we choose to use only one, which is that the reflectance is limited between 0 and 1. Unfortunately, since we don't know the scaling factor η, we can only use the condition that the reflectance should be positive. The cost function is then defined as

$$c^d = \sum_{k=1}^{K}\sum_{j=1}^{3} | f_{k,j} - \sum_n L^d(n) \times S_j(n) \times \max(\sum_{m=1}^{3} inv(A^d)\vec{F_k} B_m(n),\ 0) | \qquad (8)$$

Note that in eq. (8), the unknown value η and its inverse have canceled out each other.

The intuition behind our cost function is that for the true hypothesis, the estimated reflectance will be close to the true reflectance and thus mostly be positive. Therefore, the cost will be small. However, for false hypotheses, the reflectance will swing wildly and have negative values. By removing the negative values, the cost value will be large for these false hypotheses. It is interesting to note that

$$\max(\sum_{m=1}^{3} inv(A^d)\vec{F_k} B_m(n),\ 0)$$

can be seen as a sub-optimal solution to the constrained minimization problem:

$$\min_{R_k(n)} \sum_{j=1}^{3} || f_{k,j} - \sum_n L^d(n) S_j(n) R_k(n) ||$$

subject to the constraint $0 \le R_k(n)$, which is closely related to the problem solved in our previous work.[12]

Finally, the hypothesis with the lowest cost is declared the winner (winner-take-all) and L^d is used as the estimate of the scene illumination. Other variations to this strategy can be used too.

Computation Complexity and Pixel Selection

We give a count of multiplications needed for the algorithm. For each hypothesis in our algorithm, we need to compute the matrix A^d and its inverse. To compute A^d, 2N multiplications are needed. To invert a 3x3 matrix, several dozen multiplications are required. It should be noted that A^d and its inverse only have to be computed once for each hypothesis and thus the cost is negligible. Under each hypothesis, it costs 9 multiplications to calculate the decomposition coefficients

$$\beta_k^{d'},$$

and 3N multiplications to calculate the estimated reflectance. To compute the final cost function, only N more multiplications are needed, since $L^d(n)s_j(n)$ has been computed when we calculate the matrix A^d. Therefore, the number of multiplications for each hypothesis is approximately 4NK, and the number for the whole algorithm is 4NKD, where N is the number of wavelength samples, K is the number of pixels and D is the number of hypothesis.

Since the number of pixels used for calculation is linearly proportional to the computation cost, it is of great interests to devise efficient algorithms to select pixels to participate in the detection process. In our previous work[12] the set of possible color for various illuminations is computed from the set of Munsell chips and compared with each other. It is noticed that colors with large values are more indicative of the illumination type. Therefore, in our illumination detection algorithm, a pixel is selected only if it has a large value in at least one color channel. In our tests, we find that 1,000 pixels are often enough for 1~3 mega-pixel-size pictures. Overall, the algorithm is extremely efficient to compute.

Table 1. Description of test images.

Num	Description	Camera
1	Lightbooth; Figure 1(a); CIE A	Sony
2	Lightbooth; Figure 1(a); Horizon	Sony
3	Lightbooth; Figure 1(a); Coolwhite	Sony
4	Lightbooth; Figure 1(a); U30	Sony
5	Office; Figure 1(b); U30	Sony
6	Office; Scene not shown; U30	Nikon
7	Office; Scene not shown; U30	Nikon
8	Outdoor; Figure 1(c); Daylight	Nikon
9	Outdoor; Scene not shown; Daylight	Nikon

Experimental Results

We use the raw data output from two digital still cameras manufactured by Sony and Nikon, respectively, in our experiments. Their sensitivities are measured using a mono-chromater. Pictures are taken in light-booth, office and outdoor environments with examples shown in Figure 1. A description of these images is given in Table 1. The actual spectral power distribution of the scene illumination is measured at the same time when the images are captured. The tri-stimulus values of the true illumination can therefore be computed by multiplying the measured scene illumination with the measured camera sensitivities. We

Copyright 2000, IS&T

compare our algorithm with the gray-world algorithm and the so-called “modified gray-world” algorithm.[15]

In our algorithm, the number of wavelength samples is chosen to be 33 (N=33), 1,000 pixels are selected for each picture (K=1,000) and 18 hypotheses are used (D=18), including Daylights, Horizon, Incandescent and some typical Fluorescents.

In Figure 2, we compare our algorithm with the gray-world algorithm and the modified gray-world algorithm by using the chromaticity error. The error is computed by first normalizing the color value with respect to the sum of R, G and B, i.e. r=R/(R+G+B) and similarly for g. The error metric is computed as

$$\sqrt{(r^* - r)^2 + (g^* - g)^2}$$

where r^* and g^* are the actual illumination value as mentioned above. Note that the actual illumination is the same as one of the candidate illuminations for image 5, and our algorithm made correct detection resulting zero error.

Table 2. Comparison of minimum, mean and maximum errors for our algorithm, the gray-world algorithm and the modified gray-world algorithm.

	Minimum	Mean	Maximum
Gray-world	0.017	0.174	0.565
Modified gray-world	0.020	0.079	0.255
Our algorithm	0	0.057	0.125

The (minimum, mean, maximum) error triplets for the three algorithms are shown in Table 2. We note that our algorithm outperforms the other two algorithms on average. It especially excels when there are dominant colors in the scene such as images 1 through 4.

Finally, we look at the surface reflectance computation (eq. 7) during the illumination detection process. We use image 2 as an example. The actual illumination is Horizon. Shown in Figure 3 are the computed surface reflectances for 8 pixels selected for detection. The illumination hypotheses shown are Horizon and D40. It is seen that the computed surface reflectance has more negative values for D40 hypothesis than for Horizon hypothesis. This confirms that our premise for the new illumination detection algorithm is correct, i.e. the true hypothesis illumination will have less negative components in the computed surface reflectance. Thus, by applying the surface reflectance positivity constraint and using eq. 8, we are able to detect the correct illumination.

(a)

(b)

(c)

Figure 1. Examples of raw image data used in experiments.

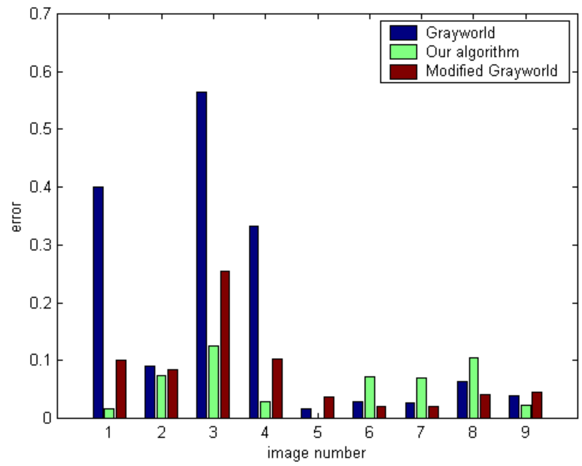

Figure 2. Comparison between the new algorithm, the gray-world algorithm and the modified gray-world algorithm. The error metric is

$$\sqrt{(r^* - r)^2 + (g^* - g)^2}\ .$$

Acknowledgments

The authors thank Ted Cooper, Naoya Katoh and Kenichi Nishio of Sony Corporation for providing some of the surface reflectance database and test images. They also thank them for many insightful discussions.

References

1. H. V. Poor. An introduction to signal detection and estimation, 2nd edition. Springer-Verlag, 1994.
2. B. Wandell. Foundations of Vision. Sinauer Associates, Inc. 1995.
3. H. Lee. Method for determining the color of a scene illuminant from a color image. U.S. Patent 4,685,071.
4. B. Funt, V. Cardei and K. Barnard. Learning color constancy. Proceedings of the Fourth IS&T/SID Color Imaging Conference, 1996.
5. V. Cardei, B. Funt, K. Barnard. White point estimation for uncalibrated images. Proceedings of the Seventh IS&T/SID Color Imaging Conference, 1999.
6. L. Maloney and B. Wandell. Color constancy: a method for recovering surface spectral reflectance. J. Opt. Soc. Am. A vol. 23, no. 1, 1986.
7. G. Finlayson, P. Hubel and S. Hordley. Color by correlation. Proceedings of the Fifth IS&T/SID Color Imaging Conference, 1997.
8. D. Marimont and B. Wandell. Linear models of surface and illuminant spectra. J. Opt. Soc. Am. A vol. 11, 1992.
9. H. Trusell and M. Vrhel. Estimation of illumination for color correction. Proc. Internat. Conf. Accous. Speech Signal Processing, 1991.
10. W. Freeman and H. D. Brainard. Bayesian decision theory, the maximum local mass estimate and color constancy. Proc. Fifth Internat. Conf. Computer Vision. 1995.
11. R. Lenz, P. Meer and M. Hauta-Kasari. Spectral-based illumination estimation and color correction. Color Research and Application, vol. 24, no. 2, 1999.
12. B. Tao and I. Tastl. Illumination estimation using linear decomposition and constrained optimization. Proceedings SPIE Color Imaging: Device-Independent Color, Color Hardcopy and Graphics, San Jose, Jan. 2000.
13. Johji Tajima. Development and analysis of standard object colour spectra database (SOCS). Proceedings of the Seventh IS&T/SID Color Imaging Conference, 1999.
14. H. Trusell. Reflectance database available at http://www.ece.ncsu.edu/people/faculty/bios/hjt.html.
15. G. Finlayson, S. Hordley and P. Hubel. Unifying color constancy. Proceedings the Seventh Color Imaging Conference, 1999.

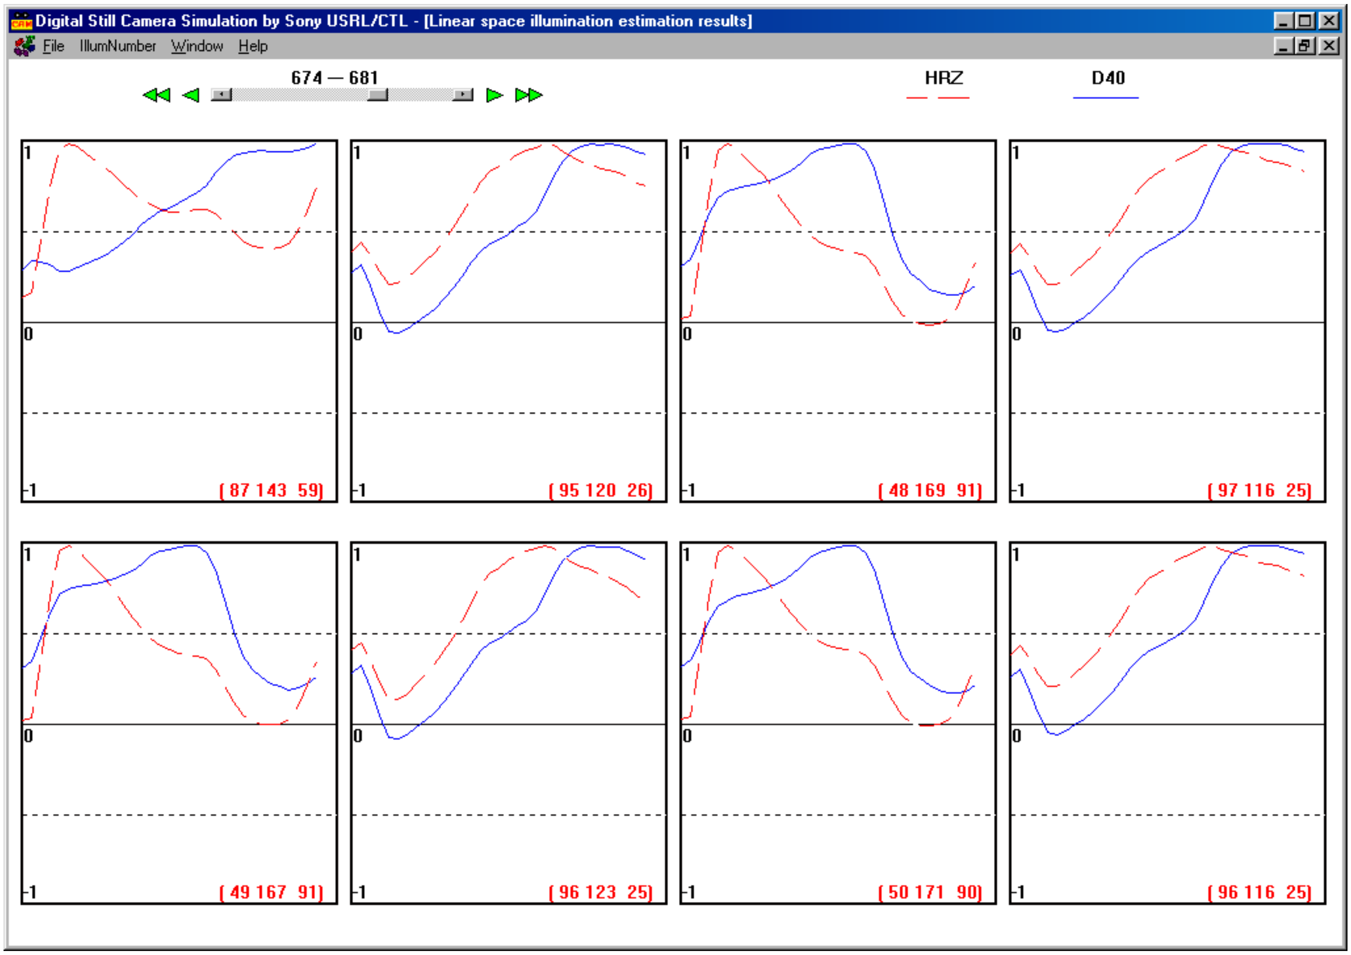

Figure 3. Surface reflectance computation for different illumination hypotheses. The actual illumination is Horizon.