Practical Scene Illuminant Estimation via Flash/No-Flash Pairs

Cheng Lu and Mark S. Drew; School of Computing Science, Simon Fraser University, Vancouver, British Columbia, Canada V5A 1S6
{*clu,mark*} *@cs.sfu.ca*

Abstract

In this paper, we present a method to estimate ambient illuminants using no-flash/flash image pairs. Accurate estimation of the ambient illuminant is useful for imaging applications. In most applications, however, this task is difficult because of the complicated combination of illuminants, surfaces, and camera characteristics during the imaging process. To estimate the scene illumination, a version of the "illuminating illumination" method suggested by Dicarlo et al. is used. The method introduces camera flash light into the scene, and the reflected light is used to estimate the ambient illuminant. The original method needs an extra step of estimating the object surface reflectance, using a 3-dimensional linear surface model and the knowledge of the spectral responsivities of camera sensors. Here we consider the problem of estimating the ambient illuminant directly, with only flash/no-flash pairs, without information on surface reflectance and camera sensors. First, the flash image is registered with the no-flash image: the difference between the two gives a pure-flash image, as if it were taken under flash only. The no-flash and pure-flash images are represented by a physically-based model of image formation which uses assumptions of Lambertian surfaces, Planckian lights, and narrowband camera sensors. We argue that first going to a "spectrally sharpened" color space, and then projecting the difference in a log domain of the pure-flash image and the no-flash image into a geometric-mean chromaticity space, gives the chromaticity of the ambient illuminant. We verify that the chromaticities corresponding to illuminants with different temperatures fall along a line *on a plane in the log geometric-mean chromaticity space. Simply by taking the nearest color temperature along this illuminant line, or classifying into one of potential illuminants, our algorithm arrives at an estimate of the illuminant.*

Remarkably, our algorithm is truly practical as it can estimate the color of the ambient light even without any prior knowledge about surface reflectance, flash light, or camera sensors. Experiments on real images demonstrate that estimation accuracy can be very good.

Introduction

Estimating the scene illumination from image data is important in many applications, including photography, color imaging and printing. Many algorithms for this problem have been suggested. Most can be described as color constancy algorithms, designed to disambiguate surface and illuminant components in images. The estimated illuminant can be either in the form of a full spectral power distribution (SPD), or classified to be one of likely illuminant types. Because of the small number of color sensor responses, estimating spectral distribution of illuminants is an underdetermined problem, usually needs physical constraints, and can lead to low estimation accuracy. On the other hand, simply classifying the unknown illuminant to be one of several potential illuminants, or defining the correlated color temperature, related to blackbody radiators, simplifies data processing, stabilizes computation, and is useful in many applications, including white balance for photography [1, 2].

Several methods have been developed using various physical or statistical models. The color-by-correlation method, suggested by Finlayson et al. [1], employs a statistical model to estimate the illuminant of a given image by assigning the most likely illuminant. This is accomplished by using a chromaticity gamut representation for the distribution of surfaces under different illuminants, and associating image chromaticities with the most likely reference gamut derived for each of several illuminants. Based on this work, Tominaga and Wandell asserted an improved estimation by using a scaled version of the red and blue sensor responses [2]. For single-surface color constancy [3], a physical dichromatic model of reflectance has been used. This model incorporates the body and highlights reflection of a single surface and predicts chromaticities of single surfaces to fall along a line. The intersecting point of this chromaticity line with the Planckian locus gives an estimate of the illuminant.

Dicarlo et al. use the camera flash to obtain an additional image for estimating the scene illuminant [4]. The flash/no-flash images are combined to produce a pure-flash image for the scene. This pure-flash image together with knowledge of the SPD of the flash is then used to estimate the surface reflectance in the scene. Finally, using the surface reflectance and the no-flash image, the most likely ambient illuminant can be determined. This approach provides a practical way to estimate ambient illuminant, though it has some limitations when applied to real world applications: it requires knowledge of the camera sensor spectral characteristics and flash SPD. Also, estimation of scene surface reflectance is required for estimating the illuminant, using a linear model approximation of the surface reflectance. The dimensionality of the linear model must be chosen to match the number of camera sensors, so a 3-dimensional model is used, for a standard camera, which is usually not sufficient for representing surfaces in the natural world. In this paper, we introduce a method that again uses the flash/no-flash image pair to estimate the ambient illuminant, but in a different way from the above that allows us to eliminate all the above requirements.

Here we examine flash/noflash still image pairs. The light impinging on a surface point is of course quite different in an image taken under ambient lighting and under a combination of both ambient plus a flash. For clarity, let us refer to the first image as "Ambient" and the second as "Both" (A and B). If we control the camera settings, or at least know them, then $(B-A)$ should yield an image as if it were taken under the flash only (assuming one adjusts overall pixel magnitudes to compensate for camera settings, as in [5, 6]). This is due to the fact that the B image consists of reflected light from the ambient sources plus from the flash. We denote this pure-flash image as "Flash" (F).

Since the unknown ambient illuminant contributes to both "A" and "B", it is hard to estimate the ambient illuminant by directly using the two images. On the other hand, the pure-flash image "F" sees reflected light from only the flash illumination, which is fixed for a given camera. So we should be able to combine F with A, to be able to estimate the ambient light. Here, we show that a simplified image formation model can greatly aid in estimating the ambient illuminant in A using F as a reference illuminant image, without any prior knowledge of camera sensor, surface reflectance, and flash. We go over to a log color space in which the log-difference between "A" and "F" under different ambient illuminants falls along a line within a geometric-mean chromaticity plane. This line coincides with the Planckian locus (the Planckian locus has a linear behavior in the log chromaticity space). By associating the chromaticity of the difference image to the nearest color temperature along the Planckian locus, we arrive at our estimate of the ambient illuminant. Once the ambient illuminant is recovered, we also carry out a simple white balance, using as a reference white a pre-determined white patch under the ambient illuminant, removing the effects of any automatic camera white balance procedure.

Image Formation

We employ a simple image formation model which assumes Planckian lighting, Lambertian surfaces, and a narrowband camera. At a Lambertian surface point, under orthography, lighting is added up into a single effective light, taking into account visibility factors for each source. Let us recapitulate how linear behavior with lighting change results from this image model: Consider the RGB color R formed at a pixel, for illumination with spectral power distribution $E(\lambda)$ impinging on a surface with surface spectral reflectance function $S(\lambda)$. If the three camera sensor sensitivity functions form a set $Q(\lambda)$, then we have

$$R_k = \sigma \int E(\lambda)S(\lambda)Q_k(\lambda)d\lambda \ , \ k = R,G,B \ , \qquad (1)$$

where σ is Lambertian shading — surface normal dotted into illumination direction — along with visibility.

We wish to go to a model that explains the change in images formed under different lights by a simple diagonal 3×3 matrix. It has been found that this illuminant-change model holds, greatly simplifying the image-formation description, if we make the above assumptions [7]. In this case, we shall find that a log-difference image for flash F and ambient A images obeys a very simple form.

If the camera sensor $Q_k(\lambda)$ is exactly a Dirac delta function $Q_k(\lambda) = q_k\delta(\lambda - \lambda_k)$, then eq. (1) becomes

$$R_k = \sigma E(\lambda_k)S(\lambda_k)q_k \ . \qquad (2)$$

Now suppose lighting can be approximated by Planck's law, in Wien's approximation [8]:

$$E(\lambda, T) \simeq I k_1 \lambda^{-5} e^{-\frac{k_2}{T\lambda}} \ , \qquad (3)$$

with constants k_1 and k_2. Temperature T characterizes the lighting color and I gives the overall light intensity.

In this approximation, from (2) the RGB color R_k, $k = 1 \ldots 3$, is simply given by

$$R_k = \sigma I k_1 \lambda_k^{-5} e^{-\frac{k_2}{T\lambda_k}} S(\lambda_k)q_k \ . \qquad (4)$$

Let us define the following short-hand notations:

$$\begin{aligned} K &= \log(Ik_1\sigma); \quad s_k = \log(S(\lambda)); \\ w_k &= \log(k_1\lambda_k^{-5}q_k); \quad e_k = -k_2/\lambda_k \end{aligned} \qquad (5)$$

Taking logarithms, eq. (4) becomes

$$\log R_k(x) = w_k + K(x) + s_k(x) + (1/T(x))e_k \qquad (6)$$

Here, we have explicitly indicated dependence on 2D pixel location x: w_k is a characteristic 3-vector for the camera, as is e_k, and so does not depend on image location. However, the intensity and shading, encapsulated in K, do depend on location, as does the surface term s_k. Lighting color is dependent on the correlated color temperature T, which depends on what lighting the surface point sees and adds up. For the pair of images F and A, both K and T are different, e.g., in Figs. 6(b,c).

To eliminate the effect of scene geometry (intensity and shading term K), let us now go over to a chromaticity space c by dividing each channel by the geometric mean [9], $\sqrt[3]{R \times G \times B}$.

Then we define the geometric-mean chromaticity as

$$c_k = R_k / \sqrt[3]{\Pi_{i=1}^{3} R_i}, \quad \equiv R_k / R_M, \qquad (7)$$

and log version [9]

$$\begin{aligned} r_k &= \log(c_k) = \log(s_k/s_M) + (e_k - e_M)/T, \text{with} \\ s_M &= \sqrt[3]{\Pi_{j=1}^{3} s_j}\ , e_M = -k_2/3 \textstyle\sum_{j=1}^{p} \lambda_j, \end{aligned} \qquad (8)$$

Notice that the 3-vector direction $(e_k - e_M)$ is *independent of the surface* — it captures the illumination-change direction. That is, if we consider a single surface in the scene, for Planckian light (or lights such as Daylights which behave as if they were Planckian), as the illuminant temperature T changes, the log chromaticity color 3-vector moves along an approximately straight line which is independent of the magnitude and direction of the lighting.

To detect the color temperature of the illuminant in an image, i.e. locate the chromaticity of the illuminant along the straight line, we can remove the surface component $s_k(x)$ via a difference of *log* chromaticity 2-vectors. Here, we use the fact that the ambient image A and the pure-flash image F have the same surface reflectance at a pixel, so that if we simply subtract the *log* image F from the *log* image A , the surface component can be removed and only illuminant remains — we arrive at an estimate of the illuminant.

Estimating Ambient Illuminant via Ambient/Flash Image Difference

Now let us investigate how this simple image formation and log-chromaticity space can aid to estimate the ambient illuminant.

Spectral Sharpening

The simplified model (4) is more closely followed if $Q(\lambda)$ approximates a Dirac delta. We form an intermediate color space, in which the sensors are optimally combined so as to form new colors that better approximate color change induced by illuminant change via a diagonal model, using Spectral Sharpening [10]. This applies a 3×3 transformation matrix M to the sensors, or directly to colors, so as to better enforce a diagonal model.

Since we mean to take logs, we need nonnegative colors from the camera data (with zero values treated specially). To do so, we carry out a "spectral sharpening with positivity" transform (cf. [11, 12]). Using calibration targets under two different lights, we find M via a new optimization [14] consisting of a constrained form of "database sharpening" [10], but with hard constraints.

Log-difference Geometric-Mean Chromaticity

From (6), we notice that a difference image which is formed by subtracting F from A in log space removes both the camera term w_k as well as the surface term $s_k(x)$. Let us form a *ratio image* by a difference image D^{A-F} in log space: subtracting eq. (6) for two images, A and F, we have

$$\begin{aligned} D_k^{A-F}(x) &= \log R_k^A(x) - \log R_k^F(x) \\ &= [K^A(x) - K^F(x)] + [1/T^A(x) - 1/T^F(x)]e_k \end{aligned} \tag{9}$$

for the difference between log pixel values under light A and light F, at pixel indexed by x. Notice that the surface term is entirely removed, leaving a type of *intrinsic illumination* difference image which arises from: (i) the intensity difference (with shading/visibility), (ii) a term proportional to the camera-dependent lighting-change 3-vector e_k. We focus our attention on the illumination temperature only. To remove the intensity difference component, we go over to the geometric log chromaticity space according to eq.(8). Notice that eq.(9) is pixel-wise: it would be different for each pixel. Here, we assume that the scene contains a single ambient illuminant temperature (for one or several sources), so the log-difference chromaticity reduces to a simple form:

$$\begin{aligned} r_k^{A-F} &= D_k^{A-F} - 1/3\sum_{j=1}^{3} D_j^{A-F} \\ &= [1/T^A - 1/T^F](e_k - e_M); \end{aligned} \tag{10}$$

The above equation explicitly gives the log-difference vector as a function of illuminant temperature. For now, we carry all three components of chromaticity. We note that, in log space, r^{A-F} is orthogonal to $u = 1/\sqrt{3}(1,1,1)^T$. I.e., r lives on a plane orthogonal to u.

To characterize the 2D space, we can consider the projector $P_u^{\perp}$ onto the plane. $P_u^{\perp}$ has two non-zero eigenvalues, so its decomposition reads

$$P_u^{\perp} = I - u u^T = U^T U, \tag{11}$$

where U is a 2×3 orthogonal matrix. U rotates 3-vectors r into a coordinate system *in* the plane:

$$\chi \equiv U r, \quad \chi \text{ is } 2 \times 1. \tag{12}$$

Straight lines for illuminant changes for r are still straight in χ, that is, in the $\{\chi_1, \chi_2\}$ plane, we expect to see the log-difference images r^{A-F} with different ambient illuminations falling along a straight line through the origin. Recall that the vector $(e_k - e_M)$ is dependent on camera properties and captures the direction of changes of illumination. The scalar $[1/T^A - 1/T^F]$ locates a color temperature position on the line. Fig. 1 illustrates a log geometric-mean chromaticity diagram for the 1931 CIE color matching functions, where the dots aligned along a line are the log-difference geometric-mean chromaticity for a set of flash/no-flash pairs (and the triangle is χ for for color-matching functions). The image pairs were synthetically formed using 9 Planckian lights, from 2500K to 14500K with interval 1500K, Macbeth ColorChecker 24 surfaces, simple sensors with single impulse responsivities, and a xenon flash.

Note that this color temperature is *not* the correlated color temperature of the ambient illuminant, as it corresponds to the inverse-temperature difference between ambient and flash lights. We know that the flash illuminant is fixed for any camera, so that the temperature position on the line is fixed and can be a *reference temperature* for the ambient illuminant, i.e. the temperature

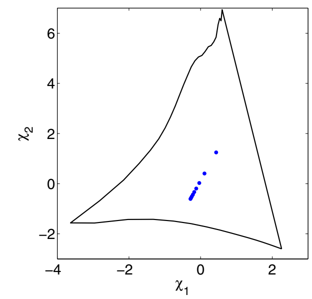

Figure 1. *CIE log geometric-mean chromaticity diagram: log-difference geometric-mean chromaticity of flash/no-flash pairs under 9 Planckian lights are shown with blue dots.*

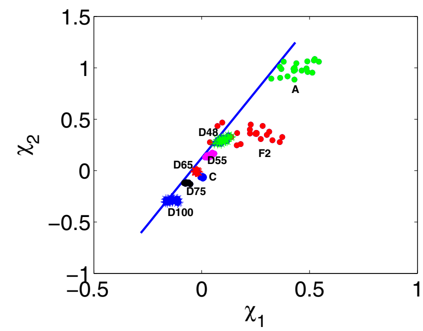

Figure 2. *SONY DXC930 camera: Log-difference geometric-mean chromaticity for Macbeth chart under 8 illuminants.*

uses the flash light as a reference. Here, we call this line the *reference illuminant temperature locus*. Thus, we can estimate the ambient light in an image by classifying it into one of a set of candidate reference color temperatures. We also carried out this log-difference chromaticity procedure for the Sony DXC930 digital camera, using flash/no-flash color patches created synthetically with Macbeth chart data under illuminants A, C, F2, and the Judd daylights. The result is shown in Fig. 2, along with the line formed by the Planckian data in Fig. 1. This camera has quite narrowband sensors. There are visibly 8 clusters, each of them corresponding to one of 8 illuminants. Among them, illuminants A, C, and the 5 daylights approximately align along the line of Planckian lights (the line is from Fig. 1); fluorescent illuminant F2 is off the Planckian lights line.

Estimating the Ambient Illumination

Fig. 3 illustrates the algorithm flow. First, in the sharpening phase, the constrained spectral sharpening process is performed and all camera responses RGB are transformed to a sharpened space. In the training step, using the sharpened camera the pairs of images of a reflectance database (e.g. the ColorChecker) under a set of sample illuminants are taken with flash turned on and turned off; the flash/no-flash pair is registered and the pure-flash image F is calculated; in log space, the difference log image $log(A) - log(F)$ projected to the $\{\chi_1, \chi_2\}$ plane in the geometric-mean chromaticity space, giving each illuminant a reference temperature along the reference temperature locus. In the estimating step for a new, test, image pair, we carry out the same process as in the training phase, then recover the temperature for the ambient light along the reference locus.

To assign an illuminant to the test image, we compute the er-

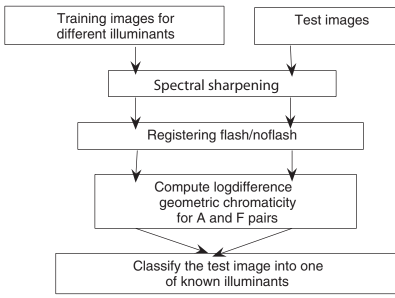

Figure 3. *Algorithm flow for estimating ambient illuminants*

ror between the the log-difference chromaticity of the test image and each illuminant cluster along the locus. Here, we use Euclidean distance between the mean of the cluster for each sample illuminant and the mean of the log-difference chromaticity of the test image as an error metric:

$$E_{ej} = \left(\sum_{i=1}^{2} (\chi_i^{mean(ej)} - \chi_i^{mean(et)})^2 \right)^{1/2} \tag{13}$$

where ej denotes the mean of the cluster for the jth illuminant, and et for the test image. Thus, the color temperature of the jth illuminant is chosen if it provides the minimum distance to the test image.

Experiments and Results

Spectral Sharpening

Our illuminant estimation algorithm is based on the assumption that camera sensors are quite narrowband. The spectral sharpening algorithm leads to a transform matrix M by which sensors can be optimally combined such that the image formation model more accurately applies. Then each T should give a smaller cluster in plot Fig. 2; i.e. after matrixing the RGB color values, the log-difference image pixels for each illuminant temperature are more separate from the image pixels under other illuminants.

For cameras like the Sony DXC930 whose sensors are quite narrowband, the clusters for different illuminants are considerably separated so that classifying amongst these illuminants can be accurately achieved. However, when we carry out the log-difference chromaticity procedure for images which are taken using Kodak DCS420, which has broader sensors, the image chromaticities for different illuminants are somewhat mixed; this would certainly lead to a failure in illuminant estimation. After matrixing the sensors with matrix M, the sensor curves are significantly narrower. Note that in our algorithm, the knowledge of camera sensors is not needed.

Again we used Macbeth surfaces under 8 illuminants to generate flash/no-flash pairs using Kodak DCS420 sensor curves. The log-difference chromaticity is shown in Fig. 4(a): the points for each illuminant are not separated enough. This is not surprising because broadband sensors make the RGB values more correlated. Fig. 4(b) plots the chromaticity after sharpening the RGB values. We see that the effect on separating illuminants is dramatic: the clusters for the 8 illuminants are much better separated.

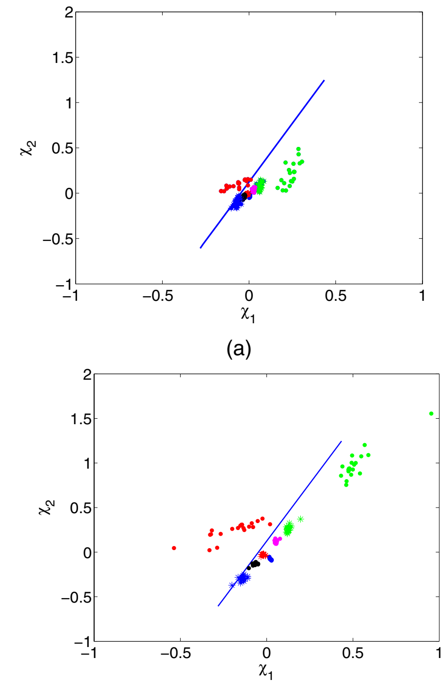

Figure 4. *Spectral sharpening (a): No sharpening, (b): With sharpening.*

Estimating Ambient Illuminants

In order to evaluate our algorithm, we first use synthetic image pairs using Sony DXC930 sensors (and Kodak DCS420 produced equally good results). In the training phase, we image 461 Munsell color patches with 102 measured light sources [13]. The means of the log-difference chromaticity of these images for each illuminant are plotted in Fig. 5(a). In the test phase, we generate images of the 24 ColorChecker surfaces under the 102 light sources. Fig. 5(b) shows the mean point of the log-difference chromaticity of the test images for each illuminant. Fig. 5(c) shows the estimate result: a 45° straight line represents a perfect estimation of the 102 illuminants. We calculate distance in (χ_1, χ_2) space from each light/Munsell combination to all the light/Macbeths ones, and show which of the 102 lights is best matched.

The dots correspond to the estimate results: we see most of the dots fall on or close to perfect, except one point which associates the 93th illuminant with the 6th illuminant.

White Balance

Estimating the ambient illuminant can guide color balance for digital imaging. To further demonstrate the performance of the algorithm, we conducted experiments for carrying out the white balance for real images based on the estimated ambient illuminant temperature. A white balance algorithm typically looks for a white patch in the image, the chromaticity of which will then be the chromaticity of the illuminant. For automatic white balance, the white patch is usually evaluated as the maximum or average found in each of the three image bands separately. The scaling coefficients are then obtained by comparing the chosen white patch with the values of the three channels of a reference white. The difficulty is that maximum (or average) values of the three color

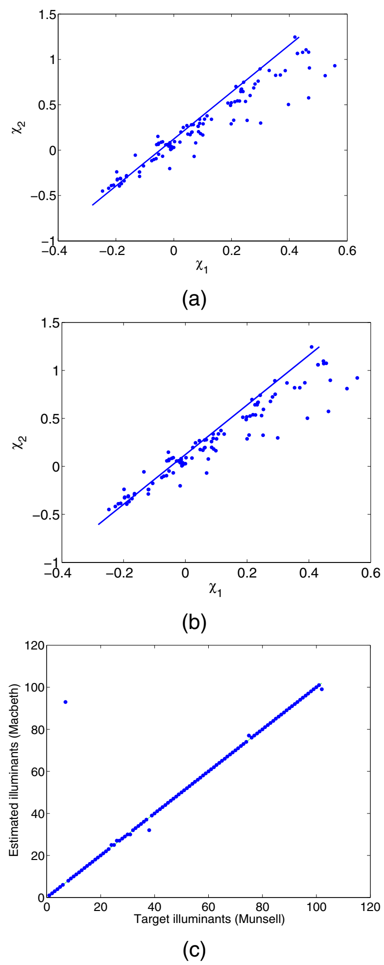

Figure 5. *Illuminant estimation. (a): Mean points of log-difference chromaticity of 461 Munsell patches for 102 illuminants, (b): Mean points of log-difference chromaticity of 24 Macbeth patches for 102 illuminants, (c): Estimate result.*

bands are not necessarily white in the scene.

This problem can be solved using our illuminant estimate approach. The key is that once we find an ambient illuminant temperature for an image, i.e. we actually classify this illuminant into one of the known illuminant clusters along the reference illuminant locus, we can explicitly know which point in the cluster corresponds to the white patch of a training ColorChecker image (supposing that we obtained the illuminant clusters using Macbeth patches). Thus, the white patch in each illuminant cluster can be taken to be the reference white color for this illuminant. In the training phase, we simply store the RGB values of the reference white patch for each known illuminant. In the testing phase, once we assign an illuminant to the test image, the corresponding reference white patch will be used to carry out white balance for the test image, as follows: $(R',G',B')^T = diag(1/R_w, 1/G_w, 1/B_w)\,(R,G,B)^T$ R_w, G_w and B_w denote the RGB values of the reference white patch, which should be normalized by the maximum value among the three values.

The advantage of this white balance scheme is that it can get around the difficulty of evaluating the white patch using color information within the image. Also it obviates computing maximum or average color values.

We captured images using a consumer HP618 camera as the imaging device. We collected image pairs of the 24 patches of the ColorChecker target under four lighting conditions: illuminant A, cool white fluorescent (CWF), the daylight D65, and the Triphosphor lamp TL84. The sharpened log-difference geometric-mean chromaticities of the images are plotted in Fig. 6(a), where four clusters corresponding to the four illuminants are shown in different colors, and each cluster has 24 dots. We then captured a test image pair for a scene with multiple objects under illuminant CWF, in Fig. 6(b,c). For speed and for display, we sampled the test images at 24 locations, evenly distributed on the images. We plot the chromaticities of these 24 sample pixels in Fig. 6(a), marked with a black star. It is obvious that these sample points mostly overlap with the CWF cluster and so the white patch of the CWF is used for white balancing the test image.

This camera has four preset white balance settings: Auto, Daylight, Fluorescent, and Tungsten. In our training and testing phases, we manually chose Daylight white balance for both flash and no-flash images, so as to effectively eliminate the white balance for our illuminant estimate because the scaling factors for the white balance can be removed in the log difference process. For comparing our white balance result, we took another image for this scene under illuminant CWF using the 'Auto' white balance function (Fig. 6(d)). The image contains subjects which are predominantly warm in color, and so the camera mistakes this for a color cast induced by a warm light source and creates a greenish color cast over the image. In contrast, our white balance result, shown in Fig. 6(e), removes the greenish effect and is closer to the Fluorescent white balance.

Summary

We have presented a practical approach for estimating ambient illuminant in a scene. The method has a number of novel features. First, it is based on a simple image formation model and obviates using complicated physical constraints on surfaces and knowledge of camera sensors and flash spectra. In our method, surfaces in the scene need not be known, and information about the camera and flash is not required. Second, in our method, a novel reference illuminant temperature locus is proposed, which specifies the path for changes of illuminant temperatures; it can be used to estimate the scene illuminant. Third, based on our estimate of ambient illuminant, an easy but more accurate white balance scheme can be carried out for automatically balancing images.

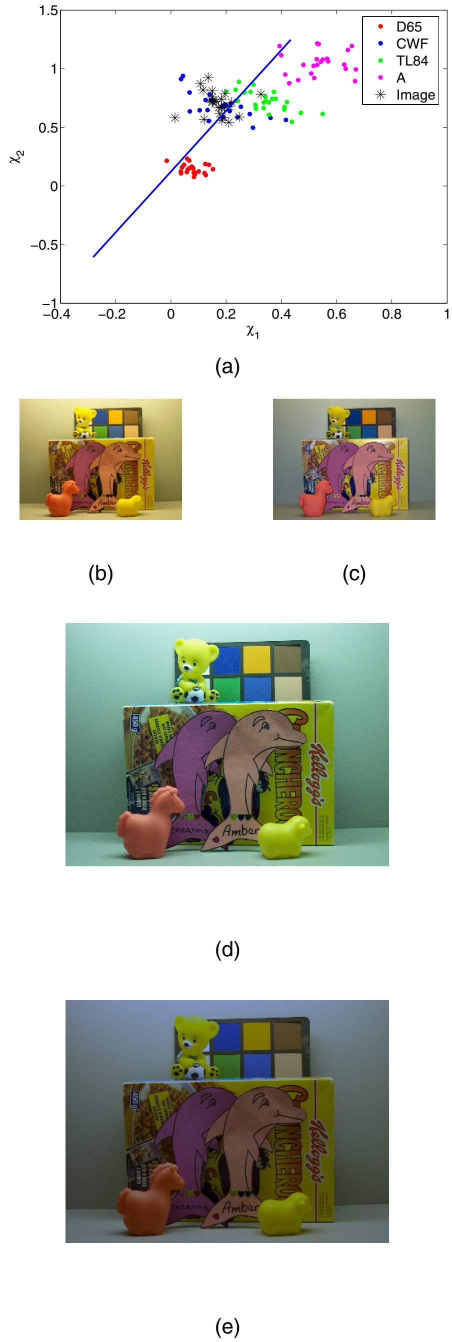

Figure 6. *White balance (HP618 camera). (a): Ambient illuminant estimate for the test image; (b): no-flash test image; (c): flashed test image; (d): Auto white balance; (e): Our white balance result based on the correct estimate of ambient illuminant.*

References

[1] G.D. Finlayson, P.M. Hubel, and S. Hordley. Color by correlation. In *5th Color Imaging Conf.*, pp.6-11, 1997.

[2] S. Tominaga and B.A. Wandell. Natural scene-illuminant estimation using sensor correlation. *Proc. of the IEEE*. 9:42-56, 2002.

[3] G. Finlayson and G. Schaefer. Single surface colour constancy. In *7th Color Imaging Conf.*, 1999.

[4] J.M. DiCarlo, F. Xiao, and B.A. Wandell. Illuminating illumination. *9th Color Imaging Conf.*, pp.27-34, 2001.

[5] G. Petschnigg, R. Szeliski, M. Agrawala, M.F. Cohen, H. Hoppe, and K. Toyama. Digital photography with flash and no-flash image pairs. *ACM Trans. Graph.* 23: 664-672, 2004.

[6] C. Lu and M.S. Drew, Automatic compensation for camera settings for images taken under different illuminants. In *14th Color Imaging Conf.*, 2006.

[7] G.D. Finlayson and S.D. Hordley. Color constancy at a pixel. *J. Opt. Soc. Am. A*. 18:253-264, 2001.

[8] G. Wyszecki and W.S. Stiles. Color Science: Concepts and Methods, Quantitative Data and Formulas. Wiley, New York, 2nd ed., 1982.

[9] G.D. Finlayson and M.S. Drew. 4-sensor camera calibration for image representation invariant to shading, shadows, lighting and specularities. In *Int. Conf. on Comp. Vision*, pp.II:473-480, 2001.

[10] G.D. Finlayson, M.S. Drew, and B.V. Funt. Color constancy: Generalized diagonal transforms suffice. *J. Opt. Soc. Am. A*, 11:1553-1563, 1994

[11] M.S. Drew and G.D. Finlayson, Spectral sharpening with positivity. *J. Opt. Soc. Am. A*. 17:1361-1370, 2000.

[12] K. Barnard, F. Ciurea, and B.V. Funt. Sensor sharpening for computational color constancy. *J. Opt. Soc. Am. A*. 18:2728-2743, 2001.

[13] K. Barnard, L. Martin, B.V. Funt, and A. Coath. A data set for colour research. Color Res. and App. 27:147-151, 2002.

[14] M.S. Drew, C. Lu, and G.D. Finlayson. Removing Shadows using Flash/Noflash Image Edges. IEEE Conf. on Multimedia and Expo, 2006.