Unsupervised illuminant estimation from natural scenes: an RGB digital camera suffices

Juan L. Nieves,* Clara Plata, Eva M. Valero, and Javier Romero

Departamento de Óptica, Facultad de Ciencias, Universidad de Granada, Campus Fuentenueva, 18071 Granada, Spain

*Corresponding author: jnieves@ugr.es

Received 4 February 2008; revised 23 May 2008; accepted 30 May 2008;
posted 9 June 2008 (Doc. ID 92252); published 3 July 2008

A linear pseudo-inverse method for unsupervised illuminant recovery from natural scenes is presented. The algorithm, which uses a digital RGB camera, selects the naturally occurring bright areas (not necessarily the white ones) in natural images and converts the RGB digital counts directly into the spectral power distribution of the illuminants using a learning-based spectral procedure. Computations show a good spectral and colorimetric performance when only three sensors (a three-band RGB camera) are used. These results go against previous findings concerning the recovery of spectral reflectances and radiances, which claimed that the greater the number of sensors, the better the spectral performance. Combining the device with the appropriate computations can yield spectral information about objects and illuminants simultaneously, avoiding the need for spectroradiometric measurements. The method works well and needs neither a white reference located in the natural scene nor direct measurements of the spectral power distribution of the light. © 2008 Optical Society of America

OCIS codes: 150.2950, 150.0150, 330.1710.

1. Introduction

The color of objects depends upon the spectral reflectance properties of their surfaces and the spectral power distribution (SPD) of the light that illuminates them [1]. Therefore the color of a natural scene captured by any conventional or digital camera can vary considerably when the ambient light changes. Most commercial digital cameras incorporate a simple white balance mechanism to solve this problem, but more sophisticated spectral imaging devices either directly measure the light impinging on the scene or derive some canonical image that is independent of the illuminant conditions. In the former situation a white reference surface has to be located within the scene and its spectral radiance measured with a telespectroradiometer [2]. In the latter case different strategies, usually referred to as color constancy algorithms, are used to obtain color-constant image descriptors [3].

Spectral imaging has been used extensively during the past decade to obtain spectral functions in each image pixel [2,4–11]. Unlike conventional imaging devices, they capture illuminant-independent images and allow accurate spectral and colorimetric reproduction of color images under any lighting conditions. The experimental setups can be very different in practice. Multispectral imaging uses several images (usually no more than a few tens) with discrete and fairly narrow bands. Hyperspectral techniques deal with imaging narrow spectral bands (up to 100) over a contiguous spectral range. Ultraspectral devices (with more than 100 bands) are typically designed for interferometer-type imaging sensors. While no formal procedure exists, in practice the spectral imaging systems used to acquire scene reflectances have to measure, immediately after acquisition, the illumination impinging upon the scene. This means that a neutral reference surface embedded in the scene should be recorded with a telespectroradiometer [6,7], which can be critical, particularly for hyperspectral devices consisting of a digital monochromatic camera with a liquid crystal tunable filter dealing with imaging over a contiguous spectral range and over a fixed exposure

0003-6935/08/203574-11$15.00/0

time for each band. If atmospheric conditions change during the band capture, the device can fail to associate the appropriate radiance with the band. Thus, spectral calibration may limit the use of spectral imaging devices to derive simplified illuminant-independent images. Different authors have relied on the underlying smoothness of signal spectra, with illuminants and spectral reflectances represented by low-dimensional models based either on principal-component analysis or independent-component analysis to develop a multispectral imaging system [12–17], but when a scene is digitally imaged under light of unknown SPD the image pixels give incomplete information about the spectral reflectances of the objects in the scene. Thus the use of spectral imaging devices to recover the spectral reflectance or spectral radiance at a single pixel requires previous measurements of the illumination impinging on a scene.

Over recent years different color-constancy algorithms have also been proposed for *a priori* illuminant estimation. In computational vision and digital photography, color constancy seeks a color-constant description of images under varying illuminant conditions instead of estimating the spectrum of the light, as spectral imaging does. The methods can usually be categorized into statistical-based models and physics-based models [18–22]. The statistical-based algorithms, which embrace the statistical knowledge of the possible image colors (possible surface reflectances and SPDs of illuminants), are robust but also very sensitive to the learning test data. Color-by-correlation and gamut-mapping algorithms, which have been widely analyzed and tested [21–23], provide excellent results, but large databases of real images are needed to obtain good algorithm performance.

The physics-based algorithms rely fundamentally on how light interacts with the object surface. Because of the complexity of this interaction it is necessary to introduce some approximations and to relax the constraints embedded within the physical phenomenon of light interaction. The dichromatic reflectance model assumes that the color signal can be decomposed into two additive terms, one originating in the interface reflectance and the other associated with the diffuse reflectance [24]. This algorithm has proved to be of particular interest in multichannel cameras in which a set of a reduced number of broadband color filters are coupled to a monochrome CCD camera [4]. A combination of statistical- and physics-based models has also been put forward as an alternative to deciding which method to apply according to the image content [20].

The target of our work here has been to estimate the SPD of the illumination of a natural scene. By using an RGB digital camera and by capturing the same image with and without a cutoff color filter in front of the camera lens the effective number of camera sensors can be increased [16]. Computational results are shown for different natural scenes and daylight illuminant conditions. The main difference from earlier multispectral imaging techniques for illuminant or radiance recoveries [16,17] is that we did not use any white reference surface in the scene. In previous studies [9,17] we found that it was possible to use a white reference surface and a linear pseudo-inverse algorithm to recover the SPD of illuminants with good spectral and colorimetric performance. A direct mapping between RGBs and illuminant spectra led to excellent spectral recoveries for daylight, which was characterized by relatively smooth spectral profiles (with the exception of some particular absorption bands beyond the visible spectrum) and was very similar for most of the hours during a clear day [9,25]. Nevertheless, spectral performance may decrease if more than three sensors are used in multispectral-illuminant estimation [17]. Mosny and Funt [26] also developed a multispectral experiment with different color-constancy algorithms and have shown that increasing the number of camera sensors from three to six or nine does not significantly improve the accuracy of illuminant estimation. For different indoor scenes under a relatively small set of artificial illuminants, their results suggest that illuminant recovery benefits from the use of a six-channel camera but that little improvement is to be gained by adding successive color filters.

In the following sections we analyze the use of a linear pseudo-inverse method for multispectral recovery of the SPD of the illuminants of natural scenes. This is particularly important for spectral imaging devices because our algorithm needs neither a white reference located in the scene nor telespectroradiometer measurement of the SPD of the light coming from this white surface. The results are discussed in terms of spectral and colorimetric performance in the analysis and synthesis of natural color images. A comparative analysis with other color-constancy algorithms that use only three color channels (max-RGB, gray-world, color-by-correlation, and dichromatic) is also presented.

2. Methods

A. Computations for Spectral Image Acquisition

When a CCD digital color camera is pointed at a surface with spectral reflectance function r^x, the response of the kth sensor for pixel x can be modeled linearly by

$$\rho_k^x = \sum_{\lambda=400}^{700} E^x(\lambda) r^x(\lambda) Q_k(\lambda) \Delta\lambda, \tag{1}$$

where $Q_k(\lambda)$ is the spectral sensitivity of the kth sensor and $E^x(\lambda)$ is the SPD of the illuminant impinging on the surface; both functions are sampled at 5 nm intervals in the visible range of 400–700 nm. But any real spectral device is affected by image noise, which can degrade both its spectral and colorimetric performance [10]. There are different noise sources, which include fixed pattern noise, dark current noise, shot noise, amplifier noise, and quantization noise.

Assuming a linear response for the CCD camera sensor outputs, the noisy sensor responses ρ' can be represented as a function of the noise-free sensor responses ρ as

$$\rho' = \rho + \eta = \mathbf{CE}, \tag{2}$$

where ρ is a $k \times N$ matrix whose columns contain the sensor responses for each of the N image pixels, η is a $k \times N$ matrix vector of uncorrelated components that affects each sensor separately, $\mathbf{C}$ is a $k \times n$ matrix whose rows contain the wavelength-by-wavelength multiplication (n means the number of wavelength samples) of the surface reflectance and the sensor sensitivities, and $\mathbf{E}$ is a $n \times N$ matrix whose rows contain the SPD of the illuminant. In what follows we simulated thermal and shot noise with standard deviations of 3% (that correspond to a signal-to-noise ratio of around 30 dB) which is an average noise value for nonrefrigerated color cameras. Although we simulated the additive noise from the calculated set of noise-free sensor responses ρ, it is straightforward from Eq. (2) that noise is also affecting the coefficients of the matrix.

An estimate of a set of test spectra may then be obtained from the corresponding set of camera responses ρ' by solving the underestimation problem, e.g., increasing the number, k, of the camera sensors by using successive cutoff filters in front of the camera lens as each cutoff filter generates three new sensitivity curves. These new sensitivities may be more or less correlated with the old ones depending on the shape of the spectral transmittance of the cutoff filters used. We first simulated digital counts by using the spectral sensitivities of a Retiga 1300 digital CCD color camera (QImaging Corporation, Canada) with 12 bit intensity resolution per channel. Camera responses ρ in Eqs. (1) and (2) are based on a set of camera sensitivities but are only used to generate a set of noise-free fictitious data. The RGB camera outputs are modified to use a three- to six-band spectral camera when no color filter ($k = 3$) or one successive color filter ($k = 6$) is used (Fig. 1). Following from previous results, [17] the color glass filter GG475 from OWIS GmbH was selected. Although the spectral transmittance of this color filter does not seem to generate an appropriate set of three new independent spectral sensitivities because it modifies the blue sensor response much more than the other two, it corresponded to that producing the best results. Some results will be cited later in case of using another cutoff filter less correlated with our camera RGB sensors.

B. Linear Pseudo-Inverse Illuminant Estimation

A priori the linear pseudo-inverse algorithm needs to apply a learning-based procedure by assigning sensor outputs to one reference surface in the scene. We begin by segmenting the images and selecting the bright areas (not necessarily the whitest) in the natural scenes. To form the sensor outputs for the training set of illuminants we use a seeded-region growing procedure. It starts with assigned luminance seeds and grows segmented areas by merging a pixel into its nearest neighboring seed area. For automatic seed selection, we first calculate the luminance component of each image as the sum of the red, green, and blue components. Next we implement a MATLAB routine that thresholds this luminance component and produces binary images containing labels for the segmented areas [27]. Similar areas are merged, and finally a single image with different seeded regions is derived for each scene and each training illuminant. Since image complexity can vary from one image to another, we further used different threshold values. Thus, if a high luminance threshold level is used, a smaller number of seeded areas is derived, while more segmented areas will be obtained for small threshold values. The segmentation procedure and its implementation is simple, computationally effective, and very fast in terms of computation times.

For spectral-illuminant recovery, it was shown that a direct relationship between the camera responses and the SPD of the illuminants can be established [17]. The method is based on a direct transformation

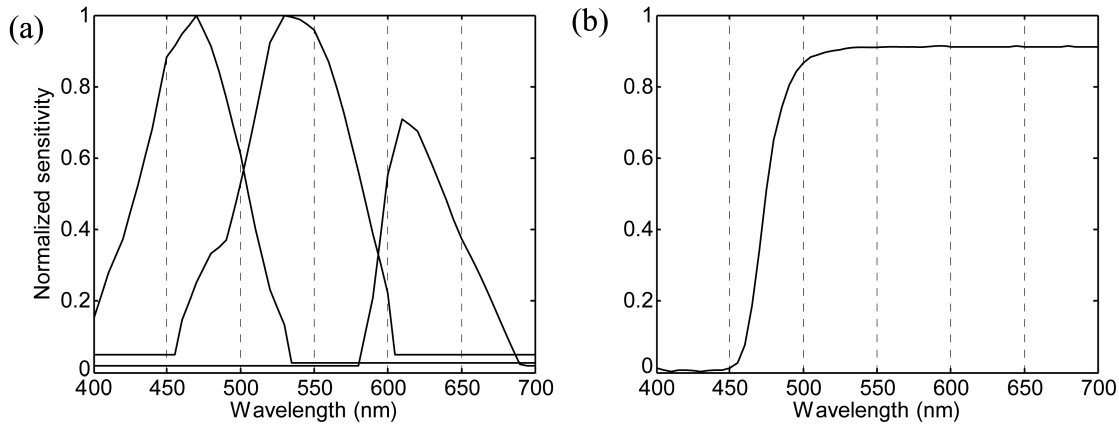

Fig. 1. (a) Spectral sensitivities of the RGB digital camera (QImaging Retiga 1300); (b) spectral transmittance of the GG475 colored-glass filter from OWIS GmbH.

between the estimated illuminant spectra E (an $n \times 1$ vector) and sensor responses ρ' (a $k \times 1$ vector) expressed by

$$E = \mathbf{F}\rho'. \tag{3}$$

In this expression, the matrix $\mathbf{F}$ is derived following a linear pseudo-inverse method by

$$\mathbf{F} = \mathbf{E_o}[(\rho'^T_{\mathbf{o}}\rho'_{\mathbf{o}})^{-1}\rho'^T_{\mathbf{o}}] = \mathbf{E_o}\rho'^+_{\mathbf{o}}, \tag{4}$$

where $\mathbf{E_o}$ (an $n \times t$ matrix, with t being the number of training illuminants) and $\rho'_{\mathbf{o}}$ (a $k \times t$ matrix) are the SPDs and their corresponding sensor outputs for the training set of illuminants, and $\rho'^+_{\mathbf{o}}$ is the pseudo-inverse matrix of $\rho'_{\mathbf{o}}$. Therefore, using this method with experimental RGB values, there is no need to use any mathematical base or to know the spectral sensitivity of the camera sensors.

Thus, suppose we are given m seeded regions for the same image under a fixed training illuminant. We can average the sensor outputs for the set of pixels belonging to the same seed and form a $k \times m$ matrix of sensor responses $\rho'_{\mathbf{o}}$ for this training illumination. Repeating this procedure for each training illuminant, a set of t sensor response matrices $\{\rho'_{\mathbf{o},1}, \rho'_{\mathbf{o},2}, \ldots, \rho'_{\mathbf{o},t}\}$ is obtained and concatenated in a rowwise manner; this procedure is iteratively repeated with each of the natural scenes to derive the overall sensor output matrix $\rho'_{\mathbf{o}}$. After the segmentation and the learning-phase procedure each $\rho'_{\mathbf{o}}$ is assigned to one individual illuminant spectrum, which is replicated m times. In this work the sensor output matrix comprised $k \times 28,834$ digital counts, as we had 28,834 pixels corresponding to the segmentation for the training scenes described below, thus yielding a matrix training set of illuminants $\mathbf{E_o}$ of $61 \times 28,834$ spectra each defined over 400–700 nm sampled at 5 nm intervals. Finally matrices $\rho'_{\mathbf{o}}$ and $\mathbf{E_o}$ are used in Eq. (4) to estimate the recovery matrix $\mathbf{F}$.

C. Hyperspectral Data, Illuminants, and Computations

A set of spectral data from nine natural scenes, seven of rural environments and two of urban environments, was used for image reproduction under different illuminant conditions. The images, which contain rocks, trees, leaves, grass, and earth, were acquired in daylight between mid-morning and mid-afternoon and were illuminated by direct sunlight from a clear sky. Additional details of the scenes, the hyperspectral imaging system, and calibration procedure can be found in [7]. For the present work different fragments of these scenes were used for the training and testing sets.

The training set of images was made up of nine different fragments, one for each scene, of 400×400 pixels and reproduced under 100 phases of daylight measured previously in our laboratory [13]. The test set comprised 12 fragments of 400×400 pixels (none of their pixels in common with the training set, but extracted from the same scenes), and the radiance for each pixel was calculated by using as illumination 25 different daylights not included in the training set.

D. Comparative Color-Constancy Methods

For comparison with the spectral-illuminant estimation method introduced in the subsection above, we implemented different standard color-constancy algorithms. Unlike spectral imaging, color-constancy models do not recover the spectrum of the illumination thrown onto a scene but seek a color image rendered under a reference light, and thus the projection of the color of the light onto the camera RGB sensors is generally enough. Because we are using bright segmented areas, we first implement a dichromatic method, which assumes that the camera sensor RGB outputs lie in a two-dimensional plane, one corresponding to the reflectance properties of the surfaces and the other describing the SPD of the interface [28].

Another successful color-constancy algorithm is color-by-correlation (CbC), which estimates the scene illuminant by determining *a priori* a set of plausible illuminant colors [22]. These two methods lead to very good color-constancy results but are computationally complex, requiring a large image data base for the training phase. Thus we have also used two additional color-constancy algorithms that estimate the color of the illuminant by using a simpler and faster, though somewhat less accurate, procedure than the previous ones. The gray-world algorithm assumes that the average reflectance in an image is achromatic and uses the fact that the mean camera sensor response for each channel gives information about the color of the light source [29]. Finally, we use the max-RGB algorithm, which estimates the illuminant color prevailing in a scene by determining the maximum sensor outputs in each channel of the color image [30]. This algorithm, which is very common in digital image processing, assumes the presence of a white point in the image and assumes that this point is that which most reflects the light over a scene.

3. Results

In evaluating the results we have analyzed the performance of the algorithm in a variety of different quality measures. We used four very common metrics to quantify both the spectral and colorimetric quality of the recovered illuminant: the goodness-of-fit coefficient (GFC), the CIELab color difference ΔE^*_{ab}, the root-mean-square error (RMSE), and the angular error (AE) [31,32]. The GFC is based on Schwartz's inequality and is defined as the cosine of the angle between the original signal $f(\lambda)$ and the recovered signal $f_r(\lambda)$; thus

$$\mathrm{GFC} = \frac{\sum_{\lambda=400}^{700} f(\lambda) f_r(\lambda)}{\left[\sum_{\lambda=400}^{700} f(\lambda)^2\right]^{1/2}\left[\sum_{\lambda=400}^{700} f_r(\lambda)^2\right]^{1/2}}. \tag{5}$$

This measurement of spectral similarity has the advantage of not being affected by scale factors. Colorimetrically accurate illuminant estimations require GFC > 0.995; GFC > 0.999 indicates quite good spectral fit, and GFC > 0.9999 an almost exact fit [13]. The CIELab color difference ΔE^*_{ab} was used to evaluate colorimetric quality and was calculated with reference to the color signal of a white patch in the scene for illuminant estimation. The third measure was the AE, defined as

$$\mathrm{AE} = \cos^{-1}(\rho_l \cdot \rho_e), \tag{6}$$

where ρ_1 and ρ_e are the color components of the light source and the algorithm's estimate of the color of the light source, respectively. Finally, the RMSE was also used.

A. Spectral and Colorimetric Performance

We obtained one recovery matrix, **F**, using the linear pseudo-inverse method, and derived two different illuminant estimation sets. The first involved only the training set of illuminants, which theoretically will give the best expected results, and the second the camera responses related to the test illuminant spectra. Table 1 summarizes the mean (and sample standard deviation) of the values of GFC, ΔE^*_{ab}, and AE where the plausible illuminant is obtained by using the camera without ($k = 3$) and with ($k = 6$) a colored cutoff filter.

There were two important results. On the one hand, for both the training and the test sets the recovery quality did not improve as the number of camera sensors increased. The GFC values are close to 0.9950, and the ΔE^*_{ab} color differences do not exceed 3 units, implying that there would be good spectral and colorimetric performance if no cutoff filter were used. The performance of the illuminant recoveries dramatically worsened with the introduction of a colored filter. These results suggest that a three-channel color camera is sufficient to make a very good estimation of the illuminant, even for the test set of images and illuminants.

On the other hand, results deriving from the control and test data sets were very similar, at least for $k = 3$ sensors. This behavior did not appear when the linear pseudo-inverse method was applied for reflectance and color signal recoveries [16], but we have to take into account that those results were derived from noise-free data. The results presented here suggest an overestimation of the learning phase of the algorithm, which is particularly important for the training set and $k = 6$ sensors.

Figure 2 shows examples of spectral-illuminant recovery for different scenes. The examples in the left-hand panels are for the digital RGB camera without a color filter, and those in the right-hand panels for the camera with one successive color filter. Even for the daylight spectra measured at dawn (not as spectrally smooth as the ones measured at midday) there are few noticeable spectral differences between the recovered and the original signals with the three-band spectral camera. Nevertheless, the plots in the right-hand panels suggest that the use of a six-band camera does not improve the recovery qualities independently of the spectral profile of the daylight.

The question immediately arises as to whether our findings for $k = 6$ sensors originate directly from a bad choice of sensor output matrix $\boldsymbol{\rho}'_{\mathbf{o}}$, thus giving a poor recovery matrix **F** in Eq. (4). Consequently we responded in two ways.

First, we started by choosing a different colored filter to generate three new distinct spectral sensitivities. An additive combination of the color glass filters BG12 and OG550 from OWIS GmbH was selected as a new successive cutoff filter, and the recovery matrix **F** was recalculated for the training set of illuminants. The correlation coefficient between the 61×3 matrices of spectral sensitivities without and with the GG475 prefilter was 0.8393. This value suggests, as expected from Fig. 1, a significant overlap among sensor sensitivities with only five of the six sensor channels producing clearly different sensors responses. In contrast, the correlation coefficient for the BG12 and OG550 combination prefilter was 0.7520, which, in principle, should be a way of increasing the effective number of spectral bands. But in the latter case the average GFC value and ΔE^*_{ab} were 0.9427 (standard deviation std = 0.0158) and 8.5 (std = 1.5), implying that there would be no improvement for $k = 6$ sensors when a very different cutoff filter was used. Although pseudo-inverting Eq. (3) can be an optimal solution to solve for the matrix **F**, the procedure involves not only $\boldsymbol{\rho}'$, which are the sensor responses for a test condition, but also $\boldsymbol{\rho}'_{\mathbf{o}}$, which are the sensor responses for a training set condition. Therefore, noise can degrade the optimal mapping between RGB and

Table 1. Linear Pseudo-Inverse Illuminant Estimation Using a 50% Luminance Threshold[a]

	GFC		ΔE^*_{ab}		AE (deg)	
Filter	Mean	Std	Mean	Std	Mean	Std
Control performance						
Without filter	0.9952	0.0080	2.9	2.0	4.9	3.3
With GG575 filter	0.9442	0.0148	8.7	1.5	18.9	2.6
Test performance						
Without filter	0.9943	0.0119	3.2	2.6	5.3	4.6
With GG575 filter	0.9616	0.0097	7.6	1.1	15.8	1.9

[a]The sensor output matrix was derived using a local mean calculated individually for each seeded area.

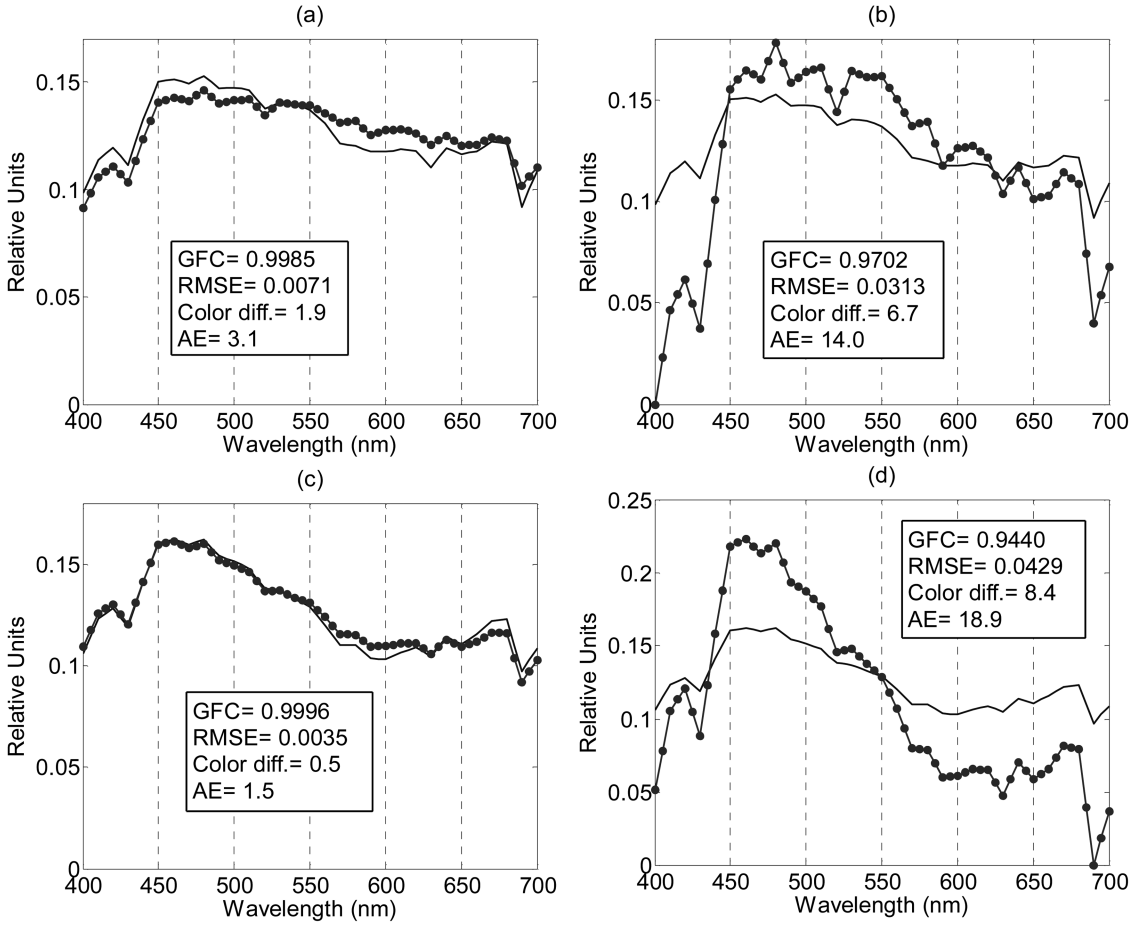

Fig. 2. Examples of the spectral-illuminant recovery (left-hand panels) with $k = 3$ sensors and (right-hand panels) with $k = 6$ sensors for different illuminants and scenes (—, original spectrum; •, recovered spectrum).

spectra after solving for Eq. (3). We should keep in mind that we are increasing the number of camera sensors by adding broadband filters in front of our digital RGB camera lens and not by adding narrowband filters as hyperspectral devices do.

Second, we recalculated the matrix $\boldsymbol{\rho}'_{\mathbf{o}}$, using different luminance threshold values in determining the seeded regions for each scene. Figure 3 shows an example of the influence of the segmentation procedure on the quality of illuminant recoveries. The upper row illustrates how the number of seeded regions (different colors in the plots) changes as the threshold value increases. The lower part of the figure shows the cumulative distribution plots for the spectral performance of this scene under the 100 training illuminants, and the whole figure shows the proportion of spectral recoveries that takes on values of less than or equal to some GFC values, without (left-hand plot) and with (right-hand plot) the prefilter. What is clear is that the performance of the linear pseudo-inverse algorithm is very good for the three-channel camera, giving very similar results for all threshold values, but the results worsen for the six-channel camera, which yields a wide dispersion in the spectral performance according to the threshold value used. Figure 4 summarizes the results for the control set of illuminants.

The effects of threshold values and number of filters have been tested by a repeated-measures analysis of variance of three factors: the number of sensors (two levels, $k = 3$ or $k = 6$), the metric (four levels, GFC, RMSE, ΔE^*_{ab}, and AE) and the threshold (five levels, $u = 10, 30, 50, 70, 90$). We obtain very good spectral and colorimetric accuracy for the three-channel camera without prefilter. Our results show that both the spectral and colorimetric performance for $k = 3$ sensors do not depend on the number of the selected seeded regions ($p \gg 0.05$). The average GFC and ΔE^*_{ab} color difference for all of the threshold values are 0.995 and 2.9, respectively. In contast, the results for $k = 6$ sensors not only worsen but are also different for the various threshold values used. We found an average GFC and ΔE^*_{ab} of 0.9420 and 8.8, respectively, indicating a poor quality of illuminant recovery with the six-channel camera. There are also significant differences for all the metrics ($p \ll 0.05$) according to the threshold value. *Post hoc* comparisons suggest significant differences between GFC, ΔE^*_{ab}, and AE for a threshold of 70% ($p \ll 0.05$) and close to significant for the RMSE ($p = 0.04$).

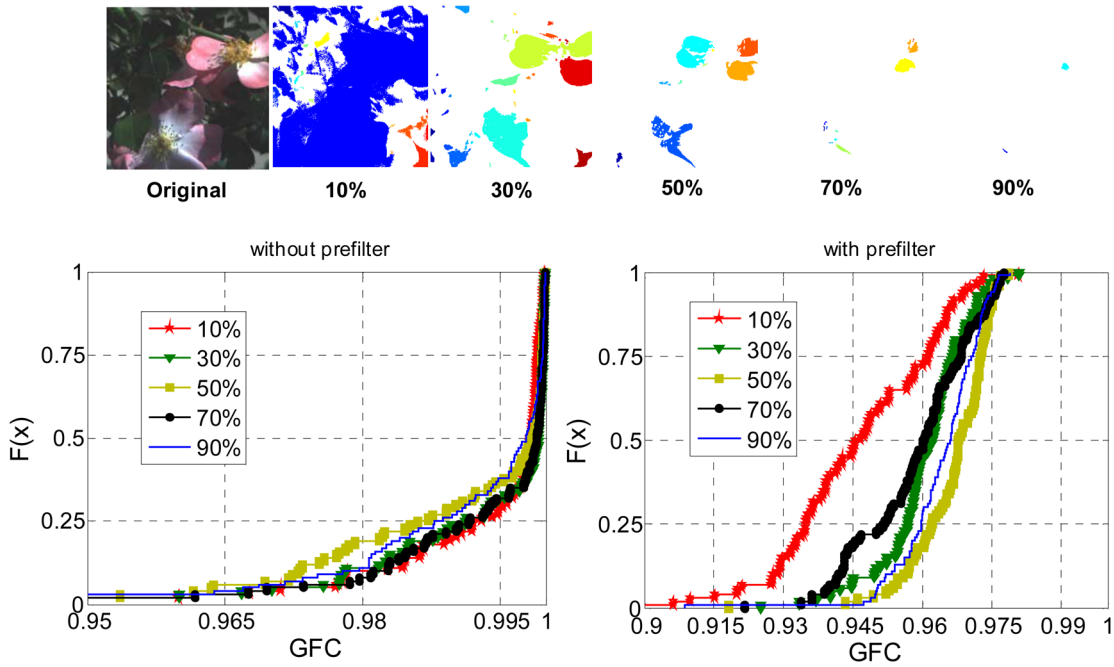

Fig. 3. (Color online) Example of the seeded regions obtained by using different luminance threshold values. Results for the cumulative distributions of GFCs with and without a prefilter are also shown for this scene.

B. Chromatic Differences in Natural Color Image Reproduction

The above results suggest that a three-channel camera is sufficient for good spectral and colorimetric illuminant estimation. But can we reproduce natural image colors with the same accuracy? To test this point we used Eqs. (1) and (2) to derive the original and recovered sensor responses when each hyperspectral image was reproduced under the original and estimated illuminant. The error in color image reproduction is calculated as

$$\mathrm{RGBerror}_x = \left[\frac{1}{3}\sum_{i=1}^{3}(\rho_i^x - \underline{\rho}_i^x)^2\right]^{1/2}, \qquad (7)$$

where ρ^x and $\underline{\rho}^x$ are the digital counts for R, G, and B of the original and recovered image, respectively, at pixel x. The RGBs were all normalized to values within the range 0–255.

Figure 5 illustrates the visual significance of the RGB color errors for different scenes when they are reproduced under the estimated illuminant spectrum with and without a prefilter. We show examples of the 95% percentile results for two very different natural images: one corresponding to a rural landscape, which contains a reduced color range centered around green hues, and an urban landscape, which contains a variety of different surfaces ranging through a more complex gamut of colors. The error histogram for the rural scene supports our previous findings and shows that we achieved good color reproduction using $k = 3$ sensors, with RGB errors below five digital counts. The use of more sensors with this image also leads to good color preproduction, with a great percentage of RGB errors around five

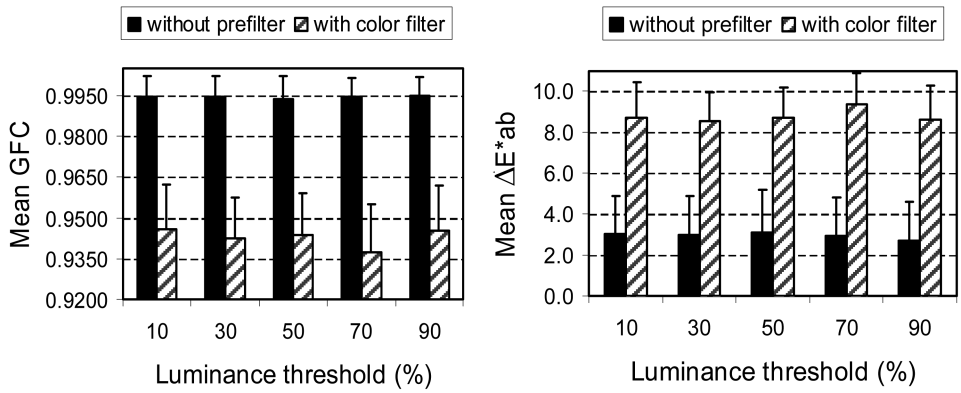

Fig. 4. Mean spectral (left) and colorimetric (right) performance for the linear pseudo-inverse algorithm using different luminance threshold values.

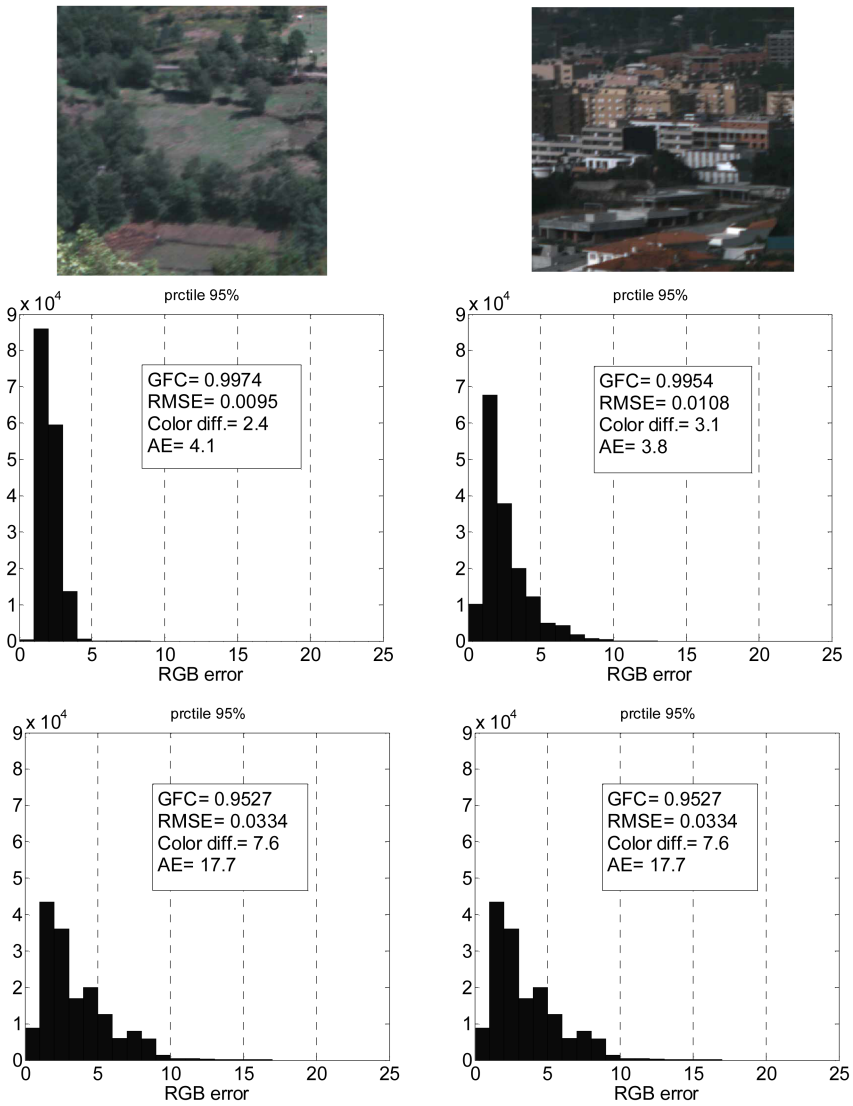

Fig. 5. (Color online) Visualized images of original scene fragments and the histogram of RGB errors. The upper histograms are for the camera without a prefilter, and the lower ones for the six-band camera.

digital counts. Nevertheless, the differences between the use of a three- or six-channel camera are more evident for the urban scene. The color difference ΔE^*_{ab} obtained for the estimated illuminant when the algorithm was applied for the two scenes was very similar in both cases (7.1 units for the rural scene and 7.6 for the urban one), but the RGB error histogram suggests more errors when the scene encompasses a complex gamut of colors.

C. Color-Constancy Comparative Results

In this section we compare the performance of the linear pseudo-inverse algorithm to other color-constancy algorithms. For comparison's sake we used three estimates of the recovery matrix, **F**, in the linear pseudo-inverse algorithm for a three-channel color camera without successive cutoff filters. The first estimate (LPI) was obtained by segmenting the test images using a 50% luminance threshold value. The second version involved only those image pixels within 90% of the maximum luminance level. We called this version the white-patch linear pseudo-inverse (WLPI) because of its similarity to the max-RGB color-constancy algorithms. The third version, or max-linear pseudo-inverse (MLPI), on the other hand, used a luminance threshold value of 10% and tended to compute the average of the RGBs in all the seeded regions.

Table 2 summarizes the performance of the different illuminant estimation algorithms. Because color-constancy approaches recover the color of the light, the results only show the AE measurements. It is clear that the linear pseudo-inverse method performs very well in comparison with the other color-constancy algorithms. Our results show that the color-by-correlation and all three versions of the linear pseudo-inverse method outperform the max-RGB, gray-world, and dichromatic approaches as far as mean RGB error and AE metrics are concerned. The color-by-correlation paradigm is slightly

Table 2. Summary of Performance (AE, degrees) for Linear Pseudo-Inverse Illuminant Estimation and Different Color-Constancy Algorithms

Criterion	Max-RGB	Gray-World	CbC	Dichromatic	LPI	MLPI	WLPI
Mean	13.7	19.2	4.8	8.1	5.4	6.1	5.1
Std	1.6	2.5	3.1	3.8	4.6	3.9	4.9

better in terms of AE and also gives very good results for the RMSE (of only 0.0019). The max-linear pseudo-inverse gives the worst results, suggesting a limitation of the method if you overload the linear pseudo-inverse algorithm. Nevertheless, with regard to color differences, the corresponding average ΔE^*_{ab} is around 3 units, indicating a good colorimetric color reproduction.

4. Discussion

We have investigated the question of how to apply a linear pseudo-inverse method for unsupervised illuminant recovery from natural scenes. The method we have introduced is a spectral-imaging learning-based algorithm that directly relates camera sensor outputs and illuminant spectra. The RGB sensor outputs can be modified by the use of a successive cutoff color filter, and thus the algorithm needs no information about the spectral sensitivities of the camera sensors or eigenvectors to estimate the SPDs of illuminants.

Our results suggest that daylight spectra can be recovered with acceptable spectral and colorimetric accuracy with a three-band camera (e.g., a digital RGB camera without a prefilter). Although linear pseudo-inverse approaches have been applied successfully for color signals and natural reflectance recoveries [16], the results suggest serious limitations of the algorithm for multispectral-illuminant recovery when more than three sensors are used. One reason for the poor results for the six-band color camera could be the inappropriate choice of the particular color cutoff filter, but we have obtained similar low spectral and colorimetric performance when other uncorrelated sets of RGB sensors were used. In a previous computational work [16] we found that a noise-free RGB camera coupled with color filters provided significantly better recovery of radiance and spectral reflectances in natural scenes than an RGB camera alone. However, later results suggest that digital cameras affected by high noise levels do not improve the performance of the algorithm for recovering illuminants [17]. This also agrees with the results anticipated by Mosny and Funt [26], who found minor improvements by extending the number of channels from RGB cameras to six and nine. Thus the influence of noise in the recovering process could be more relevant than the selection of filters, at least for spectral applications using digital RGB cameras with a prefilter.

With our results, we seek not only to establish the limits of spectral direct mapping for illuminant estimation, but also to provide a new multispectral color-constancy approach. By combining the device with the appropriate computations, spectral information about objects and illuminants can be obtained simultaneously without any spectroradiometric measurement. But will the method work well for a wider variety of indoor scenes that include fluorescent illuminants? Nothing precludes the application of our method to indoor scenes, although it can fail to describe some kind of fluorescent illuminants appropriately. These limitations were already demonstrated in previous work [17], and some examples with indoor scenes and artificial illuminants suggest similar performance here. Figure 6 illustrates the recovery results for an indoor scene and two artificial illuminants without a prefilter. In this example the training set of illuminants was composed of 82 SPDs of fluorescent-type illuminants, and the test set comprised a set of 20 commercial fluorescent and incandescent lights that were not included in the training set (see [17] for more details about these SPDs); two fragments of the toys hyperspectral scene [7] were used with the training set of illuminants, and one fragment with the test illuminants (with none of their pixels in common with the training set). Results suggest good colorimetric performance, but we find a dependency on the spectral profile of the artificial illuminant. As expected, an RGB digital camera without prefilter does not provide recoveries of artificial lighting as well as with daylight, and it can fail to describe some kind of fluorescent illuminants appropriately [17]. Nevertheless this simple example demonstrates that the linear pseudo-inverse method for unsupervised illuminant recovery could also work for indoor scenes with an adequate training.

The advantage of the linear pseudo-inverse illuminant estimation algorithm is that it recovers not only the color of the light but also the illuminant spectrum. Comparing this method with other color-constancy algorithms, the spectral and colorimetric performances surpass other approaches. Our results are close in performance to those deriving from the color-by-correlation method, but avoid a huge database for training and show that multispectral imaging fundamentals can also be used for illuminant estimation in color-constancy applications.

This work was supported by the Spanish Ministry of Education and Science and the European Fund for Regional Development (FEDER) through grant number FIS2007-60736. The authors thank A.L. Tate for revising their text.

References

1. B. A. Wandell, "The synthesis and analysis of color images," IEEE Trans. Pattern Anal. Mach. Intell. **9**, 2–13 (1987).
2. J. Y. Hardeberg, "Acquisition and reproduction of color images: colorimetric and multispectral approaches," Ph.D. dissertation (Ecole Nationale Superieure des Telecommunications, 1999).
3. M. Ebner, *Color Constancy* (Wiley, 2007).

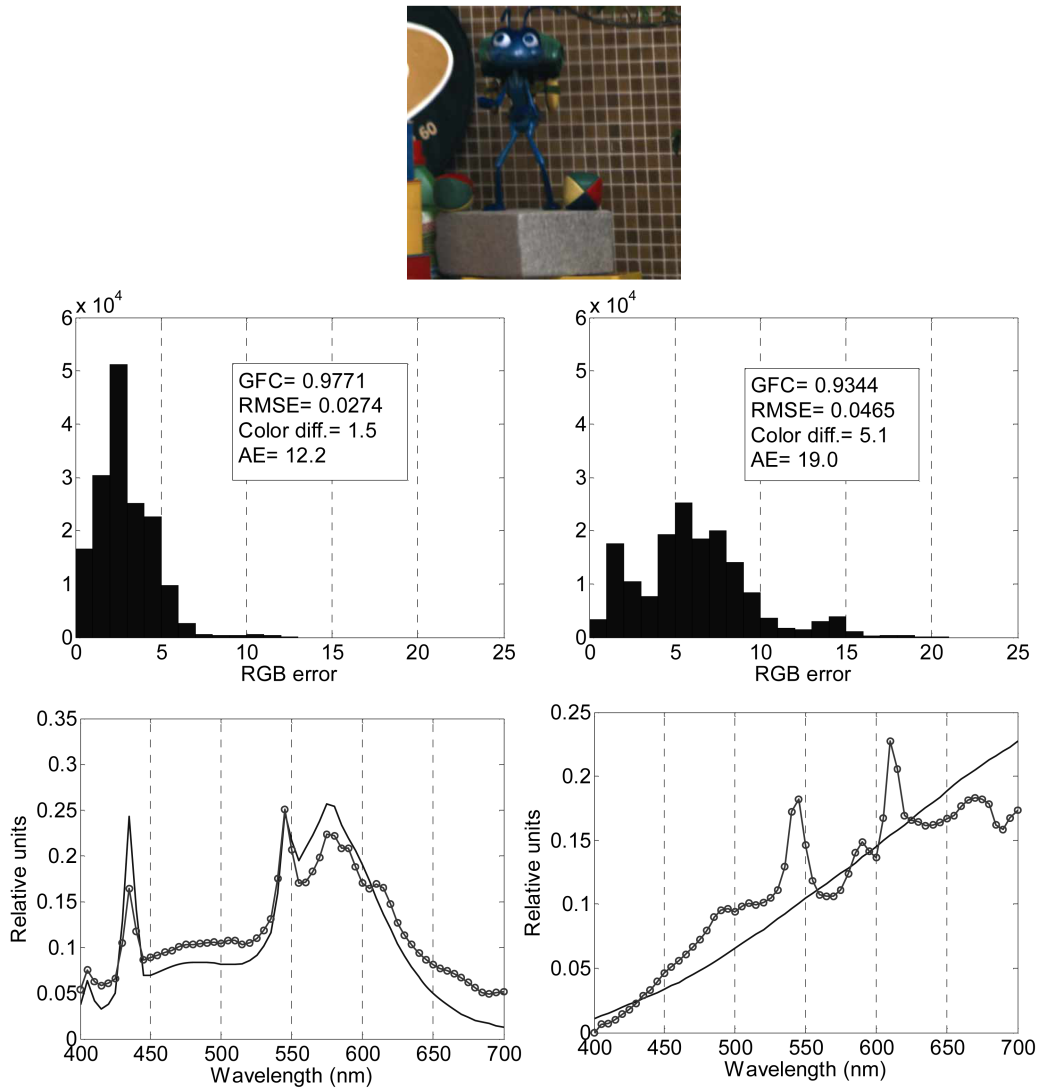

Fig. 6. (Color online) Visualized image of an indoor scene fragment and the histogram of RGB errors. Results are examples of the spectral-illuminant recovery with $k = 3$ sensors and two artificial illuminants (—, original spectrum; •, recovered spectrum).

4. S. Tominaga, "Multichannel vision system for estimating surface and illumination functions," J. Opt. Soc. Am. A **13**, 2163–2173 (1996).
5. F. H. Imai and R. Berns, "Spectral estimation using trichromatic digital cameras," in *Proceedings of the International Symposium on Multispectral Imaging and Color Reproduction for Digital Archives* (Society of Multispectral Imaging of Japan, 1999), pp. 42–49.
6. C.-C. Chiao, D. Osorio, M. Vorobyev, and T. W. Cronin, "Characterization of natural illuminants in forests and the use of digital video data to reconstruct illuminant spectra," J. Opt. Soc. Am. A **17**, 1713–11721 (2000).
7. S. M. C. Nascimento, F. P. Ferreira, and D. H. Foster, "Statistics of spatial cone excitation ratios in natural scenes," J. Opt. Soc. Am. A **19**, 1484–1490 (2002).
8. V. Cheung, S. Westland, C. Li, J. Hardeberg, and D. Connah, "Characterization of trichromatic color cameras by using a new multispectral imaging technique," J. Opt. Soc. Am. A **22**, 1231–1240 (2005).
9. J. L. Nieves, E. M. Valero, S. M. C. Nascimento, J. Hernández-Andrés, and J. Romero, "Multispectral synthesis of daylight using a commercial digital CCD camera," Appl. Opt. **44**, 5696–5703 (2005).
10. N. Shimano, "Evaluation of a multispectral image acquisition system aimed at reconstruction of spectral reflectances," Opt. Eng. **44**, 107005 (2005).
11. H.-L. Shen, J. H. Xin, and S.-J. Shao, "Improved reflectance reconstruction for multispectral imaging by combining different techniques," Opt. Express **15**, 5531–5536 (2007).
12. S. Tominaga and B. A. Wandell, "Standard surface-reflectance model and illuminant estimation," J. Opt. Soc. Am. A **6**, 576–584 (1989).
13. J. Hernández-Andrés, J. Romero, J. L. Nieves, and R. L. Lee, Jr., "Color and spectral analysis of daylight in southern Europe," J. Opt. Soc. Am. A **18**, 1325–1335 (2001).
14. S. Tominaga, "Natural image database and its use for scene illuminant estimation," J. Electron. Imaging **11**, 434–444 (2002).
15. S. Tominaga and B. A. Wandell, "Natural scene-illuminant estimation using the sensor correlation," in Proc. IEEE **90**, 42–56 (2002).
16. E. M. Valero, J. L. Nieves, S. M. C. Nascimento, K. Amano, and D. H. Foster, "Recovering spectral data from natural scenes

with an RGB digital camera," Color Res. Appl. **32**, 352–360 (2007).
17. J. L. Nieves, E. M. Valero, J. Hernández-Andrés, and J. Romero, "Recovering fluorescent spectra of lights with an RGB digital camera and color filters using different matrix factorizations," Appl. Opt. **46**, 4144–4154 (2007).
18. M. D'Zmura and G. Iverson, "Color constancy. I. Basic theory of two-stage linear recovery of spectral descriptions for lights and surfaces," J. Opt. Soc. Am. A **10**, 2148–2165 (1993).
19. D. H. Brainard and W. T. Freeman, "Bayesian color constancy," J. Opt. Soc. Am. A **14**, 1393–1411 (1997).
20. G. Shaefer and S. Hordley, "A combined physical and statistical approach to color constancy," in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)* (IEEE Computer Society, 2005), Vol. **1**, pp. 148–153.
21. S. D. Hordley and G. D. Finlayson, "Reevaluation of color constancy algorithm performance," J. Opt. Soc. Am. A **23**, 1008–1020 (2006).
22. G. D. Finlayson, S. D. Hordley, and P. M. Hubel, "Color by correlation: a simple, unifying framework for color constancy," IEEE Trans. Pattern Anal. Machine Intell. **23**, 1209–1221 (2001).
23. D. A. Forsyth, "A novel algorithm for colour constancy," Int. J. Comput. Vision **5**, 5–36 (1990).
24. S. A. Shafer, "Using color to separate reflection components," Color Res. Appl. **10**, 210–218 (1985).
25. M. A. López-Álvarez, J. Hernández-Andrés, E. M. Valero, and J. Romero, "Selecting algorithms, sensors and linear bases for optimum spectral recovery of skylight," J. Opt. Soc. Am. A **24**, 942–956 (2007).
26. M. Mosny and B. Funt, "Multispectral color constancy: real image tests" Proc. SPIE **6492**, 64920S (2007).
27. N. Otsu, "A threshold selection method from grey-level histograms," IEEE Trans. Syst. Man. Cybern. **9**, 62–66 (1979).
28. G. Schaefer, "Robust dichromatic colour constancy," in *Image Analysis and Recognition. Part II*, A. Campilho and M. Kamel, eds. (Springer 2004), pp. 257–264.
29. G. Buchsbaum, "A spatial processor model for object color perception," J. Franklin Inst. **310**, 1–26 (1980).
30. V. C. Cardei and B. Funt, "Committee-based color constancy," in *Proceedings of the IS&T/SID Seventh Color Imaging Conference* (Society for Imaging Science and Technology, 1999), pp. 311–313.
31. F. H. Imai, M. R. Rosen, and R. S. Berns, "Comparative study of metrics for spectral match quality," in *Proceedings of the First European Conference on Colour in Graphics, Imaging and Vision* (Society for Imaging Science and Technology (2002), pp. 492–496.
32. M. J. Vrhel, R. Gershon, and L. S. Iwan, "Measurement and analysis of object reflectance spectra," Color Res. Appl. **19**, 4–9 (1994).