

Trichromatic red–green–blue camera used for the recovery of albedo and reflectance of rough-textured surfaces under different illumination conditions

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

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

*Corresponding author: cplata@ugr.es

Received 18 February 2009; revised 8 May 2009; accepted 3 June 2009;
posted 4 June 2009 (Doc. ID 107721); published 22 June 2009

Photometric-stereo techniques are based on the fact that image intensity depends upon the orientation of the surface with regard to the source of the illumination and its spectral reflectance. They are of special interest when dealing with rough surfaces because they usually present shadowed regions where sudden illumination changes might be found. In the present work we introduce an extension of the four-source photometric-stereo algorithm to color images that is able to recover the surface spectral reflectance of objects captured with a red–green–blue (RGB) camera. This method allows image rendering, even for rough-textured surfaces, under different directions of the impinging illumination. In addition, the introduction of spectral recovery techniques applied to the albedo and spectral reflectance from rough surfaces offers the possibility of image rendering for scenes captured under sources of illumination differing in spectral distribution. Using albedo instead of RGB information helps to avoid any shadows or highlights that might falsify results. One of the advantages of this spectral-based photometric-stereo method is that it can recover not only the albedo values, but also the spectral reflectance spectrum of an object's surface on a pixel-by-pixel basis, as can be done with more complex hyperspectral imaging devices involving a camera coupled to an extensive set of narrowband filters. © 2009 Optical Society of America

OCIS codes: 330.1730, 330.1690.

1. Introduction

The color of a surface depends upon its spectral reflectance properties and the spectral power distribution (SPD) of the light impinging upon it [1]. Since spectral reflectance is independent of the illuminant used, it serves to characterize an object completely. Spectral imaging has been used extensively during the past decade to obtain spectral functions for each image pixel [1–7]. The main advantage of spectral imaging in comparison with conventional spectroradiometric measurements is that it can provide spectral information from which spectral radiance or reflectance can be recovered at each image pixel.

Usually a spectral system consists of an RGB or monochrome digital camera coupled to a number of wideband or narrowband color filters, ranging from just three to hundreds of components in an ultraspectral system [2–4]. The starting point for spectral imaging methods is the signals given by a CCD camera. Even if we are imaging the surface of a single material, the values at different points may be different due to the shape and/or roughness of the surface and the geometry of its illumination, even for high signal-to-noise ratio capturing devices [8]. This will result in the spectral estimation algorithms giving different reflectance values for pixels belonging to the same surface.

Photometric-stereo techniques have long been used to recover 3D information about surfaces. These techniques can be classified into what Woodham [9]

0003-6935/09/193643-11\$15.00/0
© 2009 Optical Society of America

refers to as “direct methods” and “indirect methods.” Direct methods are those that try to measure distance ranges directly, while the indirect methods attempt to determine distance by measuring parameters calculated from images of the illuminated objects. Shape from photometric-stereo techniques was conceived by Woodham in the early 1980s and has since been extensively studied both theoretically and experimentally. All approaches published since then can be classified according to the assumptions the authors make about the surface they are dealing with and the type of problem they want to solve. Thus, we might assume that a rough surface behaves in a Lambertian fashion [10,11], that we should use a more general method when the surface behaves in a more complex way [12–18], that the illumination vectors may or may not be known *a priori* [19], that single or multiple images can be used in the recovery process [10,20,21], that illumination sources may or may not be spectrally identical [14,22,23], or if there are or are not cast shadows in the images [24,25]. Another classification comes from the kind of image used by the algorithm, either grayscale or colored. Most photometric-stereo techniques only take into account grayscale images, although some reports can be found describing the extension of this method to color images [14,26–32]. One of these techniques is based on the fact that information concerning the color image of a Lambertian surface illuminated by a single light source is irrelevant since the photometric equations for individual color bands are linearly dependent. An efficient way to exploit this irrelevance is to use a conventional photometric-stereo method relying upon a single color image of a Lambertian surface under complex lighting conditions, rather than three grayscale images [20,28]. The surface should be illuminated by several spectrally distinct light sources deriving from different directions that are not coplanar. This method is called “shape from color.”

The color photometric-stereo method proposed by Christensen and Shapiro [30] uses the notion of the shading function, which maps surface normals to the color space under a given illumination. This method is implemented by means of lookup tables representing the inverse shading functions for a given imaging configuration, which were constructed using a calibrating sphere. The disadvantage of this method is that the surface should either be uniformly colored or its color should form distinct separable clusters in the color space, which severely restricts the choice of acceptable surfaces.

Barsky and Petrou [14] proposed a method based on the four-source photometric-stereo approach. The pixel information obtained with a camera with A channels can be represented by an A -dimensional vector called body color. For a Lambertian surface patch, the three color pixels corresponding to three different directions of illumination are collinear in the RGB space and differ only by a scalar factor, the shading of the patch under a particular illumina-

tion. Since any errors introduced may disturb collinearity they use principal components analysis (PCA) to find their principal direction, which is related to the chromaticity of the body’s color.

Thus photometric stereo is based on the fact that image intensity depends upon the orientation of the surface with regard to the source of the illumination and its spectral reflectance. If several images are taken from the same point of view but with different lighting directions, any variation in pixel intensity in these images will be due to changes in the relative positions of the light and the surface. Photometric-stereo techniques recover three-dimensional (3D) information from the surface, given usually in terms of a three-component normal vector. This means that at least three images will be needed to create a system that allows us to obtain these three normal components. Multiple images of the same scene under different illumination conditions create an overconstrained system, which is resolved for the surface shape by minimizing the total cost. Therefore, these constraints are used to calculate the normal vectors, which represent the surface orientation of any point on the surface, and the albedo (defined as the fraction of the incident light reflected by a surface, which in this case is filtered by the camera spectral sensitivity and can be affected by light intensity), which describes the reflection properties of the surface.

Photometric-stereo techniques are of special interest when dealing with rough surfaces because, on this kind of surface, there are shadowed regions where sudden illumination changes might be found. Previous studies [33] have shown that reflectance values can be mapped onto the surface of objects to render realistic images under different lighting and viewing conditions. But these authors do not analyze how the spectral reflectance of an object’s surface influences the sensor responses when the photometric approach is used. Their approach works well for image rendering under different incident illumination, but they do not examine the more realistic case where illumination can change not only in intensity but also in its spectral profile.

Precise 3D data and spectral data are presently acquired independently with laser scan equipment and hyperspectral systems; however, this equipment is expensive and the combination of their data to obtain a unified 3D and spectral representation is difficult. Developing a simple system that allows simultaneous acquisition of these quantities in the same sample is certainly of great practical interest. Although the technique of photometric stereo has been around for many years and in many versions, the original contribution of the present work is its extension to derive spectral reflectances.

The photometric-stereo technique presented is an extension of the four-source photometric-stereo to color images. This technique provides albedo and a normal vector in each pixel of an image of a real, colored, rough surface. Because of the shape of a surface or its roughness, its appearance can change

when the direction of illumination varies. Since the albedo values are not influenced by shadows and highlights, they will give a realistic description of the appearance of the surface.

Spectral recovery methods use RGB values as a starting point. If one pixel of a surface is affected by a shadow or a highlight, the corresponding RGB value will not be the one relative to the imaged material, and the reflectance recovered from it will be a fake. That is why, in this work, we propose to recover reflectances from albedo values instead of RGB values. The algorithm used here to recover spectral reflectance is the so-called linear pseudoinverse method, and it is used together with a supervised method to find the best set of training samples for the spectral recovery process.

Some experiments are presented, too: first we used the albedo together with normal information recovered with the photometric-stereo technique to simulate images under different illumination conditions. These simulations were then used to check the accuracy of our proposed photometric-stereo algorithm. Subsequently, reflectance information obtained from the albedo was used to change not only the geometry of illumination but also the SPD of the illumination source.

2. Method

This section is structured as follows: in Subsection 2.A we describe the basics of gray photometric stereo; in Subsection 2.B the photometric-stereo approach is extended to color images; in Subsection 2.C we introduce the method used to recover spectral reflectances from color images; and finally the acquisition procedure and datasets are described in Subsection 2.D.

A. Gray Photometric Stereo for Lambertian Surfaces

Two assumptions are usually made in the photometric-stereo approach [13]. The surface is not perfectly smooth and is composed of microfacets, which have normal vectors that are distributed about the normal vector of the approximating smooth surface, and the surface is lit from a single source. The coordinate system is chosen so that the image plane coincides with the xy plane and the z axis coincides with the viewing direction (Fig. 1). Thus the surface can be described by a two-dimensional (2D) height function $z = S(x, y)$. The gradient component can be defined for every point of the surface as

$$p(x, y) = \frac{\partial S(x, y)}{\partial x}, \quad q(x, y) = \frac{\partial S(x, y)}{\partial y}, \quad (1)$$

and the normal unit vector, \mathbf{N} , as

$$\mathbf{N} = \frac{1}{\sqrt{p^2 + q^2 + 1}}(p, q, -1)^T, \quad (2)$$

where T denotes the transpose matrix.




Fig. 1. Definition of the important vectors and reflectance angles: \mathbf{R} , viewer vector; \mathbf{L} , illuminant vector; \mathbf{N} , normal vector; i , angle of incidence; e , angle of emittance; g , phase angle.

Let us consider a Lambertian surface patch with albedo ρ and normal \mathbf{N} , lit from one source with direction \mathbf{L} (unitary vector), as shown in Fig. 1. If we assume the intensity light to be unitary, the intensity of each pixel of the sample captured by the camera can be expressed as [9]:

$$I = \rho(\mathbf{L} \cdot \mathbf{N}), \quad (3)$$

where I is the intensity in one pixel, ρ is its albedo value, \mathbf{L} is the 1×3 direction of the illumination vector, and \mathbf{N} is the 1×3 normal vector.

Our goal is to recover the normal vector of a surface at each point that is a three-component vector and thus, having three intensities in each pixel, we can create a system that allows us to obtain those components. Hence, if we light the surface with three light sources with directions L^1, L^2 , and L^3 the intensities of the pixels thus obtained can be expressed as

$$I^k = \rho(L^k \cdot \mathbf{N}), \quad (4)$$

where $k = 1, 2, 3$ represent the illuminant directions and (\cdot) the scalar product of two vectors. The pixel intensities can be stacked to obtain the 3×1 pixel intensity vector $\mathbf{I} = (I^1, I^2, I^3)^T$ and the light vectors can also be stacked row-wise to form the 3×3 illumination matrix $[\mathbf{L}] = (L^1, L^2, L^3)^T$. Equation (4) can then be rewritten in matrix form:

$$\mathbf{I} = \rho[\mathbf{L}]\mathbf{N} \quad (5)$$

If the three light directions, L^k , do not lie on the same plane, matrix $[\mathbf{L}]$ is nonsingular and can be inverted to give

$$[\mathbf{L}]^{-1}\mathbf{I} = \rho\mathbf{N}. \quad (6)$$

Since \mathbf{N} has unit length, both the normal (as the direction of the obtained vector) and albedo (as its length) can be recovered.

Controlling the lighting conditions is very important for the accuracy of this algorithm. It is necessary to avoid ambient light and to have invariant incident light intensity for the different directions of illumination. That way we can be sure that the differences between the intensity values of the same pixel in different images is due only to the change in the orientation of the light source. It can be achieved by keeping constant the distance between the light source and the imaged object. In the present work, the different orientations of illumination have been achieved by fixing the position of the source and making the sample and the camera go around together. It allows us to keep the same distance between the sample and the illumination source, setting the sample in the rotational center of the system.

Photometric-stereo techniques are based on the constraint of perfect diffuse surfaces, known as Lambertian surfaces, which reflect light equally in all directions. Real surfaces can show non-Lambertian behavior when they are illuminated in certain directions, causing highlights and leading the photometric-stereo algorithm to fail. Another source of error can be the presence of cast shadows, a situation that occurs when a point of the surface cannot be seen by the light source because another point of the same source blocks the path of the light. Surfaces with such behavior have been avoided here. In this work we have used the so-called four-source photometric-stereo technique [14], where four lighting sources are used instead of three. With four vectors, it is possible to make four different combinations of three vectors. This allows us to use only the combination in which the pixel intensities show Lambertian behavior. Hence, from each pixel we can get four albedos and normal vectors (one from each one of the four possible combinations). If there is a highlight in any of the images, its albedo value will be higher than the rest. So, calculating the standard deviation for the set of four albedo values, we can discover whether we have a non-Lambertian (highlight) behavior for any of the illumination directions. If the standard deviation is higher than a threshold, the value of which comes from previous experimental results [8], the combination that gives the higher albedo value is disregarded and the three albedo values and the normal vector will be averaged to get just one albedo and one normal in each pixel. If the standard deviation is beneath the threshold, the average will be made with all four albedo and normal vector values because it means that we do not have any highlight in any of the four images of that pixel. Once we have recovered the albedo and the normal vector we can use this information to render a scene under different directions of illumination simply by changing the illumination vector L in Eq. (4).

B. Color Photometric-Stereo Algorithm

Here we propose a generalization for color images based on the above method. Our proposal consists of applying the gray four-source photometric algorithm to each color channel (RGB sensor values) in a separate way, deriving three albedo values and three normal vectors for each pixel. Because all three channels contain information about the same surface, the difference between them is only the color information and so the three albedo components can be used as RGB components. The normal vectors can be averaged to end up with only one normal vector in each pixel.

The summary of the four-source color photometric-stereo algorithm is the following:

1. Starting with four intensity values per pixel we construct the four possible intensity vectors of size 3×1 and combine these four values into groups of three:

$$\mathbf{I}_j^x = \rho_j^x [\mathbf{L}]_j \mathbf{N}_j^x, \quad (7)$$

where $j = 1, 2, 3, 4$ and $x = 1, \dots, n$, n being the total number of pixels.

2. We apply Eq. (6) to each of the four intensity vectors to recover four albedos and normal vectors:

$$[\mathbf{L}]_j^{-1} \mathbf{I}_j^x = \rho_j^x \mathbf{N}_j^x, \quad (8)$$

3. We calculate the standard deviation $\sigma(\rho_j)$ of the four albedo values and set a threshold, t , and:
 - a. if $\sigma(\rho_j) \leq t$, average the four albedo values and the components of the normal vector to get just one albedo and normal vector in this pixel;
 - b. if $\sigma(\rho_j) > t$, look for the maximum albedo value and discard the combination of intensities that has provided it. We then average the other three albedos and components of the normal vector to get just one albedo and normal vector in this pixel.

4. It only remains to repeat all the above steps for each color channel (R, G, and B) and use the albedo value recovered for the first channel as ‘‘R albedo’’ and so on. The final normal vector in each pixel will be the average of the normal vectors recovered for all three channels:

$$\mathbf{N}^x = \frac{1}{3} \sum_i \mathbf{N}_i^x, \quad (9)$$

where i is the number of channels, $i = 1, 2, 3$.

C. Spectral Reflectance Estimation via Linear Pseudoinverse Method

The goal of spectral imaging is to recover spectral radiance or reflectance for each pixel of any scene in question. In the so-called linear pseudoinverse method [4], given a set of training spectra, S , (which can be the spectral radiance of reflectance) and the

corresponding set of experimental camera responses, q , a recovery transformation matrix, \mathbf{D} , is defined by:

$$\mathbf{D} = \mathbf{S}\mathbf{q}^+ \quad (10)$$

where \mathbf{q}^+ is the pseudoinverse of \mathbf{q} and, if \mathbf{q} has full rank, then $\mathbf{q}^+ = (\mathbf{q}^T\mathbf{q})^{-1}\mathbf{q}^T$. An estimation set of spectra, S_1 , may then be obtained from the corresponding set of camera responses, \mathbf{q}_1 , by applying the transformation, \mathbf{D} , i.e.,

$$\mathbf{S}_1 = \mathbf{D}\mathbf{q}_1 \quad (11)$$

Thus, once we have recovered reflectance we can now deal with more parameters because it is possible to model the albedo in each pixel as

$$\rho = \sum E(\lambda)S(\lambda)Q(\lambda), \quad (12)$$

where $E(\lambda)$ is the spectrum of the illuminant, $S(\lambda)$ is the reflectance of the pixel, and $Q(\lambda)$ represents the camera's sensitivity for each channel. Once we have the reflectance of each pixel of the sample, we need only the SPD of one illuminant and the camera's sensitivity in each channel to simulate the sample, not only under any direction of illumination but also under any illuminant, too.

This approach can be sensitive to the size and composition of the training set of reflectances [34,35]. To choose an appropriate training set, a supervised training sample selection is applied pixel-wise for a given image. As a training set we used all 1269 samples from the Munsell Color Chart [36] to take advantage of its classification of colors in terms of hue groups. In the 3D space determined by the R, G, B channels the distribution of the coordinates of the samples of each hue group more or less forms a plane (Fig. 2).

Hence, a reduction in the dimensions of the distribution of each hue group to a plane would provide us with an easy way of representing them in the RGB




Fig. 2. Example of the distribution of one Munsell hue group's sample in the RGB space in digital counts as captured by our Retiga camera.

space. PCA is the most usual approach for reducing multidimensional datasets to a lower number of dimensions. We have applied PCA to all the sets of samples contained in each Munsell page in the RGB space in order to reduce their 3D distribution to a plane. The choice of the supervised training sample involves calculating the Euclidean distance in the RGB space between each pixel of the image and these planes. The closest plane to each pixel tells us the hue group that we must use to recover the reflectance of this particular pixel.

The summary of the linear pseudoinverse algorithm is:

1. First we calculate the transformation matrix, \mathbf{D} , associated to each of the ten Munsell Color Chart hue groups.
2. On the basis of the albedo values recovered with the four-source photometric-stereo method in each pixel of the image we calculate the distance between its components in the RGB space and the ten planes. The closest plane to the albedo coordinates of each pixel will tell us the transformation matrix to use.
3. Once we know the closest plane we use the corresponding transformation matrix to recover reflectance using Eq. (11).

D. Acquisition and Dataset

Images were captured with a Retiga 1300 CCD camera (12 bit intensity range per channel) from QImaging, Canada, with a LINOS MeVis-C lens with a fixed 5.6 aperture and focal length of 25 mm. The camera was incorporated into a setup like the one shown in Fig. 3, in which the relative position between the camera and the sample is fixed. This setup allows the camera-sample couple to rotate together around the z axis. In this way, by fixing one lamp, it is possible to capture the sample under different illumination directions simply by rotating the sample and the camera.

Three different commercial lamps were used as illumination sources throughout the experiments, two fluorescent and the third incandescent. Their spectral power distributions are shown in Fig. 4.

Our sample set contained 24 samples of frontage cover with two different kinds of texture; an example is shown in Fig. 5. The first ones were smoother, terrainlike, without big differences over the sample, while the second ones were more abrupt, a mixture of terrain and little stones. Samples were captured under eight different illumination directions: 0° , 90° , 180° , and 270° , which is the *training set*; and 45° , 135° , 225° , and 315° , which is the *test set*. All the samples were captured under the eight directions of illumination and under the three illumination sources at an image size of 100×100 pixels. The Munsell samples used as a training set in the pseudoinverse method were captured with the same CCD camera under one of the fluorescent lamps.




Fig. 3. Experimental setup. Camera and sample can go around together, maintaining the relative position between them. By fixing the position of a source and making the camera–sample set move, it is possible to obtain any desired direction of illumination.

3. Results and Discussion

To quantify the quality of the results we used the following metrics: the RGB error, the CIE Lab color difference, ΔE^*ab , and the angular error (AE). The RGB error values [37] were calculated pixel by pixel via the differences between the values of the R, G, and B coordinates of the compared images using the expression

$$\text{RGBError}_x = \sqrt{\frac{1}{3}(\Delta R_x^2 + \Delta G_x^2 + \Delta B_x^2)}, \quad (13)$$

where ΔR_x , ΔG_x , and ΔB_x are the differences at any pixel x . CIE Lab color differences were obtained employing as reference the tristimulus values calculated from RGB values of a white patch captured under the same illuminant as the samples. The angular error measurement was defined as




Fig. 4. Normalized SPDs of the illumination sources.




Fig. 5. Example of the two different kinds of textures of the samples used in this work. The first image represents an example of the smooth, terrainlike samples, and the second one is an example of the more abrupt texture.

$$\text{AE}_x = \cos^{-1}(\rho_o \cdot \rho_e), \quad (14)$$

where ρ_o and ρ_e are the RGB color components of the original and the algorithm's estimation of the images at pixel x , respectively.

A. Albedo Recovery and Surface Rendering

We used half of the images captured with one fluorescent lamp to recover both the albedo and normal vectors using the color photometric algorithm, and went on to employ this information to simulate the samples under the eight illumination directions to check the performance of our algorithm. Thus, only the images captured with one of the fluorescent lamps was used here and, from the set of eight different illumination directions, four were used in the four-source photometric-stereo process (the training set of samples) to recover the albedo and normal vectors. To check the accuracy of this recovery a simulation set was made using Eq. (3). To simulate the test images we changed the illumination vector, L , to get the eight illumination directions mentioned above. Hence the 24 samples were simulated under the eight illumination directions, making a total number of 192 simulated samples.

These simulated images were compared to the real ones using all three metrics mentioned above pixel by pixel over the images. First, we analyzed separately the results obtained for the images used to recover the albedo and normal vectors (training sample set at 0° , 90° , 180° , and 270°) (see Table 1). The effects of the direction of incident light and the metric were tested by a repeated-measures analysis of the variance of two factors: the direction used to capture the images (0° , 90° , 180° , and 270°) and the metric used to evaluate the recovery quality (RGBError, ΔE^*ab , and AE). The results from the training set of samples suggest very good color recovery for the rendering algorithm. The average RGB error is 78, which means only a 1.9% total error, and decreases to 57 (1.39% total error) when the median is taken into account. These values correspond to AE values of 0.74 and 0.53, respectively. When analyzing RGB error values we must remember that our image capturing device is a 12 bit CCD camera that allows RGB values ranging from 0 to 4095. The colorimetric

Table 1. Pixel-by-Pixel Statistics Obtained When Original and Rendered Images From the Albedo and Normal Recovery Step are Compared

	Mean	Standard Deviation	90 Percentile
Training set of samples			
RGBerror	78	70	165
ΔE^*ab	2.9	2.3	5.8
AE	0.74	0.74	1.51
Test set of samples			
RGBerror	87	89	177
ΔE^*ab	3.2	2.8	6.1
AE	0.76	0.75	1.53
Averaged results			
RGBerror	82	80	171
ΔE^*ab	3.1	2.6	6.0
AE	0.75	0.75	1.52

results are satisfactory, too, with a mean ΔE^*ab of 2.9 and a median of only 2.3. The values are under 5 ΔE^*ab units, which is acceptable enough. Although 1 ΔE^*ab is considered to be a just-noticeable color difference, the accepted ΔE^*ab values in the literature depend upon the applications and the field of interest. Thus the mean ΔE^*ab color-difference tolerances in printing applications and television are usually around 5 to 6, and even above these limits for industry [37]. Statistical results show that the performance depends upon the direction selected and shows significant differences for all the metrics ($p \ll 0.05$). In order to complete this analysis, we tested our algorithm using samples that were not included in the training set. So, images with directions of illumination of 45°, 135°, 225°, and 315° were simulated using the recovered normal vectors and albedo from the training set, and captured, too, to check results. In this case we found significant differences for the geometry of illumination and the metric used ($p \ll 0.05$), but color recovery was also very good, with a mean RGBerror of 82 (2.1% total error) and mean ΔE^*ab of 3.2. Table 1 resumes all of these results and also includes the 90 percentile values.

The visual significance of the RGB color error for a rough surface not included in the training and set at 45° incident illumination is illustrated in Fig. 6, in which the histogram of the RGBerror differences obtained over the whole image (10,000 pixels) is shown first. In the second row of the figure, the first image is the original one and the second one is a color-scale image showing the distribution of the RGB differences in the image. We also present an example of the median results, obtaining very good color image rendering. Furthermore, it can be seen how some of the image pixels belonging to shadows in the image lead to considerable color errors. This reflects one of the limitations of the method, which does not take any shadows cast into account.

B. Spectral Reflectance Recovery and Lack of Homogeneity

We selected materials with heterogeneous surface textures. This becomes a problem when we want to




Fig. 6. (Color online) Example of simulated sample not included in the training set. First line, histogram of RGB differences. Second line, original sample and color-scale image showing the distribution of RGB differences in the images.

recover spectral information with a device such as a spectroradiometer because this kind of instrument uses a spot with a fixed area and averages the incoming radiation from this area. If we have a textured sample, where reflectance varies from one point to another, it is possible to find areas with different reflectance values within the spot. This is the reason why pixel-by-pixel recovery is very important with this kind of sample.

With the aim of verifying this fact, we have included one experiment where measures from a spectroradiometer and from our algorithm are compared. It is not possible to exactly relate the area measured by the spectroradiometer with the pixels chosen in an image to recover reflectance, but our goal here is to show that, if reflectance inside the spot of a spectroradiometer varies, the reflectance given for this area cannot be appropriate for all pixels inside of it.

Figure 7 shows an example of spectral recovery as compared to spectroradiometric measurements on small surface areas. The image shows one of the samples where the reflectance of one pixel within the three white circles has been chosen. These pixels belong to three different kinds of stone, so their reflectances are different, too. The first graph shows reflectances recovered with our algorithm from albedo information and the second one shows reflectances measured with a PR650 spectroradiometer. The area measured with the PR650 was centered on the pixel whose reflectances were recovered with our algorithm. As we said, it is not possible to measure exactly the same pixel because of the size of the




Fig. 7. (Color online) Comparison of reflectance recovered with our algorithm and measured with a spectroradiometer. (a) Original image with the measured areas marked. (b) Reflectances recovered with our algorithm. (c) Reflectances measured with a PR650 spectroradiometer.

spectroradiometer's spot, which shows the limitations of this kind of device when pixel-by-pixel information is required.

Measured reflectances have a more homogeneous shape due to the average of the incoming radiance from the area inside the spot of the instrument, but in our recovered reflectances a change in shape corresponding to the different materials included in the spectroradiometer's spot can be seen.

C. Albedo Recovery and Image Rendering under Spectrally Different Lights

The results described above suggest that a spectral-based photometric-stereo algorithm is sufficient to recover the albedo satisfactorily and then to render color. This has also been found in previous results for different light intensities [33], but what would happen if the source of illumination changed in its spectral content? In this section we show the preliminary results deriving from our method for different spectral lights and incidence geometry. We have started from the estimated albedo, as shown in Subsection 2.C, and recovered the spectral reflectances at a pixel using Eqs. (9) and (3) to simulate the samples under the three spectrally different light sources and the eight illumination directions. We used the sensitivities of the Retiga camera shown in Fig. 8. Thus the 24 samples were simulated under three lamps and eight illumination directions, which adds up to a total of 576 test images.




Fig. 8. Spectral sensitivities of the RGB digital camera.

Figures 9–11 contain the same elements as Fig. 6 and show examples of the quality of image recovery when the spectral profile of the illumination under which the photometric stereo runs (i.e., albedo estimation) is different from that under which the image is rendered (i.e., spectral reflectance estimation from linear pseudoinverse). As might be expected, the RGB errors are now larger than those found previously, with averaged values of around 435 (i.e., 10% total error). Table 2 contains separate summaries of the results of the training and the test sets of samples. The first three rows set out the results for each of the lamps separately and the last row shows the average values for all of them. For the training set of samples we found average RGB error values of 435 (10.6% total error). For the test set of samples the color errors were almost the same, with an average RGB error of 436. When only the direction of





Fig. 9. (Color online) Example of simulation made with the albedo calculated using the reflectance obtained by the pseudoinverse method and the SPD of the Digilite lamp, showing the same information as Fig. 6.





Fig. 10. (Color online) Example of simulation made with the albedo calculated using the reflectance obtained by the pseudoinverse method and the SPD of the incandescent lamp, showing the same information as Fig. 6.





Fig. 11. (Color online) Example of simulation made with the albedo calculated using the reflectance obtained by the pseudoinverse method and the SPD of the Trilite lamp, showing the same information as Fig. 6.

illumination was changed, we obtained a clear difference between results from test and training sets. But, when changing the SPD of the illuminant also, those differences between training and test sets results are lower, finding even a case where results are slightly better for the test set (mean values of RGBError and AE for Trilite).

Besides, the ΔE^*ab values now exceed what might be considered as being of acceptable colorimetric quality, showing mean ΔE^*ab values of 14.8 (± 9.9 standard deviation). Although those values are not satisfactory according to usual tolerances, we should consider that there is no appropriate metric to use with samples like ours, i.e., textured samples, so we do not know for sure that the obtained color differences really describe the observed differences.

Could it be that the CIELab is not the appropriate metric for these comparisons? The final answer to this question would be suitable after future analysis and results.

The AE values corroborate these results with mean values of around 3.46, which lead to an acceptable image rendering, independently of the spectral content of the light. The cosine values associated to the AE are clearly above 0.9981, which is usually considered as being a colorimetrically accurate color estimations when Eq. (13) is interpreted in terms of a goodness-of-fit coefficient [3]. Finally, the results for both the test and training sets are shown in Table 3.

Results in this section shows that this method can be used to render images, changing both the orienta-

Table 2. Pixel-by-Pixel Statistics Obtained When Original and Rendered Images From the Spectral Recovery Step are Compared

		Training Set			Test Set		
		Mean	Standard deviation	90 percentile	Mean	Standard deviation	90 percentile
Digilite	RGBError	323	293	719	328	296	719
	ΔE^*ab	11.0	9.0	22.7	11.2	9.2	23.8
	AE	2.10	1.37	3.96	2.33	1.81	4.22
Incand	RGBError	414	277	780	419	282	791
	ΔE^*ab	15.7	9.1	28.0	15.9	9.4	28.6
	AE	4.01	2.79	7.73	4.15	3.06	8.42
Trilite	RGBError	567	324	994	561	322	979
	ΔE^*ab	17.8	10.2	31.1	22.0	16.0	41.1
	AE	4.26	2.87	8.48	4.16	2.62	7.76
Average	RGBError	435	315	855	436	315	852
	ΔE^*ab	19.4	15.1	37.3	14.9	9.9	38.0
	AE	3.46	2.63	6.79	3.55	2.69	7.20

Table 3. Overall Statistic Results From the Spectral Recovery Step

		Mean	Standard Deviation	90 Percentile
Digilite	RGBerror	326	295	719
	ΔE^*_{ab}	11.1	9.1	23.2
	AE	2.21	1.61	4.08
Incand	RGBerror	417	280	785
	ΔE^*_{ab}	15.8	9.2	38.3
	AE	4.08	2.93	8.07
Trilite	RGBerror	564	323	986
	ΔE^*_{ab}	17.7	10.1	30.9
	AE	4.21	2.75	8.04
Average	RGBerror	435	315	854
	ΔE^*_{ab}	14.9	9.9	28.1
	AE	3.50	2.66	7.01

tion and the spectral composition of the illumination source, but with low colorimetric accuracy.

4. Conclusions

The image reconstruction method described here is a photometric-based algorithm that directly relates the camera sensor outputs to albedo values. Using an RGB digital camera, the algorithm recovers the surface spectral reflectance of objects to allow image rendering even for rough, colored surfaces under different lighting conditions. Although photometric-stereo methods have been applied successfully for the 3D shape of objects, few results are available concerning what might happen when the lighting conditions vary not only in intensity but also in their spectral content. This is the case, for example, in art galleries and museums, where paintings and other objects must be shown under fluorescent or incandescent light. Thus we have extended the four-source photometric-stereo approach to color images to first recover albedo and normal vectors in each pixel of a color image. This information allows us to simulate the captured samples under different illumination directions. Subsequently a linear pseudoinverse method has been used to recover spectral reflectances in single pixels from the albedo values. The linear pseudoinverse method can be combined with a supervised algorithm to choose the most appropriate set of training samples. By adding this spectral reflectance information to the surface shape recovered beforehand, it is possible to simulate a sample under any illumination direction and different spectral source.

Spectral reflectance obtained by the linear pseudoinverse method allows us to recover spectral reflectances pixel by pixel from RGB scenes but, by using albedo instead of RGB information, it is possible to avoid any shadows or highlights that might falsify results. This information can be used to make computational simulations of real samples under different conditions. Using albedo and normal vector data it is possible to simulate a sample under different directions of illumination and, by adding the spectral reflectance information, the source of illumination can be changed as well. Therefore, starting with four

images captured under different illumination directions and a determined source of illumination we can simulate a sample under any direction and source of illumination just by knowing its SPD. In the former case we found good results with all the metrics. The latter case was not so accurate and we will introduce several improvements in the algorithm for future works.

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 their colleague A. L. Tate for revising their English text.

References

1. B. A. Wandell, "The synthesis and analysis of color images," *IEEE Trans. Pattern Anal. Machine Intell.* **pami-9**, 2–13 (1987).
2. J. Hardeberg, F. Schmitt, H. Brettel, "Multispectral color image capture using a liquid crystal tunable filter," *Opt. Eng.* **41**, 2532–2548 (2002).
3. F. H. Imai and R. Berns, "Spectral estimation using trichromatic digital cameras," in *International Symposium on Multispectral Imaging and Color Reproduction for Digital Archives* (Society of Multispectral Imaging of Japan, 1999), pp. 42–49.
4. N. Shimano, "Evaluation of a multispectral image acquisition system aimed at reconstruction of spectral reflectances," *Opt. Eng.* **44**, 107005 (2005).
5. 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 and colored filters," *Color Res. Appl.* **32**, 352–360 (2007).
6. 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).
7. 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).
8. C. Plata, J. L. Nieves, and J. Romero, "Combining spectral and photometric stereo techniques for reflectance estimation using an RGB digital camera," in *Color in Graphics, Imaging and Vision (CGIV) '08 and Multispectral Colour Science (MCS) '08 Final Program and Proceedings* (International Science and Technology, 2008), pp. 516–518.
9. R. J. Woodham, "Reflectance map techniques for analyzing surface defects in metal castings," Technical Report AI-TR-457 (MIT, Artificial Intelligence Laboratory, 1987).
10. R. J. Woodham, "Photometric method for determining surface orientation from multiple images," *Opt. Eng.* **19**, 139–144 (1980).
11. G. McGunnigle and M. Chantler, "Rough surface description using photometric stereo," *Meas. Sci. Technol.* **14**, 699–709 (2003).
12. K. Ikeuchi, "Determining surface orientations of specular surfaces by using the photometric stereo method," *IEEE Trans. Pattern Anal. Machine Intell.* **pami-3**, 661–669 (1981).
13. E. Coleman Jr. and R. Jain, "Obtaining 3-Dimensional shape of textured and specular surfaces using four-source photometry," *Comp. Graph. Image Process.* **18**, 309–328 (1982).
14. S. Barsky and M. Petrou, "The 4-source photometric stereo technique for 3-dimensional surfaces in the presence of highlights and shadows," *IEEE Trans. Pattern Anal. Machine Intell.* **25**, 1239–1252 (2003).

15. H. Tagare and R. de Figueiredo, "A theory of photometric stereo for a class of diffuse non-Lambertian surfaces," *IEEE Trans. Pattern Anal. Machine Intell.* **13**, 133–152 (1991).
16. B. Kim and P. Burguer, "Depth and shape from shading using the photometric stereo method," *Comp. Vis. Graph. Image Process.* **54**, 416–427 (1991).
17. M. Oren and S. K. Nayar, "Generalization of Lambert's reflectance model," *Comp. Graph.* **28**, 239–246 (1994).
18. C. Hernández, G. Vogiatzis, and R. Cipolla, "Multiview photometric stereo," *IEEE Trans. Pattern Anal. Machine Intell.* **30**, 548–554 (2008).
19. A. Spence and M. Chantler, "On capturing 3D isotropic surface texture using uncalibrated photometric stereo," in *3rd International Workshop on Texture Analysis and Synthesis* (TextureLab, 2003), pp. 83–88.
20. M. S. Drew, "Photometric stereo without multiple images," *Proc. SPIE* **3016**, 369–380 (1997).
21. T.-P. Wu, K.-L. Tang, and T.-T. Wong, "Dense photometric stereo: a Markov random field approach," *IEEE Trans. Pattern Anal. Machine Intell.* **28**, 1830–1846 (2006).
22. G. Healey and L. Wang, "Three-dimensional surface segmentation using multicolored illumination," *Opt. Eng.* **37**, 1553–1562 (1998).
23. C. Hernández, G. Vogiatzis, G. J. Brostow, B. Stengar, and R. Cipolla, "Non-rigid photometric stereo with colored lights," in *Proceedings of IEEE International Conference on Computer Vision* (IEEE, 2007), pp. 1–8.
24. M. Chandraker, S. Agarwal, and D. Kriegman, "ShadowCuts: photometric stereo with shadows," in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, 2007), pp. 1–8.
25. C. Hernández, Vogiatzis, and R. Cipolla, "Shadows in three-source photometric stereo," in *Proceedings of IEEE European Conference on Computer Vision* (IEEE, 2008), pp. 1–14.
26. K. Schl and O. Witting, "Photometric stereo for non-Lambertian surfaces using color information," in *Proceedings of 5th International Conference on Computer Analysis of Images and Patterns* (Springer, 1993), pp. 444–451.
27. L. L. Kontsevich, A. P. Petrov, and I. S. Vergelskaya, "Reconstruction of shape from shading in color images," *J. Opt. Soc. Am. A* **11**, 1047–1052 (1994).
28. M. S. Drew, "Shape from color," Technical Report CSS/LCCR TR 92-07 (Simon Fraser University School of Computing Science, 1992).
29. M. S. Drew, "Optimization approach to dichromatic images," *J. Math. Imaging Vis.* **3**, 187 (1993).
30. P. H. Christensen and L. G. Shapiro, "Three dimensional shape from color photometric stereo," *Int. J. Comput. Vis.* **13**, 213–227 (1994).
31. S. Barsky and M. Petrou, "Color photometric stereo: Simultaneous reconstruction of local gradient and color of rough textured surfaces," in *Proceedings of Eighth IEEE International Conference on Computer Vision* (IEEE, 2001), pp. 600–605.
32. B. Bringier, D. Helbert, and M. Khoudeir, "Photometric reconstruction of a dynamic textured surface from just one color image acquisition," *J. Opt. Soc. Am.* **25**, 566–574 (2008).
33. C.-Yen Chen, R. Klette, and C.-F. Chen, "Recovery of colored surface reflectances using the photometric stereo method," in *Proceedings of International Conference on Information Systems* (Association for Information Systems, 2003), pp. 969–974.
34. M. de Lasarte, J. Pujol, M. Arjona, and M. Vilaseca, "Influence of the size of the training set on color measurements performed using a multispectral imaging system," in *Color in Graphics, Imaging and Vision (CGIV) '08 and Multispectral Colour Science (MCS) '08 Final Program and Proceedings* (International Science and Technology, 2008), pp. 437–440.
35. C. Plata, E. M. Valero, J. L. Nieves, and J. Romero, "Supervised training sample selection for the estimation of spectral reflectance using a RGB camera," in *Color in Graphics, Imaging and Vision (CGIV) '08 and Multispectral Colour Science (MCS) '08 Final Program and Proceedings* (International Science and Technology, 2008), pp. 519–522.
36. *Munsell Book of Color—Matte Finish Collection* (Munsell Color, 1976).
37. Corbalan, Millan, and Yzuel, "Color measurement in standard CIE Lab coordinates using a 3CCD camera: correction for the influence of the light source," *Opt. Eng.* **39**, 1470–1476 (2000).