Recovering the Spectral Distribution of the Illumination from Spectral Data by Highlight Analysis

Harro Stokman and Theo Gevers

Intelligent Sensory Information Systems,
University of Amsterdam,
Kruislaan 403, 1091 SJ Amsterdam, The Netherlands

ABSTRACT

The problem of color constancy for discounting illumination color to obtain the apparent color of the object has been the topic of much research in computer vision. By assuming the neutral interface reflection and dichromatic reflection with highlights (i.e. highlights have the same color as the illuminant) various methods have been proposed aiming at recovering the illuminant color from color highlight analysis. In general, these methods are based on three color stimuli to approximate color. In this contribution, we estimate the spectral distribution from surface reflection using spectral information obtained by a spectrograph. The imaging spectrograph provides a spectral range at each pixel covering the visible wavelength range. Our method differ from existing methods by using a robust clustering technique to obtain the body and surface components in a multi-spectral space. These components determine the direction of the illumination spectral color. Then, we recover the illumination spectral power distribution by using principal component analysis for all wavelengths. To obtain the most reliable estimate of the spectral power distribution of the illuminant, all possible combinations of wavelengths are used to generate the optimal averaged estimation of the spectral power distribution of the scene illuminant. Our method is restricted to images containing a substantial amount of body reflection and highlights.

Keywords: Color constancy, dichromatic reflection model, segmentation, estimation of spectral distribution of illuminant

1. INTRODUCTION

The problem of color constancy for discounting or recovering the illumination color to obtain object color has been the topic of much research in computer vision. Our goal is to recover the spectral distribution of the illuminant from highlight analysis in multi-spectral images based on the dichromatic reflection model. Various methods have been proposed for recovering the illuminant color from highlight analysis. These methods assume that the neutral interface reflection model holds i.e. highlights have the same color as the illumination source. Lee's[1] method is based on normalized color to recover the chromaticity of the scene illuminant. Healey[2] estimates spectral reflectance using highlights within a multiplicative constant. Further, Tominaga et al.[3] uses a vector analysis technique to estimate the illuminant. Our method differ from existing methods by using a robust clustering technique to obtain the body and surface components in a multi-spectral space. These components determine the direction of the illumination spectral color. Then, we recover the illumination spectral power distribution by using principal component analysis for all wavelengths. To obtain the most reliable estimate of the spectral power distribution of the illuminant, all possible combinations of wavelengths are used to generate the optimal averaged estimation of the spectral power distribution of the scene illuminant. Our method is restricted to images containing a substantial amount of body reflection and highlights.

The paper is organized as follows. In section 2 we give a mathematical formulation of the dichromatic reflection model. In section 3 we describe the underlying principle of our method for estimation of the color of the illuminant. In section 4 we present our experimental results. The performance is discussed in section 5.

2. DICHROMATIC REFLECTION MODEL

Consider an image of an infinitesimal surface patch. Using N narrow-band filters with spectral sensitivities given by $f_1(\lambda)...f_N(\lambda)$ to obtain an image of the surface patch illuminated by a spectral power distribution of the incident light denoted by $e(\lambda)$, the measured sensor values will be given by Shafer[4]:

$$C_n = m_b(\vec{n}, \vec{s}) \int_\lambda f_n(\lambda)e(\lambda)c_b(\lambda)d\lambda + m_s(\vec{n}, \vec{s}, \vec{v}) \int_\lambda f_n(\lambda)e(\lambda)c_s(\lambda)d\lambda \qquad (1)$$

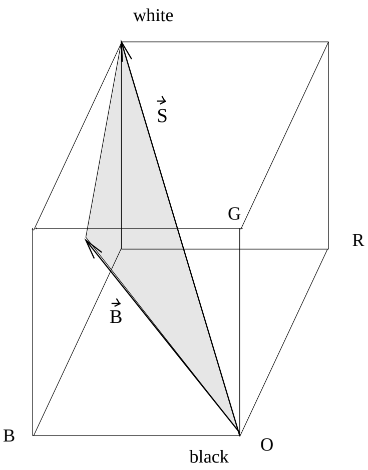

Figure 1. Color cluster vector $\vec{B}$ of a matte surface in RGB, and color cluster vector $\vec{S}$ of a shiny surface in RGB.

for C_n giving the nth sensor response. Further, $c_b(\lambda)$ and $c_s(\lambda)$ are the albedo and Fresnel reflectance respectively. λ denotes the wavelength, $\vec{n}$ is the surface patch normal, $\vec{s}$ is the direction of the illumination source, and $\vec{v}$ is the direction of the viewer. Geometric terms m_b and m_s denote the geometric dependencies on the body and surface reflection respectively. The body reflection occurs for light that crosses the object's interface. Once the light enters the pigment colorant layer, scattering of the light among the pigment particles in this layer occurs. The light is observed only after it crosses the surface boundary again into the air. Some light is reflected at the interface between the object's surface and the air. This type of reflection is mirror like; light reflected in this way from a point source can be seen only over a narrow range of viewing geometries. This light is said to be surface reflection.

At the moment an image is captured, the spectral sensitivities of the color filters, the spectral power distribution of the illuminant, and the spectral transmission of the object are constant. The term $\int_\lambda f_n(\lambda)e(\lambda)c_b(\lambda)d\lambda$ can therefore be abbreviated as the constant k_n, and the term $m_s(\vec{n},\vec{s},\vec{v})\int_\lambda f_n(\lambda)e(\lambda)c_s(\lambda)d\lambda$ as constant l_n. This allows Eq. 1 to be written as

$$\begin{pmatrix} C_1 \\ C_2 \\ . \\ . \\ C_N \end{pmatrix} = m_b(\vec{n},\vec{s}) \begin{pmatrix} k_1 \\ k_2 \\ . \\ . \\ k_N \end{pmatrix} + m_s(\vec{n},\vec{s},\vec{v}) \begin{pmatrix} l_1 \\ l_2 \\ . \\ . \\ l_N \end{pmatrix} \tag{2}$$

A homogeneously colored matte object with constant $k_1, k_2, .., k_N$ but with varying surface normal $\vec{n}$, and thus $0 \leq m_b(\vec{n},\vec{s}) \leq 1$ and $m_s(\vec{n},\vec{s},\vec{v}) = 0$, will draw a half-ray, emanating from the origin, in sensor space. A homogeneously colored shiny object with varying surface normal $\vec{n}$, and thus $0 \leq m_b(\vec{n},\vec{s}) + m_s(\vec{n},\vec{s},\vec{v}) \leq 1$, will draw two half-rays in sensor space. Assuming that the object is illuminated by a single light source, all pixels that belong to the homogeneously colored shiny object are located in the N- dimensional plane spanned by these two vectors. This is visualized in fig. 1 for an RGB camera where $N = 3$. The observed colors of a matte, dull surface with fixed k_n can be represented by the color vector $\vec{B}$ where the direction of $\vec{B}$ is based on k_n and its extent by $m_b(\vec{n},\vec{s})$. A uniformly painted surface may give rise to a broad variance of sensor values due to the varying circumstances induced by the image-forming process such as a change in object orientation, illumination intensity and position. The observed colors of a shiny surface with fixed k_n and l_n is composed by the color vectors $\vec{B}$ and $\vec{S}$. Here, the direction of $\vec{S}$ is based on l_n. In fig. 1, the direction of $\vec{S}$ is equal to the achromatic axis and thus, for this example, the color of the illuminant is equal for all wavelengths. The extend of $\vec{S}$ is based on $m_s(\vec{n},\vec{s},\vec{v})$. Thus, all pixels that belong to the homogeneously colored shiny object are located in the triangle spanned by $\vec{B}$ and $\vec{S}$.

3. METHOD FOR ESTIMATION OF THE SPECTRAL DISTRIBUTION OF THE ILLUMINANT FROM BODY AND SURFACE REFLECTION

Consider two differently colored shiny objects made from the same material that are illuminated by a single light source. According to the dichromatic reflection model, discussed in section 2, the body reflection component of the reflected light of the two objects will draw two half-rays in the N-dimensional sensor space. The half-rays emanate from the origin, and the orientation depends on the spectral transmittance of the narrow-band filters, the illuminant and the objects. Since the objects have different colors, the orientation of the half-rays is different. Because the objects are illuminated by a single light source, each object draws a single line in the sensor space. Both objects are shiny as well. According to the dichromatic reflection model, the surface reflection component of the objects will draw a half-ray in the N-dimensional sensor space. This half-ray also emanates from the origin, but its orientation depends on the spectral transmittance of the narrow-band filters, the illuminant and the Fresnel reflectance of the objects. As both objects are made out of the same material, and thus have the same Fresnel reflectance, and because both objects are illuminated by the same light source, the orientation of the surface reflection half-ray is identical for both objects, in contrast to the orientation of the body reflection half-rays.

According to the dichromatic reflection model, for a single shiny object the reflected light is a mixture of body and surface reflection. The exact amounts of these components depends on the orientation of the surface normal, the direction of the light source and the viewing direction of the camera. As a result, the colors observed for a single homogeneously colored, shiny object are located on a hyper-plane in N-dimensional sensor space. The plane is bounded by the body and surface reflection half-rays. To see this in three dimensions, for example for the case of an RGB camera, the observed colors are located on a triangular plane, as shown in fig. 1.

Our goal is to detect the spectral distribution of the illuminant. Thus, the aim is to estimate the orientation of the surface reflection half-ray, which depends only on the spectral distribution of the illuminant, and not on the color of the objects itself. Given two shiny objects, which draw two hyper-planes in sensor space, the intersection of these planes corresponds to the orientation of this surface reflection component. Prior to computing this intersection, the orientation of the surface and reflection half-rays need to be retrieved for both objects, since these half-rays define the hyper-planes. In the next step, these planes are then intersected.

In order to extract the orientation of the surface and reflection half-rays, the N-dimensional image, containing the spectral data of the two shiny objects, needs to be partitioned in two regions. Each region then corresponds to the spectral data of one shiny, homogeneously colored object. This partitioning is achieved through the segmentation method described in.[5] In short, the method is based on k-means based clustering, with k the number of clusters. K-means clustering is based on the minimization of the sum of squared distances of N-dimensional pixels to their N-dimensional cluster means. First an initial partition is given, obtained by assigning each pixel to its closest cluster mean, on basis of the Euclidean distance measure, and new partitions are obtained by moving pixels from one cluster to another until the summed square distance is reduced to its minimum. According to the dichromatic reflection model, homogeneously colored matte objects draw half-rays in N-dimensional sensor space. The segmentation method exploits this knowledge by using lines as cluster model, instead of centroids ($k-lines$ clustering). The lines emanate from the origin in the N-dimensional sensor space. The distance measure is defined as the shortest distance from the N-dimensional pixel to the N-dimensional cluster line. For segmentation of shiny objects, which do not draw lines in sensor space but which draw planes, the method is not optimal. However, under the assumption that the body reflection component of the reflected light is more dominantly present then the surface reflection component, and thus under the assumption that the majority of pixels that belong to the shiny object will be located close to the body reflection half-ray, the method is still capable to partition the image in two regions.

Now, under the assumption that the N-dimensional image, that depicts two shiny objects, is correctly segmented into two regions by this k-lines cluster method, the next step is to estimate the spectral distribution of the illuminant. First, consider this task for 3-dimensional spectral data, for instance for the case of recovering the spectral distribution of the illuminant in RGB images. The image is partitioned in two regions, each region represents a homogeneously colored, shiny object. Only the 3-dimensional spectral data is considered, the spatial information of the image can be ignored. Principal component analysis now computes, separately for each segment, a new orthogonal basis in the 3-dimensional sensor space. The resulting eigenvectors are ordered, such that the first eigenvector corresponds to the maximal variance in the spectral data. As it is assumed that the body reflection component of the reflected light is more dominant than the surface reflection component, the first eigenvector will point in the direction of the body reflection. Since the object is shiny, the second eigenvector will point in the direction of the surface reflection.

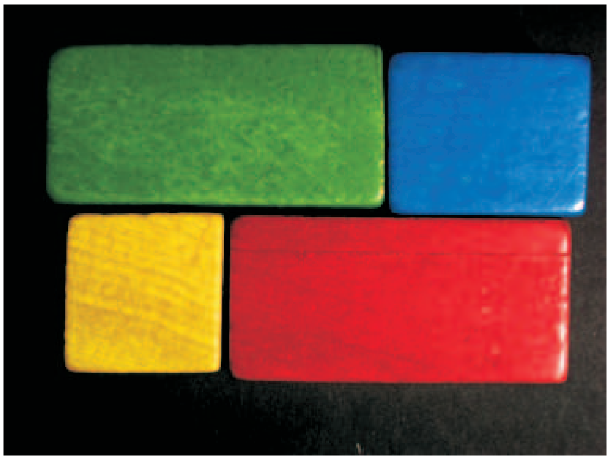

Figure 2. Experiments are conducted on homogeneously painted blocks of wood in the colors red, green, yellow and blue. This particular image is taken under halogen illumination. Due to the texture of the wood, the images will exhibit different amounts of body-reflection. Depending on the orientation of the illuminant, camera and objects, the painted blocks of wood will reflect highlights.

The third eigenvector is perpendicular to the first and second eigenvector, and has no physical meaning. This last eigenvector is of interest, solely because it has the property of being oriented perpendicular to the plane spanned by the body and surface reflection vector. For both segments of the 3-dimensional image, this third eigenvector is now computed. The outer product of these two third eigenvectors now corresponds to the 3-dimensional intersection of the two 3-dimensional planes. This way, the spectral distribution of the illuminant is estimated for 3-dimensional spectral data.

However, this technique can not be applied for N-dimensional data, if N is greater than 3, because the outer product is undefined if $N > 3$. Given the N-dimensional data, we can still take three different wavelengths, and estimate the spectral distribution of the illuminant for these wavelengths. Having two wavelengths fixed, and by consecutively using a third, different wavelength, it is possible to retrieve all N distributions. Because each of such computation yields a new, possibly slightly different solution for the two fixed wavelengths, the averaged estimation for these two fixed wavelengths is more robust than for the non-fixed wavelengths. Thus, to generate the most reliable estimate of the spectral distribution of the illuminant for all wavelengths, all possible combinations of wavelengths are used to obtain the best possible, averaged estimation of the spectral distribution of the illuminant.

4. EXPERIMENTS

Experiments are conducted on homogeneously painted blocks of wood in the colors red, green, yellow and blue. The blocks are shown in fig. 2. Due to the texture of the wood, the images will exhibit different amounts of body-reflection. Depending on the orientation of the illuminant, camera and objects, the painted blocks of wood will reflect highlights. In the experiments, the blocks are illuminated by differently colored light sources. The goal is to recover the spectral distribution of these light sources. The actual spectral distribution of these illuminants is obtained using a calibrated Spectralon white reference[6] with almost constant reflectance in the visible wavelength range of $0.990 \pm .002$. The image taken from this white reference, which is illuminated by differently colored lights, yields the spectral distribution of these illuminants.

4.1. Recovery of the Spectral Distribution of the Illuminant in RGB Images

RGB images are taken with a Sony 3CCD XC-003P color camera. Experiments are conducted on images containing two differently colored objects. The first experiment is shown in fig. 3. Two objects, with red and blue colors, are illuminated with a halogen light source, where the orientation of the objects, the camera and the illumination is such, that the objects show a significant amount of highlight. The figure also shows the actual and recovered spectral distribution of the illuminant in the bar graph. Although the transmission of the objects in the second (green)

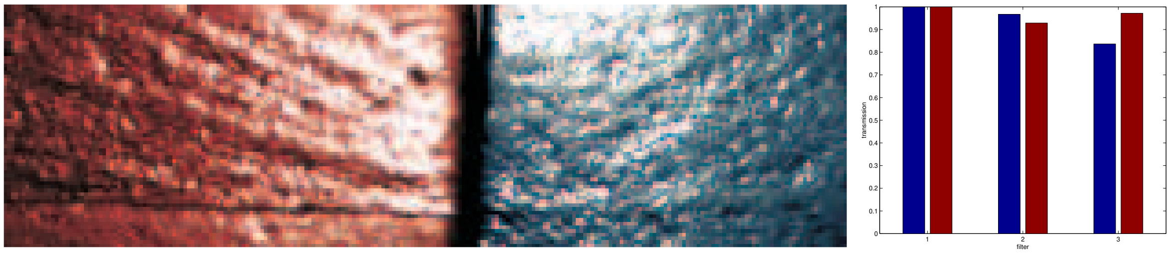

Figure 3. Left: Wooden blocks reflecting highlights. Right: Actual vs. recovered spectal distribution of the illuminant. The recovered spectral distribution of the illuminant is displayed in blue bars, the actual spectral distribution of the illuminant is displayed in red bars. The χ^2 error between the actual and recovered distribution is 0.020 for the three spectra. The experiment shows that, if both objects contain highlights, the color of the illuminant can be estimated.

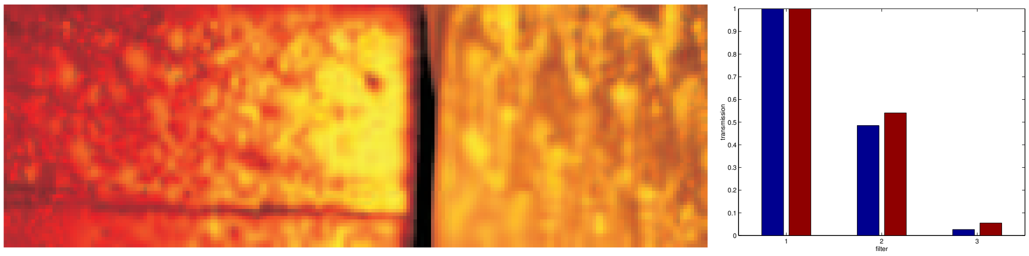

Figure 4. Left: Wooden blocks reflecting highlights. Right: Actual vs. recovered spectal distribution of illuminant. The recovered spectral distribution of the illuminant is displayed in blue bars, the actual spectral distribution of the illuminant is displayed in red bars. The χ^2 error between the actual and recovered distribution is 0.004 for the three spectra. The experiment shows that, if both objects contain highlights, the color of the illuminant can be estimated.

wavelength range is minimal, the spectral distribution of the illuminant is correctly recovered for this wavelength. The experiment shows that, if both objects contain highlights, the color of the illuminant can be estimated.

The second experiment is shown in fig. 4. Two objects, with red and yellow colors, are illuminated with a yellow light source. The objects show a significant amount of highlight. The figure also shows the actual and recovered spectral distribution of the illuminant in the bar graph. Although the transmission of the yellow object in the blue wavelength range is minimal, and although the transmission of the red object is minimal in both the blue and green wavelength range is minimal, the spectral distribution of the illuminant is correctly recovered. The experiment shows that, if both objects contain highlights, the color of the illuminant can be estimated.

The third experiment is shown in fig. 5. Two objects, with red and blue colors, are illuminated with a halogen light source. In contrast to the previous experiments, the objects do not exhibit highlights. The figure also shows the actual and recovered spectral distribution of the illuminant in the bar graph, which are unsimilar. The experiment shows that the spectral distribution of the illuminant can not be recovered if the objects do not reflect highlights.

4.2. Recovery of the Spectral Distribution of the Illuminant in Multi-Spectral Images

Multi-spectral images are recorded using an Imspector V7 spectrograph with the aid of a Jain CV-M300 gray value camera. Experiments are conducted on images containing two differently colored objects. The multi-spectral images contain spatial information on the horizontal axis, and spectral information on the vertical axis. The top of the vertical axis corresponds with transmittance of the sample of wavelengths of 400 nm, the bottom of the spectral axis corresponds to 700 nm. Due to the low sensitivity of the spectrograph and CCD camera to low wavelengths, the observed transmittance in the 400 to 450 nm range is low compared to higher wavelengths.

The fourth experiment is shown in fig. 6. Two objects, with blue and green colors, are illuminated with a halogen light source. The objects are are positioned in such a way that the blocks reflect highlights. The occurrence of these

Figure 5. Left: Wooden blocks reflecting highlights. Right: Actual vs. recovered spectal distribution of illuminant. The recovered spectral distribution of the illuminant is displayed in blue bars, the actual spectral distribution of the illuminant is displayed in red bars. The χ^2 error between the actual and recovered distribution is 0.238 for the three spectra. The experiment shows that the spectral distribution of the illuminant can not be recovered if the objects do not reflect highlights.

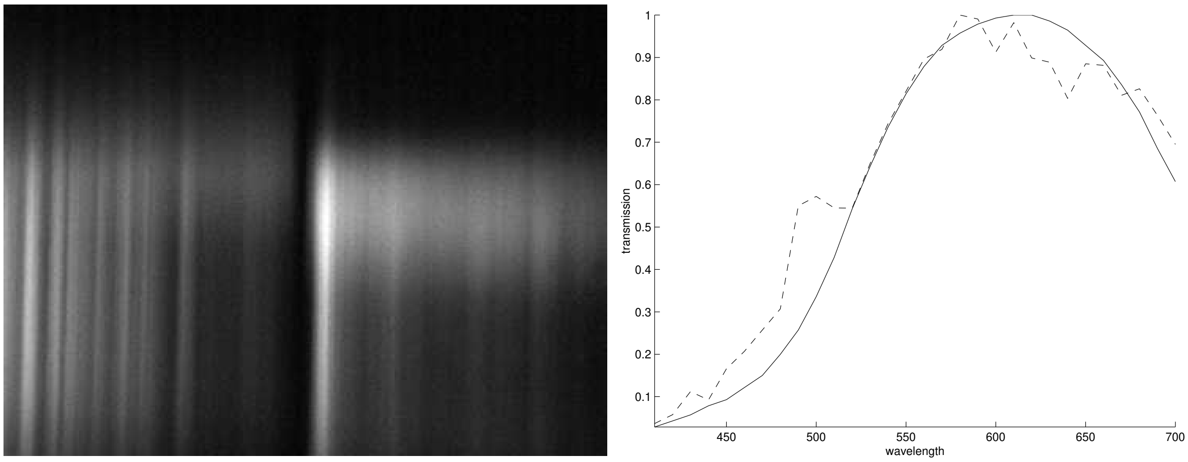

Figure 6. Left: Multi-spectral images of blue and green colored wooden blocks reflecting highlights. Right: Actual vs. recovered spectal distribution of illuminant. The recovered spectral distribution of the illuminant is displayed as a dashed line, the actual spectral distribution of the illuminant is displayed as a solid line. The χ^2 error between the actual and recovered distribution is 0.2465 for 30 spectra. The experiment shows that the spectral distribution of the illuminant can be recovered if the objects reflect highlights.

highlights can be seen in the image: The vertical bright stripes correspond to highlights. The figure also shows the actual and recovered spectral distribution of the illuminant in the graph. The experiment shows that, if both objects contain highlights, the color of the illuminant can be estimated for multi-spectral data.

The fifth experiment is shown in fig. 7. This time, the same objects of the fourth experiment are illuminated with a red light source. The objects are positioned in such a way that the blocks reflect highlights. The figure also shows the actual and recovered spectral distribution of the illuminant in the graph. The experiment shows that, if both objects contain highlights, the color of the illuminant can be estimated for multi-spectral data.

The sixth experiment is shown in fig. 8. Two objects, with red and yellow colors, are illuminated with a halogen light source. The objects are positioned in such a way that the blocks reflect highlights. The figure also shows the actual and recovered spectral distribution of the illuminant in the graph. The experiment shows that, if both objects contain highlights, the color of the illuminant can be estimated for multi-spectral data.

The seventh experiment is shown in fig. 9. The same objects as shown in fig. 8 are illuminated with a red light source. The objects are positioned in such a way that the blocks do not reflect highlights. The figure also shows the

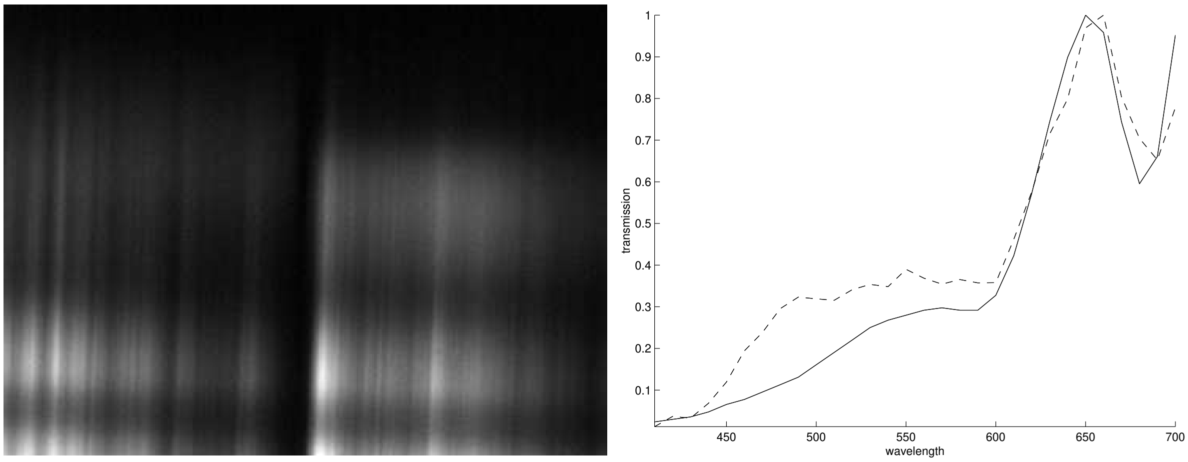

Figure 7. Left: Multi-spectral images of blue and green colored wooden blocks reflecting highlights under red illumination. Right: Actual vs. recovered spectal distribution of illuminant. The recovered spectral distribution of the illuminant is displayed as a dashed line, the actual spectral distribution of the illuminant is displayed as a solid line. The χ^2 error between the actual and recovered distribution is 0.596 for 30 spectra. The experiment shows that the spectral distribution of the illuminant can be recovered if the objects reflect highlights.

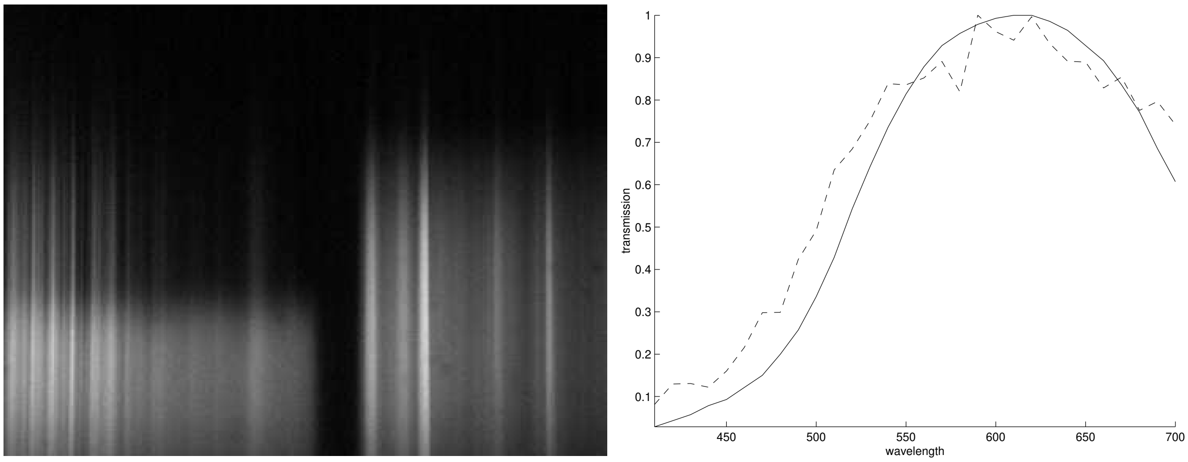

Figure 8. Left: Multi-spectral images of red and yellow colored wooden blocks reflecting highlights. Right: Actual vs. recovered spectal distribution of illuminant. The recovered spectral distribution of the illuminant is displayed as a dashed line, the actual spectral distribution of the illuminant is displayed as a solid line. The χ^2 error between the actual and recovered distribution is 0.305 for 30 spectra. The experiment shows that the spectral distribution of the illuminant can be recovered if the objects reflect highlights.

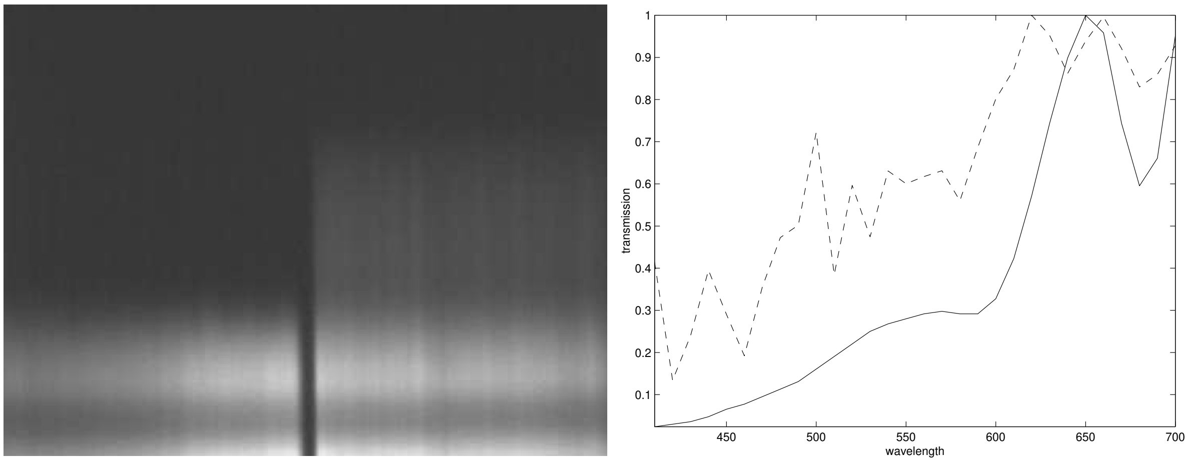

Figure 9. Left: Multi-spectral images of red and yellow colored wooden blocks reflecting highlights. Right: Actual vs. recovered spectal distribution of illuminant. The recovered spectral distribution of the illuminant is displayed as a dashed line, the actual spectral distribution of the illuminant is displayed as a solid line. The χ^2 error between the actual and recovered distribution is 2.929 for the 30 spectra. The experiment shows that the spectral distribution of the illuminant can not be recovered if the objects do not reflect highlights.

actual and recovered spectral distribution of the illuminant in the graph. The experiment shows that, if no highlights are present, the color of the illuminant can not be estimated for multi-spectral data.

5. DISCUSSION

Our goal is to recover the spectral distribution of the illuminant from highlights analysis in multi-spectral images based on the dichromatic reflection model. We employ a robust clustering technique to obtain the body and surface components in a multi-spectral space. These components determine the direction of the illumination spectral color. Then, we recover the illumination spectral power distribution by using principal component analysis for all wavelengths. To obtain the most reliable estimate of the spectral power distribution of the illuminant, all possible combinations of wavelengths are used to generate the optimal averaged estimation of the spectral power distribution of the scene illuminant. Experiments conducted on multi-spectral images show that the method yields a good estimate of the spectral distribution of the illuminant. Our method is restricted to images containing a substantial amount of body reflection and highlights.

REFERENCES

1. H.-C. Lee, "Method for computing the scene-illuminant chromaticity from specular highlights," *Journal of the Optical Society of America* **3**(10), pp. 1694–1699, 1986.
2. G. Healey, "Estimating spectral reflectance using highlights," *International Journal of Computer Vision* **2**(1), pp. 7–32, 1988.
3. S. Tominaga and B. A. Wandell, "Standard surface-reflectance model and illuminant estimation," *Journal of the Optical Society of America* **6**(4), pp. 576–584, 1989.
4. S. A. Shafer, "Using color to separate reflection components," *Color Research Applications* **10**, pp. 210–218, Winter 1985.
5. H. Stokman and T. Gevers, "Photometric invariant region detection in multi-spectral images," in *Proceedings of Vision Interface '99*, (Trois Rivieres, Quebec, Canada), 19-21 May 1999.
6. Labsphere Inc., Reflectance Calibration Laboratory, *Calibration Certificate*, September 1998. Spectralon Reflectance Target, Sample I.D.: SRT-99-050.