Practical Total Appearance Imaging of Paintings

Roy S. Berns and Tongbo Chen, Munsell Color Science Laboratory, Rochester Institute of Technology, Rochester, NY USA

Abstract

The total appearance of a painting is defined by its spatially varying spectral reflectance factor, surface macrostructure (depth or surface normal), and surface microstructure (bi-directional reflectance distribution function, BRDF). For paintings with uniform BRDF (e.g., varnished), their total appearance can be measured using equipment commonly found in a photographic studio. Such a system was built and tested for several acrylic-dispersion paintings. The system consisted of three strobes affixed with triacetate film linear polarizers and a Dual-RGB camera also affixed with a linear polarizer in order to achieve cross polarization. Using the principles of photometric stereo, images of each light source taken sequentially from 30° from the normal and 120° apart annularly were used to measure surface normal. A learning-based algorithm was used to measure colorimetry and spectral reflectance factor. Software, Artviewer*, was written to render images for specific geometries and for studio lighting. The system produced images that approximated, but not equaled, conventional studio photography. Because diffuse data were collected, these images are useful for the long-term evaluation of color changes. Evaluation of the surface normal provide new information for the technical examination of artwork.*

Introduction

Imaging artwork for documentation and reproduction has a long and rich history. The vast majority of such imaging reduces an illuminated three-dimensional object onto a two-dimensional plane, rendering a specific observing experience. For this reason, successful museum photographers have backgrounds in art history and aesthetics enabling them to produce images that convey important appearance phenomena about the object [1]. The drawback to this practice is a need for reshooting when the appearance criteria change. It would be ideal if we can separate the image capture and rendering. In this manner, the object can be re-rendered for various criteria. In other words, we first image the artwork to define its physical characteristics, or total appearance. Then computer graphics techniques are used to produce an image for specific lighting, observing or publishing criteria.

For artwork such as paintings and drawings, a complete physical description would include spatially varying spectral reflectance factor, surface macrostructure (depth), and surface microstructure (bi-directional reflectance distribution function, BRDF) [2]. Spectral reflectance factor and BRDF can be measured with a spectral camera and a single light source positioned about the object (or multiple light sources). Depth can be measured with a variety of techniques including laser scanning, confocal microscopy, and structured light. An alternative to measuring depth is to measure surface normal, which is a vector direction perpendicular to the surface. Measurements of spectral reflectance factor, BRDF, and surface normal can be accomplished with a single imaging system. However, if the object surface has appreciable impasto and is not matte, hundreds of images may be required to assure that both diffuse and specular reflections have been captured for every point on the object. Furthermore, these many images need to be captured at every wavelength within the visible spectrum. Obviously, such a system would remain within a research domain.

Beginning in 2006 a research program was initiated to develop an abridged approach to measuring the total appearance of paintings. One goal was developing an artist material database [3] where 600 samples were produced [4] and their BRDF's estimated using the Ward model [5]. If an object is assumed to have uniform BRDF characteristics, the database can be used to define the object's BRDF, reducing the imaging system complexity significantly. The data capture methodology developed in this research used three strobe lights affixed with linear polarizers and a cross-polarized Dual-RGB camera to measure surface normal and spectral reflectance factor. This approach will be referred to as "3LI": three light imaging.

Basic Assumptions

Several assumptions are made to achieve practicality. First, the object is flat. For most paintings and drawings, any lack of planarity is small relative to both the size of the object and the distance between the object and camera. Second, all the object's reflections can be described using a dichromatic reflection model where the color information is contained in the body reflection and the gloss information is contained in the surface reflection [6]. As described below, the system only measures body reflection. For paintings with colored surface reflections, such as gold leaf or coatings containing metallic particles and interference pigments, the rendered images will not be color accurate in these passages.

Measuring Surface Normal

One of the most straightforward techniques to measure surface normal is known as photometric stereo where an object is imaged using lights positioned in different locations [7]. For objects that are completely matte (Lambertian) and for lighting that is a point source at infinite distance, the surface normal is calculated using Eq. 1 for a three-light system:

$$\mathbf{n} = \mathbf{L}^{-1}\mathbf{I} \tag{1}$$

where

$$\mathbf{n} = \begin{pmatrix} n_x \\ n_y \\ n_z \end{pmatrix} \quad \mathbf{L} = \begin{pmatrix} x_1 & y_1 & z_1 \\ x_2 & y_2 & z_2 \\ x_3 & y_3 & z_3 \end{pmatrix} \quad \mathbf{I} = \begin{pmatrix} I_1 \\ I_2 \\ I_3 \end{pmatrix}.$$

Vector **n** is the surface normal defined by orthogonal directions *x*, *y*, and *z* where *x* and *y* are within the plane of the painting and *z* is perpendicular to the painting; **L** defines

normalized light directions for lights 1, 2, and 3, and **I** defines the object's intensities (reflected light) as a function of position (pixels). Since **n** consists of three unknowns, we need at least three different lighting directions. Quite often, dozens of light directions are used so that specular reflections can be avoided and to provide an over-determined approach to estimating **n**. In this research, specular reflections were eliminated using cross-polarization and the three directions selected that balanced accuracy and dimensional limitations (e.g., room size).

Three Broncolor PulsoG 1600 J strobes with P70 reflectors, barn doors, and affixed triacetate film linear polarizers were used in this research. Exploratory experiments were carried out where the three strobes were positioned in several different configurations. Measurements of **n** were compared with those obtained for a small painting using a rotating system designed for measuring both BRDF and **n** and with measurements of depth using a laser scanner. At first, two of the strobes were positioned 45° from the normal on either side of the object plane, typical of many museum imaging setups, and the third strobe was above the object, also at 45° from the normal. In this manner, such a lighting system could be used for both conventional imaging and for measuring **n**. However, this configuration produced poor results because the lighting was over-weighted about the object plane's annulus. By positioning the lights symmetrically 120° apart within an annulus (azimuth) and 30° from the object plane normal, accuracy improved dramatically.

The camera used in this experiment was a Canon Mark II, modified by replacing its blue-green cover glass with BK7 glass. It was affixed with a Canon 100mm 1:2.8 macro lens. Attached to the front of the lens was a filter wheel that rotated two filters in the beam path, sequentially. A B&W 77mm linear polarizer with multi-resistant coating was screwed onto the front of the filter wheel assembly. The imaging system was positioned along the normal to the object plane. Each filter was a glass absorption filter with an AR coating on the camera side and a hot-mirror coating on the opposite side resulting in a visible bandpass filter.

The first step in the calibration process was to achieve cross polarization. A pair of square linear polarizers were mounted in a frame as reference vertical and horizontal polarizers. An ancillary tungsten source was aimed through the square polarizer and the camera polarizer was rotated to define vertical and horizontal positions. These were marked on the filter edge. The camera was set to vertical polarization.

A glossy black ball (cue ball) was mounted on a tripod with a boom extension and positioned at the center of the object plane. Images were taken at various rotations of the light source polarizer using the strobes. Cross polarization was achieved when the specular reflection from the ball was minimized, evaluated visually when the images were displayed. An image of the ball and its three barely visible specular highlights is retained. (Note that the modeling incandescent lights should not be used: They will melt the polarizers.)

A diffuse white board, made by mounting Color-aid White 10 paper on a flat metal substrate, was placed at the object plane. The energy of each light was optimized for dynamic range and images were captured of the white board using each strobe and each filter, sequentially, resulting in six images. These images were low-pass Gaussian filtered to reduce any local non-uniformities and the board's surface normal was assumed to be perpendicular to the object plane, formally $\mathbf{n} = (0\ 0\ 1)^{\mathrm{T}}$.

A set of color targets were mounted on the flat metal substrate and imaged using each strobe and each filter with the identical camera and strobe settings as the diffuse white board.

The artwork was positioned at the object plane and imaged using each strobe and each filter.

A highlight position is defined as $\boldsymbol{H} = [H'_x, H'_y, H'_z]^T$ where H' is normalized highlight position: $H' = H/\|H\|_2$. The user interactively selected three positions on the edge of the cue ball, used to calculate the ball's radius and center, and the center of each specular reflection, was used to define H_x and H_y. Direction H_z was calculated by $\sqrt{R^2 - H_x^2 - H_y^2}$. Since the distance between the camera and the ball is much larger than the size of the ball, we approximated the viewing direction as $\boldsymbol{V} = [0, 0, 1]^T$. According to the law of reflection, the direction of the light, L was computed as:

$$\boldsymbol{L} = \mathbf{2}(\boldsymbol{V}^{\boldsymbol{T}}\boldsymbol{H})\boldsymbol{H} - \boldsymbol{V} \tag{2}$$

Because the strobes were not point sources at infinite distance, flat fielding was performed including cosine falloff compensation by multiplying the dot product of each i_{th} light direction and the surface normal of the white board for each j_{th} filter:

$$I_{i,j} = z_{i,j} I'_{i,j} / W_{i,j} \quad i = \{1,2,3\}; j = \{1,2\} \tag{3}$$

where I is the light intensity of the object after flat-fielding and cosine compensation, z, I' is the raw maximum camera signal, and W is the white board camera signal.

A single process was written to convert CR2 to TIFF using dcraw, flatfield using Eq. 3, align the images from the two filters using correlation-based translation, and write a single six-channel 16-bit linear photometric TIFF file.

Depending on the spectral reflectance of the artwork at a given pixel, one of the six channels had a maximum value. This channel's signal defined the image intensity.

Using Eq. 1, the surface normal was calculated for each pixel of the artwork and stored as a floating-point image file.

Measuring Diffuse Spectral Reflectance Factor and Colorimetry

There are a number of approaches to building an imaging spectrophotometer for artwork [8]. In this research, an RGB color-filter array camera was used where replacing its blue-green cover glass with clear glass was first performed. A filter wheel positioned two custom filters in front of the camera system, sequentially. The filters were designed to optimize spectral accuracy, colorimetric accuracy, and spatial image quality (by maximizing light throughput). The final filters were selected such that one of the filters resulted in spectral sensitivities nearly identical with the camera before cover-glass replacement, enabling the camera to be used in its standard RGB configuration and as a spectral camera. This approach has been named the Dual-RGB approach [9, 10].

A learning-based calibration was used where images were collected of the X-rite ColorChecker Classic using each strobe, sequentially. The images were flatfielded by dividing by the low-

pass Gaussian filtered images of the white board and normalized by the white board's average spectral reflectance factor. The images from the three different lighting geometries were averaged. Each color patch was measured using a Macbeth 7000 integrating sphere spectrophotometer in specular excluded mode. The average digital counts for each of the 24 color patches were calculated.

Two transformations were generated. The first, Eq. 4, was a matrix transformation from the six camera signals to CIE tristimulus values for illuminant D50 and the 1931 standard observer. The matrix coefficients were optimized to achieve high color accuracy and spatial image quality. The color accuracy of the camera system was an average of $1.4\Delta E_{00}$.

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = \mathbf{M}_{XYZ} \begin{pmatrix} R_{filter\,1} & G_{filter\,1} & B_{filter\,1} & R_{filter\,2} & G_{filter\,2} & B_{filter\,2} \end{pmatrix}^T \tag{4}$$

$$\begin{pmatrix} R_{\lambda=380} \\ \dots \\ R_{\lambda=730} \end{pmatrix} = \mathbf{M}_R \begin{pmatrix} R_{filter\,1} & G_{filter\,1} & B_{filter\,1} & R_{filter\,2} & G_{filter\,2} & B_{filter\,2} \end{pmatrix}^T \tag{5}$$

The second, Eq. 5, was a matrix transformation from the six camera signals to spectral reflectance factor where weighted linear regression was used. The weighting coefficients were defined to improve the accuracy of predicting the gray scale.

The artwork images were flatfielded and averaged in the same manner as the calibration target. The matrices derived for the ColorChecker were used to transform the camera data to either colorimetry or spectral reflectance factor. The tristimulus values were used to calculate sRGB images for display following a chromatic adaptation transformation using CIECAM02 from D50 to D65.

The average of the three lights produce nearly diffuse illumination and when combined with cross-polarization removing any specular highlights, the diffuse image does not have the range of spatial frequencies that are typical of conventional imaging. (Even a bounce flash produces a slight amount of highlights and shadowing.) As a consequence, the diffuse images may appear slightly blurry. To improve the appearance of these images, unsharp masking was used on the achromatic channel during rendering.

Software was written, *Artviewer*, to render images for specific geometries. A screenshot is shown in Figure 1 where the user can interactively define lighting geometry, adjust diffuse and BRDF characteristics, and view either color or rendered normal maps as the dot product with a specific light direction, $\boldsymbol{n} \cdot \boldsymbol{L}$. The artist material database can be used to define BRDF. One unique rendering feature is simulating studio lighting for a pair of point sources with adjustable geometries and intensities.

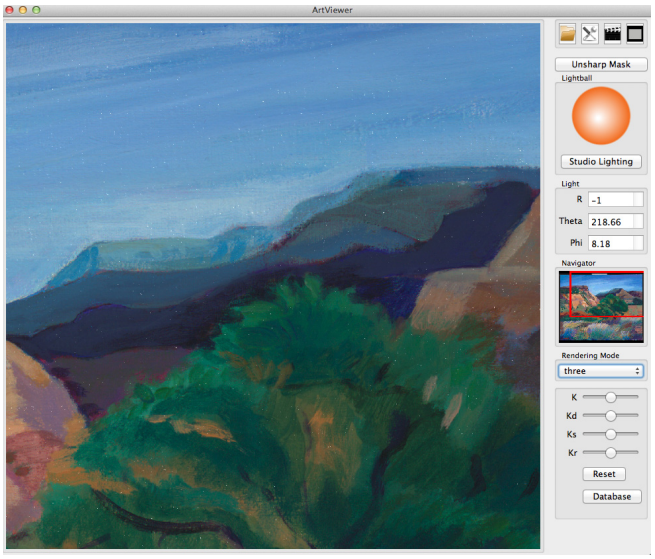

Figure 1*. Screenshot of Artviewer.*

Results and Discussion

Two paintings were evaluated. The first, *Trees*, (8" x 10") has been evaluated a number of times and used for a variety of rendering applications including tangible displays, spectral rendering, spatially varying BRDF (SVBRDF) measurements, and gloss analyses. In addition to 3LI, Trees was imaged conventionally using strobes on either side of the painting at 45° from the normal, 180° apart annularly (Figure 2). This painting has significant SVBRDF properties where a gloss medium was applied selectively to areas of moderate impasto as well as a large color gamut. This painting was laser scanned using an Arius RGB scanner and from these data, **n** was calculated. It was also imaged using a CHI dome system at the Museum of Modern Art, referred to as RTI, for reflectance transformation imaging.

Figure 2. Trees*: conventional imaging.*

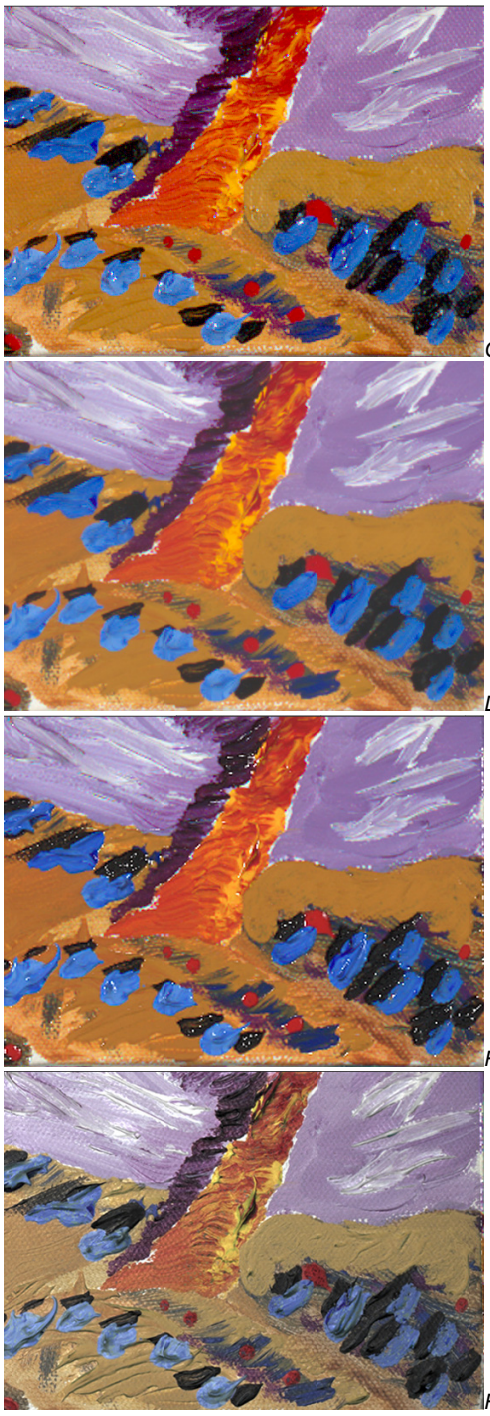

Figure 3. Trees*: conventional, diffuse, rendered, and RTI imaging.*

Details of Trees comparing conventional, diffuse, rendered for studio lighting, and RTI are shown in Figure 3. The two closest RTI images to this geometry were averaged. Lightness histogram equalization was performed on the RTI image as an attempt to improve the color rendering. As expected, the diffuse image does not have any specular highlights and is blurrier than conventional imaging. This image was sharpened and a varnished acrylic brushed canvas's BRDF was defined as the painting's surface. A slight tonal adjustment was also added to account for the diffuse illumination. This rendering was a reasonable approximation of the conventional lighting. Because Trees has spatially varying gloss, it isn't possible for 3LI to accurately render this painting since only a uniform gloss can be imposed. The RTI image produced the most visible surface topography. It also had the greatest occlusion artifacts (black areas) since the lighting is quite directional compared to strobe illumination.

It was of particular interest to determine whether the lack of visible topography was a result of 3LI or the studio rendering in Artviewer. The rendered surface normal maps for 3LI and laser scanning are compared in Figure 4 (the 3LI image was sharpened to accentuate the data). The topography not visible in the rendered image is seen in the normal map. The 3LI map has more detail than achieved by the laser scanning, the result also found by MacDonald when comparing photometric stereo using dome illumination and laser scanning [11].

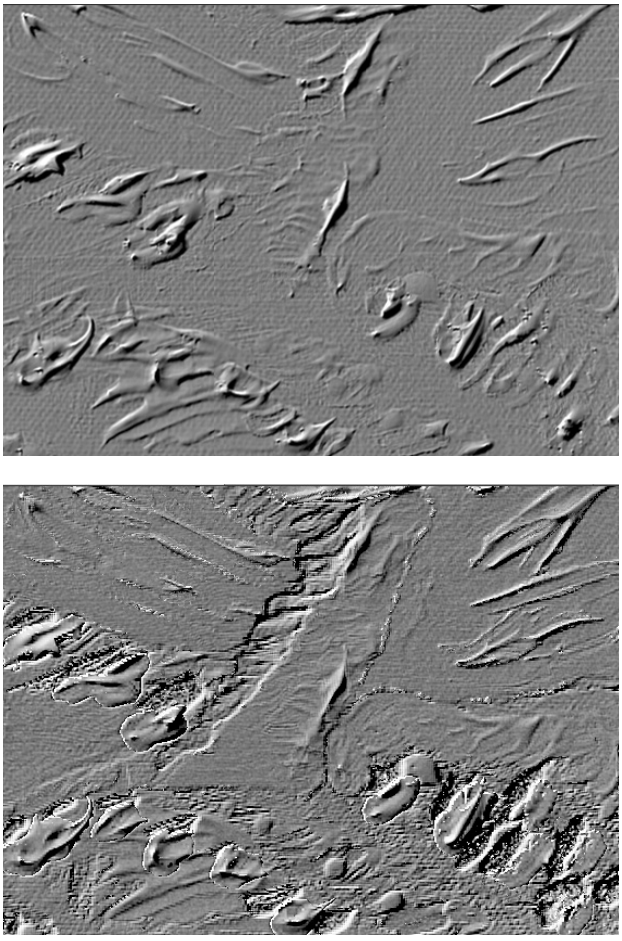

Figure 4. Trees*: surface normal rendering for laser scan (top) and 3LI (bottom).*

The second, *Landscape*, (24" x 30") was made using acrylic-dispersion paints and was varnished with a medium-gloss polymer, shown in Figure 5 using conventional imaging. Because of the size

of our imaging studio, the right side strobe was closer to the painting then best practices. As a consequence, there are considerable specular reflections on the right side of the painting.

Figure 2. Landscape: *conventional imaging.*

Image close-ups are shown in Figure 6. The diffuse image, by definition, does not contain any specular reflections. From a conservation perspective, this image would be better to evaluate long-term color changes. The rendered image does not reproduce the specular reflections seen in the conventional image. The surface normal map reveals that the structure has been captured.

Conclusions

A practical approach to measuring the total appearance of painting has been developed, 3LI, three-light imaging. The approach uses polarization-enhanced photometric stereo to measure surface normal and Dual-RGB imaging to measure colorimetry and spectral reflectance factor. The system has a number of advantages. First, it uses conventional lighting, strobe illumination. Second, it is not size limited. Any size object can be imaged. Although a Dual-RGB modified commercial camera was used to improve color accuracy and enable spectral capture, this approach would be equally effective using an RGB camera to measure surface normal. Third, because diffuse reflectance is measured, it produces colorimetric data that are not confounded by specular highlights. As such, it provides excellent image data to evaluate long-term color stability. By removing the lighting geometry, issues such as matching geometry and ambiguity between white passages and specular highlights are avoided.

One disadvantage of the system is its inability to render an image if the surface has spatially varying gloss.

These qualitative results indicate that the rendering software should be improved when simulating studio lighting. Rather than a point source, an environment map may prove more appropriate. This may reduce the lack of specularity for *Landscape*.

Future research will be aimed at performing psychophysics to quantify these differences.

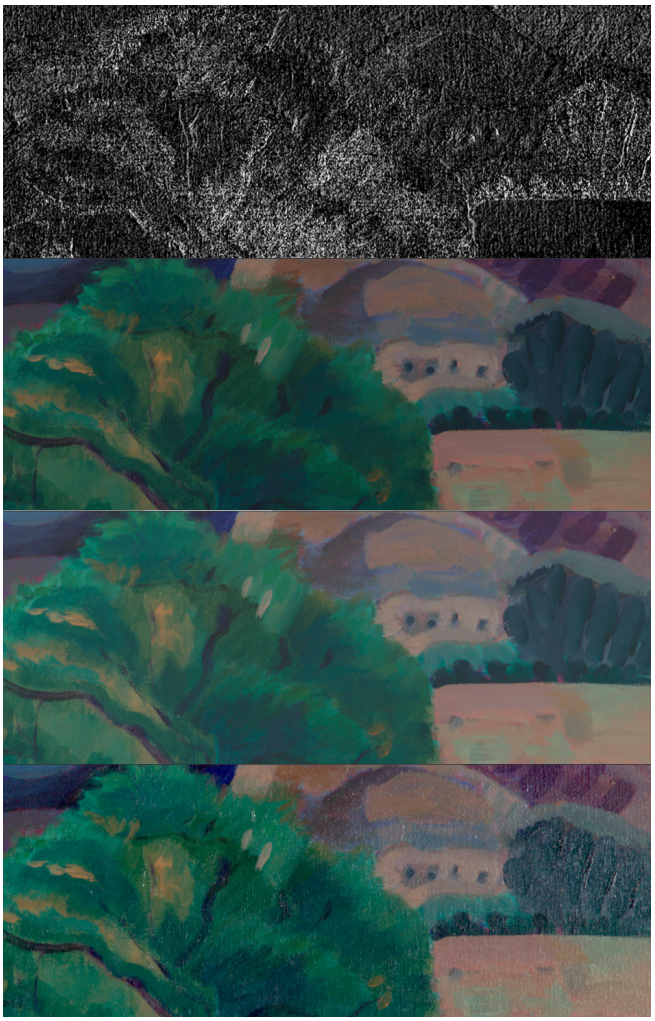

Figure 2. Landscape: *surface normal (top), diffuse (second from top), rendered (third from top), and conventional (bottom).*

References

[1] R. S. Berns, F. S. Frey, M. R. Rosen, E. P. Smoyer and L. A. Taplin, Direct Digital Capture of Cultural Heritage - Benchmarking American Museum Practices and Defining Future Needs Final Report, RIT, http://www.cis.rit.edu/museumSurvey/documents/Benchmark_Final_Report_Web.pdf (2005).

[2] J. Dorsey, H. E. Rushmeier, and F. X. Sillion, *Digital modeling of material appearance*, Kaufmann/Elsevier, Amsterdam (2008).

[3] Artist Material BRDF Library, http://art-si.org/.

[4] J. C. Ashbaugh, R. S. Berns, B. A. Darling, and L. A. Taplin, "Artist material BRDF database for computer graphics rendering," Proc. IS&T/SID Seventeenth Color Imaging Conference, 62-68 (2009).

[5] G. J. Ward, Measuring and modeling anisotropic reflection, SIGGRAPH Comput. Graph. 26, 265–272 (1992).

[6] S. Shafer, Using color to separate reflection components, Color Research and Application 10, 210–218 (1985),

[7] R. J. Woodham, Photometric method for determining surface orientation from multiple images, Optical Engineering 19, 139–144 (1980).

[8] C. Fischer and I. Kakoulli, Multispectral and hyperspectral imaging technologies in conservation: current research and potential applications, Reviews in conservation, 7, 3-16 (2006).
[9] 7,554,586 Francisco H. Imai and Roy S. Berns, System and method for scene image acquisition and spectral estimation using a wide-band multi-channel image capture, June 30, 2009, Assignee: Rochester Institute of Technology (Rochester, NY).
[10] R. S. Berns, L. A. Taplin, M. Nezamabadi, M. Mohammadi, "Spectral imaging using a commercial color-filter array digital camera," Proc. 14th Triennial Meeting The Hague, ICOM Committee for Conservation, 743-750 (2005).
[11] L. MacDonald and S. Robson, Polynomial texture mapping and 3d representations, International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, 38 (5) 422-427 (2010).

Acknowledgments

This research was supported by the Andrew W. Mellon Foundation.

Author Biography

Roy S. Berns is the Richard S. Hunter Professor in Color Science, Appearance, and Technology and Director of the Munsell Color Science Laboratory within the Center for Imaging Science at Rochester Institute of Technology, USA. He received B.S. and M.S. degrees in Textiles from the University of California at Davis and a Ph.D. degree in Chemistry from Rensselaer Polytechnic Institute (RPI). Berns has received scientific achievement awards from the Inter-Society Color Council, the Society of Imaging Science and Technology, and the Colour Group of Great Britain. He is the author of the third edition of "Billmeyer and Saltzman's Principles of Color Technology.

Tongbo Chen is a Postdoctoral Fellow in the Munsell Color Science Laboratory. He has B.S. (Harbin Institute of Technology), M.S. (Beijing University of Technology), and Ph.D. (Max-Planck-Institut Informatik and Saarland University) degrees in Computer Science. Before joining RIT, he was a Postdoctoral Research Associate at the USC ICT Graphics Laboratory.

.