

Real-Time Multispectral Rendering with Complex Illumination

Benjamin A. Darling, James A. Ferwerda, Roy S. Berns, Tongbo Chen; Munsell Color Science Laboratory, Rochester Institute of Technology, Rochester, NY

Abstract

Techniques for rendering with illumination maps captured from real-world environments have provided the means to produce complex, realistic patterns of reflections in real-time computer graphics simulations. While the directional aspects of light emission and surface reflection are modeled comprehensively with these methods, the spectral component is typically represented with a three-channel RGB workflow. In this paper, we present a multispectral framework for capturing environmental illumination maps and incorporating them into a real-time rendering pipeline. The framework maintains two sets of trichromatic data through a real-time rendering workflow, using a six-channel representation optimized to allow for color-accurate relighting under a range of illuminants. We present the design of an end-to-end computer graphics color pipeline from capture to display, provide simulation results on the pipeline's colorimetric accuracy, and develop a prototype to demonstrate how the framework can be implemented in a GPU-based graphics shader architecture.

Introduction

Computer graphics techniques that allow for realistic image synthesis have enabled significant advances in a wide range of fields including simulation and training, product design, and entertainment. The basic image synthesis pipeline simulates the interactions of light sources with material surfaces. Accurate modeling of local light reflection (BRDF) and light transport has significantly improved the accuracy and realism of synthetic images [1]. This realism has been further enhanced by image-based lighting techniques [2] for capturing illumination maps from real-world environments. With recent advances in graphics technologies, environmental illumination maps can be used to light complex materials at interactive frame rates, allowing for realistic patterns of reflections in real-time applications [3].

Traditionally, imaged-based lighting techniques have focused on capturing and simulating the complex geometric properties of light emission and surface reflection. The spectral nature of the light-surface interaction has typically been represented with the three-channel RGB workflow common to traditional imaging systems. The use of a trichromatic workflow is computationally efficient, which is an important consideration in real-time rendering, but can lead to color error in the rendered surfaces. A trichromatic representation can produce accurate results when only one lighting spectral distribution must be considered, but higher-dimensional spectral information is required to accurately determine the color of a surface over a range of other lighting distributions. In the case of image-based lighting capture of real environments, it is common that a range of spectral distributions will be present, either from multiple primary light sources in the scene (as in the example of Figure 1) or from the indirect illumination reflected from chromatic environmental objects. A

full-spectral rendering approach would provide the highest color accuracy for these conditions, but would be too resource intensive for a real-time pipeline. Abridged spectral rendering representations [4, 5, 6, 7], however, have been shown to retain much of the information needed to accurately relight surfaces, while only requiring a small number of additional color channels.

In this paper, we present a multispectral framework for real-time rendering with environmental illumination. It is designed to provide a high degree of colorimetric accuracy across illuminants while maintaining compatibility with standard methods for rendering with environmental illumination maps. We adopt an abridged spectral rendering approach, develop a set of optimized primaries for performing the rendering calculations, and develop methods to incorporate real-world data captured using multispectral imaging systems into the rendering pipeline. In an effort to develop a real-time system that is both color-accurate and practical, the framework was designed with the following principal requirements:

- Accurate colorimetry for a range of typical illuminants
- A maximum of six channels, so that processing can be performed by repeating a standard RGB workflow
- A single set of rendering primaries, so surface and illumination data can be maintained in one standardized form
- Support for illumination and reflectance data captured using a range of spectral measurement and imaging systems

In the following sections, we describe background and related work, present the design of the end-to-end rendering pipeline, provide simulation results on the pipeline's accuracy, and develop a prototype to demonstrate how the framework can be implemented for real-time rendering using a modern GPU-based graphics shader architecture.



(a) Three-Channel Rendering

(b) Six-Channel Rendering

Figure 1. 3D models lit using a multispectral illumination map captured in a mixed-lighting environment. A window, left, and tungsten lamp, right, can be seen reflected in the gloss of each teapot. The bands below the models show their expected color based on their spectral reflectance and full spectral measurements taken of the two light sources. With three channels (a), the change from daylight to tungsten introduces color error. The six-channel workflow (b) is able to better maintain accuracy across the mixed illumination.

Background and Related Work

The physical process simulated by computer graphics rendering, that of light interacting with material surfaces and producing reflections, is fundamentally a spectral process. As in standard colorimetry, it is necessary to multiply the illumination and the (bidirectional) surface reflectance on a per-wavelength basis to fully determine the color of rendered reflections (Figure 2a). However, the use of a high-dimensional spectral representation is often not practical in computer graphics, and so surface and illumination data are often reduced to small number of channels before light-surface calculations are performed.

The data representations used in computer graphics are typically three-dimensional and, in many cases, rendering calculations are performed in device-specific RGB. To improve color processing, trichromatic representations based on standardized color spaces were introduced into computer graphics systems [8, 9]. Borges [10] evaluated the accuracy of performing rendering calculations directly in the CIE XYZ primaries. Ward and Eydberg-Vileshin [11] found the accuracy could be improved by transforming the XYZ primaries to a sharpened RGB space and pre-filtering the surface representation based on the spectral properties of the lighting.

For more accurate color simulations, spectral rendering methods have been incorporated. Johnson and Fairchild [12] developed a 93-band full spectral rendering workflow to model wavelength-based color phenomena. Abridged spectral rendering methods have been developed to provide increased color accuracy over three-channel methods, while reducing the computational load relative to full spectral rendering. Meyer [4] developed methods to select optimal sets of point-sampled wavelengths for use in rendering. Peercy et al. [5, 6] created frameworks for abridged spectral rendering based on linear orthonormal basis functions, using a general set if spectral distributions were sufficiently smooth, and or ones tailored to the illumination for spiky distributions. Drew and Finlayson [7] developed an abridged spectral rendering approach based on the use of sharpened basis functions, and demonstrated that performing rendering calculations by multiplying the basis function coefficients of the lighting and reflectance could produce results similar to full spectral calculations. We adopt a similar abridged rendering approach, where color calculations are performed by multiplying the coefficients of a set of basis functions, though we derive a different set of spectral curves for use as our optimized basis functions.

Image-based Environmental Illumination

Techniques for capturing and rendering with illumination maps from real-world environments have provided the means to produce complex, realistic patterns of reflections in computer graphics simulations. Debevec [2] developed a method to capture the radiance information from real-world environments by taking a sequence of photographs of a mirror sphere at a range of different exposure settings. The photographs are merged to a high dynamic image and used to produce an omni-directional illumination map of the environment [2]. Colbert et al. [3] introduced a real-time importance sampling method for illumination maps, allowing them to be used to light surfaces with complex BRDF models at interactive frame rates.

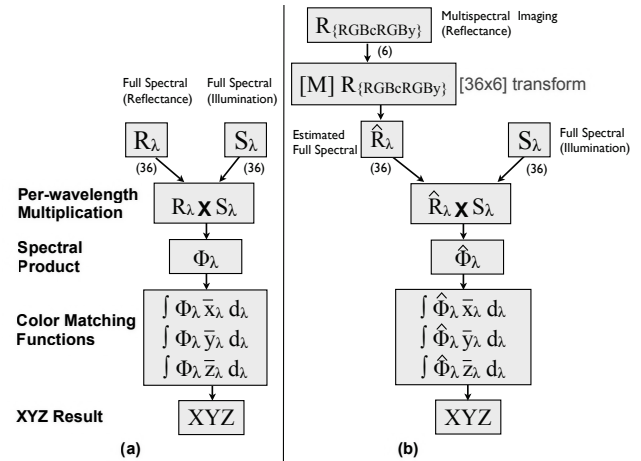


Figure 2. Left (a), a flowchart depicting the standard calculation of XYZ values from spectral reflectance factor and illumination data. Right (b), in a typical multispectral workflow, surface reflectance captured in a smaller number of channels (6) is first transformed to an estimate of full spectral reflectance.

Spectral Imaging

Spectral imaging methods allow for the capture of higher-dimensional illumination and surface data than provided by traditional trichromatic imaging systems. Hardeberg et al. [13] developed a system for estimating spectral reflectance from multispectral images acquired using a monochrome CCD camera and a set of optimized spectral filters, to more accurately predict the color of the captured surfaces under different illuminants. Multispectral imaging methods using trichromatic color-filter array cameras, shot sequentially through two additional filters, have been developed to provide six spectral bands for estimating surface reflectance [14]. While multispectral capture is typically used to estimate surface reflectance data, recently Tominaga and colleagues [15, 16] applied multispectral imaging techniques to capture illumination maps of real-world environments. By imaging through multiple filters when capturing mirrored spheres or with panoramic cameras, they have been able to estimate the spectral distribution of illumination in real-world environments. These spectral illumination maps were used for image-based lighting in an offline rendering process.

In multispectral imaging workflows, color calculations are typically performed by first transforming the abridged multispectral data to a full-spectral representation. A typical case, illustrated in Figure 2b, is to capture surface reflectance data with a multispectral imaging system (e.g. 6-channel), and then estimate the reflectance factor at the full spectral resolution of the lighting (36 channels), so that tristimulus values can be calculated with the color matching functions of a CIE standard observer. This is feasible for an off-line imaging calculation and has the advantage that it preserves all the available spectral information when estimating the surface color. This approach, however, is less practical for real-time computer graphics rendering where multiple lighting calculations may be required per pixel and computational time is a critical factor.

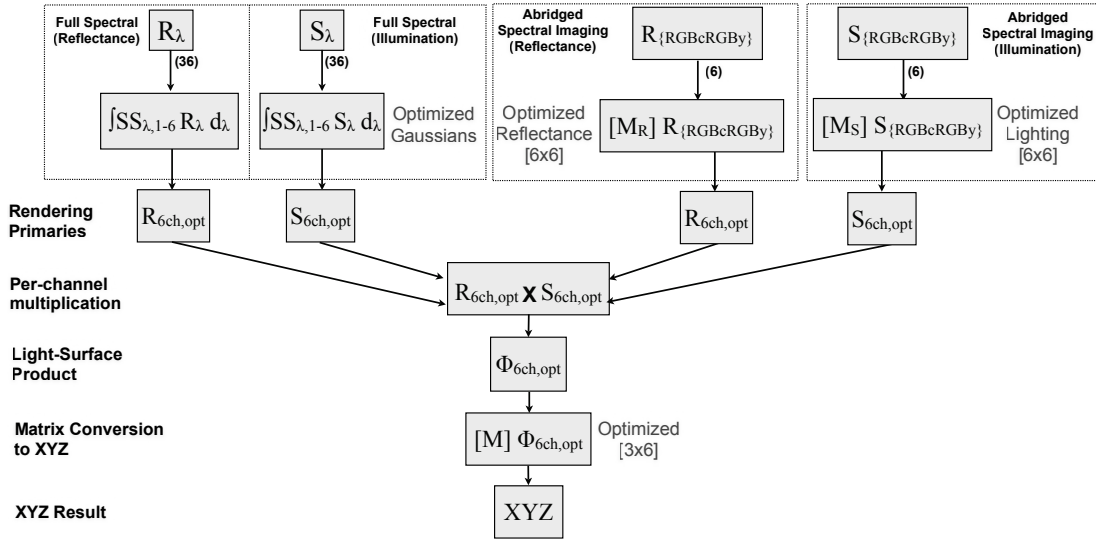


Figure 3. Calculation of tristimulus values using a six-channel rendering workflow. In a pre-processing stage (dashed boxes), full spectral or abridged multispectral data on surface reflectance and illumination are converted to the six rendering primaries. Once all data are in this common form, lighting calculations are performed in the six channels and the results converted to XYZ values.

Framework for Multispectral Rendering

Design Objectives

In developing the framework, our two principal objectives were to maintain a high level of colorimetric accuracy for a range of illumination, while still performing lighting calculations in a small enough number of spectral channels to be suitable for real-time rendering. In particular, by limiting the number of rendering channels to six, the rendering workflow can closely follow a typical trichromatic RGB process and merely repeat the rendering calculations for a second set of RGB data. The input data, typically stored as RGB images, can still be represented in this form, by maintaining just one an additional image per reflectance map or illumination map.

An additional design objective was to base all the rendering calculations on a single set of rendering primaries. Certain methods for color-accurate rendering, such as point sampled wavelengths or basis vectors derived from characteristic vectors analysis, customize the rendering primaries to the lighting in each virtual environment. While this can be advantageous for color accuracy, it requires that a full spectral representation of surface reflectance factor be maintained so that the abridged representation can be recalculated for each new lighting environment. With one fixed set of rendering primaries, once the data is converted to this representation, it can stay in this final form and does not require additional data to be maintained or re-processing to be performed.

The final objective is to allow input data captured or specified using different methods to be used within the rendering pipeline. Regardless of how surface reflectance and illumination data are initially captured, whether by multispectral imaging or full spectral devices like spectrophotometers and spectroradiometers, the goal is to allow the data to be used interchangeably when rendering.

Design

The color-processing pipeline, illustrated in Figure 3, has two main components. There is an offline pre-processing portion

where surface and lighting data are converted to the set of rendering primaries, and a real-time rendering portion where the surface shading calculations are performed by multiplying illumination and surface reflectance.

By converting all the data to a common form, the pre-processing stage allows the input data captured with different measurement devices to be used interchangeably within the rendering system. To support full spectral data input, a spectral sensitivity curve has been derived for each of the six primaries. Full spectral data are converted by summing the product of the data and derived curves:

$$R_{i,OPT} = \sum_{\lambda=380}^{730nm} SS_{\lambda,i} R_{\lambda} \quad (1)$$

where $SS_{\lambda,i}$ is the normalized spectral sensitivity of the i^{th} primary, R_{λ} is the spectral reflectance factor, and $R_{i,OPT}$ is the resulting scalar for the i^{th} primary. A similar calculation is performed to convert light source spectra, S_{λ} , to a lighting representation, $S_{i,OPT}$, specified in terms of the rendering primaries

The rendering pipeline is also intended for use with surface reflectance and lighting data captured with abridged multispectral imaging systems. The example in Figure 3 is based on a six-channel imaging system that combines an RGB camera with two additional filters, cyan and yellow. The reflectance data, $R_{\{RGBcRGBy\}}$, is represented in a form similar to spectral reflectance factor and determined by dividing captured surface data by a similarly illuminated white diffusing standard. Captured illumination data, $S_{\{RGBcRGBy\}}$, is represented in an absolute form proportional to radiance. Surface reflectance data and lighting data are converted to the rendering primaries using $[6 \times 6]$ matrix transforms, one for reflectance data (M_R) and one for lighting data (M_S), that are optimized for the specific camera system. Other spectral imaging systems, with different spectral sensitivities or a different number of channels (N) can be incorporated by providing a set of optimized $[6 \times N]$ matrices for that particular system.

Once all the data are represented in the set of six rendering primaries, surface shading calculations are performed by

multiplying surface reflectance and lighting on a per-channel basis and summing the results over all the light sources:

$$\Phi_{i,OPT} = \sum_{j=1}^L (R_{i,OPT})(S_{i,OPT,j}), \text{ for } i = 1 \text{ to } 6 \quad (2)$$

where $\Phi_{i,OPT}$ represents the total surface reflection for the i^{th} rendering channel under the set of L light sources. The final surface reflection, specified in the six rendering primaries, is multiplied by an optimized matrix to convert to CIE XYZ tristimulus values for display:

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} m_{1,X} & m_{2,X} & m_{3,X} & m_{4,X} & m_{5,X} & m_{6,X} \\ m_{1,Y} & m_{2,Y} & m_{3,Y} & m_{4,Y} & m_{5,Y} & m_{6,Y} \\ m_{1,Z} & m_{2,Z} & m_{3,Z} & m_{4,Z} & m_{5,Z} & m_{6,Z} \end{bmatrix} \begin{bmatrix} \Phi_{1,OPT} \\ \Phi_{2,OPT} \\ \vdots \\ \Phi_{6,OPT} \end{bmatrix} \quad (3)$$

To allow the calculations to be performed using a trichromatic rendering pipeline, the six-channel data can be separated into two sets of trichromatic data and the final XYZ results combined:

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} m_{1,X} & m_{2,X} & m_{3,X} \\ m_{1,Y} & m_{2,Y} & m_{3,Y} \\ m_{1,Z} & m_{2,Z} & m_{3,Z} \end{bmatrix} \begin{bmatrix} \Phi_{1,OPT} \\ \Phi_{2,OPT} \\ \Phi_{3,OPT} \end{bmatrix} + \begin{bmatrix} m_{4,X} & m_{5,X} & m_{6,X} \\ m_{4,Y} & m_{5,Y} & m_{6,Y} \\ m_{4,Z} & m_{5,Z} & m_{6,Z} \end{bmatrix} \begin{bmatrix} \Phi_{4,OPT} \\ \Phi_{5,OPT} \\ \Phi_{6,OPT} \end{bmatrix} \quad (4)$$

Selection of Rendering Primaries

An objective for the rendering pipeline was that it be based on a single set of six rendering primaries, and not ones customized for each particular lighting environment. A non-linear optimization was performed to select a general set of rendering primaries to provide color accuracy for a range of different illuminants.

The optimization was performed using 100 surface spectral reflectance factor curves from a database of artist materials [17] and a range of spectral power distributions for lighting: illuminant A; daylight illuminants: D40, D50, D65, D75, D90; fluorescent illuminants: F1 through F12; and the measured SPD of a D50 light booth. Full-spectral data were specified at 36 wavelengths from 380 to 730 nm. The XYZ tristimulus values for the 100 surfaces under the 19 illuminants were calculated spectrally using the 1931 standard observer, to serve as a baseline for comparison to the spectrally-abridged rendering workflow. The resulting 1900 XYZ values were converted to corresponding colors for D65 using the chromatic adaptation transform from CIECAM02 (MCAT02) before calculating CIELAB values.

The optimization was performed using the full-spectral data input path depicted on the left side of Figure 3. The spectral sensitivity curves for the six rendering primaries and the $[3 \times 6]$ matrix transform from the rendering primaries to XYZ were simultaneously optimized with constrained non-linear optimization in MATLAB (fmincon). To limit the number of parameters involved in optimizing the spectral curves for the primaries, each was represented as a Gaussian function with two parameters: peak wavelength and width (σ). As the Gaussian parameters were updated, $R_{i,OPT}$ and $S_{i,OPT}$ were recalculated and the results multiplied to determine $\Phi_{i,OPT}$ for each light-surface pair. XYZ values were estimated from $\Phi_{i,OPT}$ with the $[3 \times 6]$ matrix, then transformed to D65. The sum-squared error over the set of CIEDE2000 color differences, calculated between the abridged spectral results and baseline full-spectral results, was minimized by the optimization.




Figure 4. Spectral curves for the six Gaussian rendering primaries.

Results

The optimization resulted in the selection of six rendering primaries corresponding to Gaussians functions with peak wavelengths of (447.0, 481.5, 519.6, 543.1, 572.9, 622.4 nm) and width parameters of (16.9, 4.3, 7.8, 9.4, 15.5, 18.0 nm). The resulting set of spectral curves is shown in Figure 4.

The colorimetric accuracy of performing lighting calculations with the optimized set of rendering primaries and matrix transform was evaluated for the 100-patch target (S100) used in the optimization, and two additional targets, the 24-patch Xrite ColorChecker (CC) and 240-patch ColorCheckerDC (CCDC). The three targets were evaluated for the 19 illuminants used during the optimization, and also five measured real-world lighting environments: a cloudy outdoor environment ($x=0.317, y=0.329$), a desk with a lamp (0.456, 0.404), a lightbooth simulating horizon light (0.500, 0.417), an office environment (0.405, 0.404), and a classroom with fluorescent lighting (0.393, 0.391).

For comparison, results were also calculated using three-channels rendering methods, either directly in XYZ [10] or in a sharpened RGB rendering space [11]. In the XYZ lighting case, the three-channel representation of surface reflectance was calculated by assuming a D65 taking illuminant, calculating XYZ values under D65, and then dividing by the XYZ of the lighting. Surfaces were lit by multiplying this surface representation by the XYZ of the new light source. A similar methodology was used for the sharpened RGB space. For all methods, the final calculated XYZ values were transformed to corresponding colors under D65 and a CIEDE2000 color difference was calculated between the abridged lighting result and a full spectral calculation. Statistics for the three different methods are shown in Table 1.

Under the 19-light training set, the six-channel workflow was

Table 1: Comparison of Six-Channel, XYZ, and RGB Lighting

		Optimization Lighting				Measured Test Lighting			
		mean	s.d.	max	90th	mean	s.d.	max	90th
S100	6-Channel	0.5	0.5	3.6	0.9	0.9	0.7	5.0	1.7
	Sharp RGB	1.9	2.1	14.9	4.8	2.0	1.9	9.8	4.7
	XYZ	2.4	2.1	13.8	5.3	3.2	3.0	15.5	7.2
CC	6-Channel	0.6	0.6	3.1	1.2	0.8	0.6	2.8	1.5
	Sharp RGB	1.7	2.0	13.3	4.4	1.7	1.6	7.5	3.9
	XYZ	2.3	2.2	12.1	5.4	3.2	3.1	13.5	7.8
CCDC	6-Channel	0.5	0.5	4.9	1.0	0.8	0.5	4.4	1.4
	Sharp RGB	1.6	2.1	15.9	4.5	1.6	1.8	12.5	4.1
	XYZ	1.9	2.1	14.0	4.6	2.5	2.6	13.4	6.2

found to provide greater colorimetric accuracy than either of the three-channel methods, for both the S100 data and for the two test targets. Under the measured test illumination, the advantage relative to sharpened RGB was not as large as in the training illumination case, but the six-channel method still provided a higher level of accuracy, with mean color differences generally half the magnitude of those produced by the RGB method.

Simulation of Multispectral Input

The six channel rendering workflow is intended for use with multispectral-imaging data, in addition to full spectral data. To evaluate the capabilities of the workflow for this type of input, a simulation was performed using the spectral response curves from a real-world multispectral imaging system [14]. The imaging system, developed for use in artwork conservation, combines a Sinarback 54H color-filter-array camera with cyan and yellow filters, used sequentially to provide capture in six channels.

In the simulation, the six spectral response curves were used to calculate idealized six-channel camera signals for illumination ($S_{\{RGB\rightarrow RGB\}}$) and surface reflectance ($R_{\{RGB\rightarrow RGB\}}$) from spectral illumination and reflectance data. The two $[6 \times 6]$ matrix transforms required to convert from the simulated six-channel camera data to the six rendering primaries, \mathbf{M}_R and \mathbf{M}_S , were estimated using non-linear optimizations that minimized the CIEDE2000 of the final output of the pipeline for the S100 target data and the training illumination dataset.

Three multispectral input cases were evaluated in the simulation: (1) multispectral-imaging data on reflectance only (MSI-Reflectance) with the lighting data from the Gaussian spectral input workflow, (2) multispectral-imaging data on lighting only (MSI-Lighting) with reflectance data from the Gaussian spectral input workflow, or (3) both reflectance and lighting data from multispectral imaging (MSI-Both). Case 1 corresponds to the typical use of a multispectral-imaging system, to capture reflectance information over a surface (such as a painting), which is then re-lit using the spectral power distribution of an illuminant. Case 2 corresponds to capturing an illumination map of an environment with a multispectral system, which is then used light a uniform object of known spectral reflectance. Case 3 is a combination of these two types of multispectral capture.

The color accuracy results for the three multispectral input cases are shown in Table 2. For comparison, baseline results are provided for the case where full-spectral input data are processed through the optimized Gaussian curves for both illumination and reflectance (FS-Both). CIEDE2000 statistics on the S100 training

data, the ColorChecker, and ColorCheckerDC, under both the initial training lighting and the test set of five measured lighting environments, are shown. The results for rendering in the six optimized channels are also compared to a trichromatic rendering workflow, where multispectral input data were converted to sharpened RGB primaries with an optimized $[3 \times 6]$ matrix before lighting calculations were performed.

The simulation results indicate that case 1, where only the reflectance data is entered through the multispectral path, produced the highest overall accuracy of the three multispectral cases, with results similar to the baseline case of full-spectral inputs to the six-channel pipeline. In practice, with real captured data from a camera system, noise and quantization error would likely reduce the performance relative to the full-spectral input method. Multispectral input of the light sources (case 2) introduced a greater degree of color error into the rendering pipeline, and multispectral-imaging input of both gave the lowest accuracy. The mean and 90th percentile color difference results followed this pattern for nearly all the conditions, but a less consistent pattern was found for the maximum error, where in some instances (when reflectance increased sharply beyond 650 nm), the full-spectral input method exhibited larger maximum errors than some corresponding multispectral-input simulations.

For the optimization lighting, the six channel rendering workflow produced mean, max, and 90th percentile color errors that were generally 50% smaller than the CIEDE2000 values found for the RGB workflow. Under the measured test lighting, the six-channel workflow still produced higher accuracy than the three channel RGB calculations, though its benefit was not as large, in particular for the multispectral lighting case.

End-to-End System Prototype

A prototype has been developed to demonstrate the potential for the six-channel rendering approach described in the preceding sections. It implements a preliminary version of the full end-to-end pipeline including multispectral-imaging capture of real-world illumination, real-time GPU-based 3D rendering, and display to a colorimetrically-characterized monitor.

Multispectral-Imaging Input: Illumination Map Capture

A set of illumination maps were captured with a multispectral-imaging system to provide real-world data for use in the rendering workflow. A six-channel imaging system consisting of a modified Canon 60D camera with the IR-filter removed, shot sequentially through cyan and yellow filters, was used for data

Table 2: Comparison of Six-Channel and Sharpened RGB Rendering for Multispectral and Full-Spectral Input

(CIEDE2000)	Optimization Lighting (19)								Measured Lighting (5)								
	Optimized Six Channels				Sharpened RGB				Optimized Six Channels				Sharpened RGB				
	mean	s.d.	max	90th	mean	s.d.	max	90th	mean	s.d.	max	90th	mean	s.d.	max	90th	
S100	FS-Both	0.5	0.5	3.6	0.9	1.9	2.1	14.9	4.8	0.9	0.7	5.0	1.7	2.0	1.9	9.8	4.7
	MSI-Reflectance	0.7	0.4	3.4	1.2	1.7	1.4	9.9	3.7	0.9	0.5	2.9	1.7	1.9	1.5	9.2	4.0
	MSI-Lighting	0.8	0.6	3.4	1.6	2.0	1.9	13.4	4.8	1.4	0.9	5.2	2.7	2.1	1.7	8.7	4.6
	MSI-Both	0.9	0.6	3.7	1.7	2.6	1.2	10.3	4.2	1.4	0.9	5.7	2.7	2.6	1.2	8.9	4.1
CC	FS-Both	0.6	0.6	3.1	1.2	1.7	2.0	13.3	4.4	0.8	0.6	2.8	1.5	1.7	1.6	7.5	3.9
	MSI-Reflectance	0.6	0.4	2.3	1.1	1.6	1.3	9.1	3.6	0.9	0.5	2.7	1.5	1.8	1.4	6.8	4.0
	MSI-Lighting	0.9	0.7	3.4	1.9	2.0	1.8	12.8	4.3	1.5	1.0	5.1	2.7	2.1	1.4	7.5	3.6
	MSI-Both	1.0	0.6	3.5	1.7	2.6	1.2	8.6	3.9	1.5	1.1	5.5	3.0	2.8	1.0	6.6	4.2
CCDC	FS-Both	0.5	0.5	4.9	1.0	1.6	2.1	15.9	4.5	0.8	0.5	4.4	1.4	1.6	1.8	12.5	4.1
	MSI-Reflectance	0.7	0.5	6.0	1.2	1.6	1.4	10.6	3.4	0.8	0.5	4.3	1.5	1.6	1.3	11.9	3.4
	MSI-Lighting	0.9	0.7	4.4	1.9	2.0	1.9	15.4	4.5	1.7	1.1	5.2	3.2	2.1	1.6	12.5	4.1
	MSI-Both	1.0	0.7	5.3	2.0	2.6	1.2	9.9	4.0	1.7	1.2	5.7	3.5	2.6	1.2	12.0	4.1



Figure 5. Images of the mirror sphere captured through the cyan (left) and yellow (right) filters of the multispectral camera system.

acquisition. High-dynamic range illumination maps (shown in Figure 5) were captured by photographing a stainless steel sphere at a range of exposure times. The mirror-sphere images were geometrically transformed to cubic environment maps using the HDRShop software package [18].

The illumination map data were converted to the six rendering primaries with an \mathbf{M}_S matrix, then divided by an estimate of the spectral reflectance of the sphere, also represented in the rendering primaries, to give the final illumination maps. The $[6 \times 6]$ \mathbf{M}_S transform needed to convert from the camera's six-channel representation of the captured lighting, $S_{[RGBcRGBv]}$, to the optimized rendering primaries, S_{iOPT} , was estimated by imaging a target under five different light sources and measuring the corresponding spectral radiance of each patch with a spectroradiometer. These spectral radiance curves were converted to S_{iOPT} using the Gaussian input path and the $[6 \times 6]$ matrix was estimated with a pseudo-inverse calculation.

Rendering System

The prototype system was developed in OpenGL® to allow for rendering of the 3D computer graphics models at interactive frame rates using GPU-based hardware acceleration. The Ward BRDF Model [19] is used to simulate the material properties (glossiness and color) of the 3D objects:

$$\rho_{brdf}(\theta_i, \phi_i, \theta_r, \phi_r) = \frac{\rho_d}{\pi} + \rho_s \frac{1}{\sqrt{\cos \theta_i \cos \theta_r}} \frac{\exp(-\tan^2 \delta / \alpha^2)}{4\pi\alpha^2}. \quad (5)$$

where the ρ_d parameter represents diffuse reflectance, ρ_s represents the amount of specular reflectance, and α is a roughness parameter representing the width of the specular lobe. Spectral reflectance information is represented in the framework of the Ward model by providing a set of six ρ_d and ρ_s parameters, one for each of the six rendering primaries. The Ward model was selected because it is physically plausible and computationally simple enough to be rendered in real-time. However, the color-rendering pipeline is not contingent on the use of this model and could be implemented for a range of different rendering methods. Surfaces are illuminated



Figure 6. Rendering of a 3D model for a range of specular roughness properties using real-time importance sampling of environmental illumination.




Figure 7. Rendering of a virtual ColorChecker of glossy spheres in the captured lighting environment shown in Figure 5. Left, when facing the daylight from the window and right, when facing the incandescent lamp.

with cubic maps of the captured lighting environment using the real-time importance sampling technique developed by Colbert et al. [3], where the filtering level of the environment map is determined by the probability density function of the sample direction, and the Ward model sampling equations described by Walter [20]. The diffuse reflectance is calculated based on a uniform sampling, around the surface normal direction, of the filtered environment map. A teapot model rendered with a range of different surface specular roughness properties (Ward α parameter values) using real-time importance sampling of a six-channel environment map is shown in Figure 6.

Cubic environment maps are typically represented as trichromatic high-dynamic range images and used to perform rendering calculations in three channels. In our workflow, the six channels of illumination data are divided into two separate trichromatic environment maps. Using multi-texturing support in the OpenGL Shading Language, both maps can be used in the same fragment shader program. During importance sampling, the set of calculated lighting directions are used to index each of the maps and lighting values from each are applied only to the corresponding three channels of the surface model. As in Eq. (4), the results for each set of three primaries are multiplied separately by a $[3 \times 3]$ matrix transform to XYZ, and in the final stage the XYZ results from the two sets are combined. The final XYZ results are transformed for display to the screen with a colorimetric characterization model [21]. In the current prototype, the resulting images are not adjusted to compensate for chromatic adaptation.

Screen-captured results for a virtual 3D ColorChecker model, rendered in the mixed lighting environment of Figure 5, are shown in Figure 7. In Figure 8, rendered results are shown demonstrating the ability of the six-channel workflow to predict the metameric match between two models with different spectral reflectance curves, for the expected light source, in a captured environment.

Conclusions

In this paper, we have presented a multispectral framework for computer graphics rendering, designed to provide colorimetric accuracy while maintaining compatibility with real-time rendering methods that use complex illumination from real-world environments. The framework, based on an optimized six-channel representation, allows input data captured from different spectral measurement devices and multispectral imaging systems to be used interchangeably within the rendering pipeline. We described how the framework can be implemented for real-time rendering using a modern GPU-based graphics shader architecture and developed a prototype capable of rendering virtual objects at interactive frame rates using real-world illumination maps captured with a multispectral imaging system.

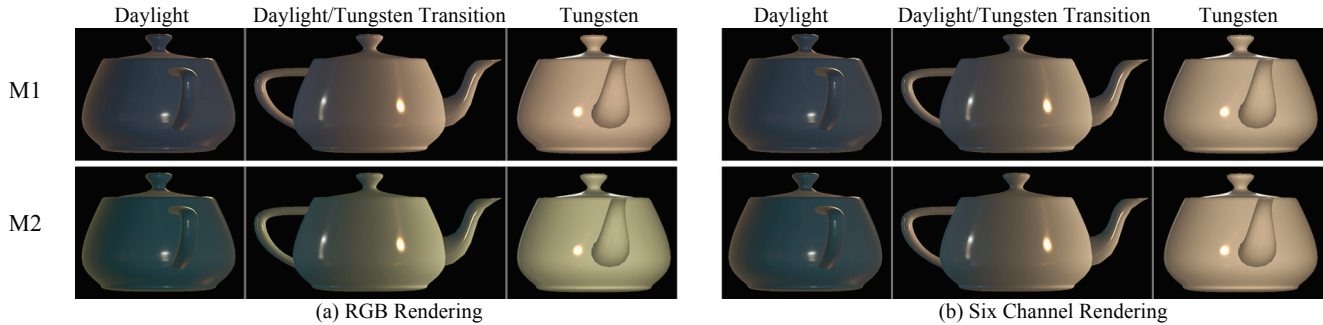


Figure 8. Two models (M1, top; M2, bottom) were assigned metameric spectral reflectance curves that are a colorimetric match under tungsten. In the captured lighting environment, the back of each teapot is lit primarily by a window and the front primarily by a tungsten desk lamp. With three-channel rendering (a), a color mismatch between the two models, expected for daylight, is still present for the tungsten lamp. With the six-channel workflow (b) maintaining additional spectral information, there is a color difference for the side facing daylight, and a metameric match for the side facing the tungsten lamp.

While the current results are promising, the framework still has limitations that suggest future work. In the simulation results, the six-channel pipeline performed well for full-spectral input and maintained much of its accuracy for multispectral reflectance input, but there was some decline in accuracy for multispectral illumination input. Since the simulations were all based on a multispectral imaging system designed for reflectance capture, it would be worthwhile to investigate whether incorporating a separate multispectral system specifically optimized for capturing illumination could improve the accuracy of the lighting input path.

There are also limitations in the current implementation of the prototype system that suggest directions for future work. Currently, the conversion from the multispectral camera signals to the optimized rendering primaries is a simple least squares minimization, but color accuracy could likely be improved by optimizing the transform to minimize overall color error. An additional limitation is that the illumination maps and the rendering pipeline are high-dynamic range, but currently the results are clipped to allow them to be shown on a low-dynamic range screen. In future work, we would like to incorporate real-time tone-mapping or present the results on a high-dynamic range display system. Finally, the rendering system does not account for the color appearance effects that result from the chromatic adaptation state of the observer. In future work, this could be addressed by incorporating a chromatic adaptation transform into the final stages of the rendering workflow. With continued development, a multispectral framework combining data capture and rendering has great potential for creating computer graphics systems that are both color-accurate and practical for providing complex environmental illumination in real-time applications.

Acknowledgments

This work was supported in part by NSF grant 0811680 to James A. Ferwerda, with additional support through a grant from the Andrew W. Mellon foundation to Roy S. Berns.

References

[1] D.P. Greenberg, K.E. Torrance, P. Shirley, J. Arvo, J.A. Ferwerda, S. Pattanaik, E. Lafortune, B. Walter, S. Foo, and B. Trumbore, "A Framework for Realistic Image Synthesis," *ACM Transaction on Graphics (ACM SIGGRAPH)*, 477 (1997).
 [2] P. Debevec, *Rendering Synthetic Objects into Real Scenes: Bridging Traditional and Image-based Graphics with Global Illumination and High Dynamic Range Photography*, *Proc. SIGGRAPH*, 189 (1998).

[3] M. Colbert, S. Pattanik, and J. Krivanek, "BRDF-Shop: An Artistic Tool for Creating Physically-Correct BRDFs," *IEEE Computer Graphics and Applications*, 26(1), 30 (2006).
 [4] G. Meyer, "Wavelength Selection for Synthetic Image Generation," *Computer Vision, Graphics, and Image Processing*, 41, 57 (1988).
 [5] M.S. Peercy, *Linear Color Representations for Full Spectral Rendering*. *Proc. ACM SIGGRAPH '93*, pg. 191. (1993).
 [6] M. S. Peercy, D.R. Baum, and B.M. Zhu, "Linear Color Representations for Efficient Image Synthesis," *Color Research & Application*, 21(2), 129 (1996).
 [7] M.S. Drew and G.D. Finlayson, "Multispectral Processing Without Spectra," *Jour. Opt. Soc. Am. A*, 20(7), 1181 (2003).
 [8] G.H. Joblove and D. Greenberg, "Color Spaces for Computer Graphics," *Computer Graphics (ACM SIGGRAPH)*, 12, 323 (1978).
 [9] R. Hall, *Illumination and Color in Computer Generated Imagery* (Springer-Verlag, New York, NY, 1989).
 [10] C.F. Borges, "Trichromatic Approximation Method for Surface Illumination," *Jour. Opt. Soc. Am. A*, 8(8), 1319 (1991).
 [11] G. J. Ward and E. Eydelberg-Vileshin, *Picture Perfect RGB Rendering using Spectral Prefiltering and Sharp Color Primaries*, *Proc. Eurographics Workshop on Rendering*, pg. 117. (2002).
 [12] G. M. Johnson and M.D. Fairchild, *Computer Synthesis of Spectroradiometric Images for Color Imaging Systems Analysis*, *Proc. IS&T 6th Color Imaging Conference*, pg. 150. (1998).
 [13] J. Y. Hardeberg, F. Schmitt, H. Brettel, J.-P. Crettet, and H. Maitre, *Multispectral Image Acquisition and Simulation of Illumination Changes*, In *Colour Imaging: Vision and Technology* (John Wiley & Sons, Chichester, England, 1999) pg. 145.
 [14] R.S. Berns, L.A. Taplin, M. Nezamabadi, M. Mohammadi, and Y. Zhao, *Spectral Imaging Using a Commercial Color-Filter Array Digital Camera*, *Proc. 14th Triennial ICOM-CC*, pg. 743 (2005).
 [15] S. Tominaga and T. Fukada, *Omnidirectional Scene Illuminant Estimation using a Multispectral Imaging System*. *Proc. SPIE Color Imaging*, 6493, 649313-1. (2007).
 [16] S. Tominaga and N. Tanaka, "Spectral Image Acquisition, Analysis, and Rendering for Art Paintings," *Journal of Electronic Imaging*, 17(4), 043022 (2008).
 [17] Y. Okumura, *Developing a Spectral and Colorimetric Database of Artist Paint Materials*, M.S. Thesis, Roch. Inst. of Tech. (2005).
 [18] USC Institute for Creative Technologies, (<http://www.hdrshop.com/>).
 [19] G. J. Ward, "Measuring and Modeling Anisotropic Reflection," *Computer Graphics*, 26, 265 (1992).
 [20] B. Walter, *Notes on the Ward BRDF*, Tech. Rep. PCG-05-06, Cornell University, (2005).
 [21] E. Day, L. Taplin, and R. Berns, "Colorimetric Characterization of a Computer-Controlled Liquid Crystal Display," *Color Research & Application*, 29(5), 365 (2004).