

Efficient Multispectral Reflectance Function Capture for Image-Based Relighting

Chloe LeGendre, Xueming Yu, Paul Debevec; USC Institute for Creative Technologies, Playa Vista, CA, USA

Abstract

Image-based relighting (IBRL) renders the appearance of a subject in a novel lighting environment as a linear combination of the images of its reflectance field, the appearance of the subject lit by each incident lighting direction. Traditionally, a tristimulus color camera records the reflectance field as the subject is sequentially illuminated by broad-spectrum white light sources from each direction. Using a multispectral LED sphere and either a tristimulus (RGB) or monochrome camera, we photograph a still life scene to acquire its multispectral reflectance field – its appearance for every lighting direction for multiple incident illumination spectra. For the tristimulus camera, we demonstrate improved color rendition for IBRL when using the multispectral reflectance field, producing a closer match to the scene’s actual appearance in a real-world illumination environment. For the monochrome camera, we also show close visual matches. We additionally propose an efficient method for acquiring such multispectral reflectance fields, augmenting the traditional broad-spectrum lighting basis capture with only a few additional images equal to the desired number of spectral channels. In these additional images, we illuminate the subject by a complete sphere of each available narrow-band LED light source, in our case: red, amber, green, cyan, and blue. From the full-sphere illumination images, we promote the white-light reflectance functions for every direction to multispectral, effectively hallucinating the appearance of the subject under each LED spectrum for each lighting direction. We also use polarization imaging to separate the diffuse and specular components of the reflectance functions, spectrally promoting these components according to different models. We validate that the approximated multispectral reflectance functions closely match those generated by a fully multispectral omnidirectional lighting basis, suggesting a rapid multispectral reflectance field capture method which could be applied for live subjects.

Introduction

Image-based relighting (IBRL) techniques [12, 4] allow computer graphics practitioners to compute the appearance of a subject under novel lighting conditions from how it appears in a set of basis illumination images. With image-based lighting environments [3], the scene can then be rendered to appear as it would in a real-world environment for realistic compositing into a background photograph. Typically, the appearance of the subject lit from each incident lighting direction, called a *4D reflectance field* [4], is recorded with an omnidirectional lighting sphere comprised of individual white LED light sources. However, white LEDs have emission spectra different from many target illumination spectra commonly found in the real world, such as incandescent or fluorescent illumination, or indirect illumination such as bounce light from vegetation or a brick wall.

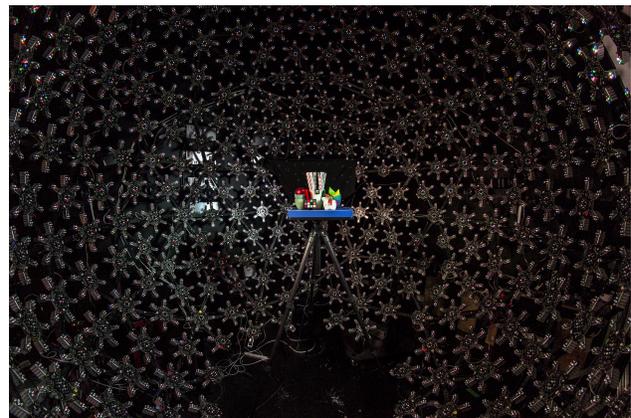


Figure 1. A wide-angle image peering into the "Light Stage X" LED sphere used in this work, with the still life we use for the results placed at the center.

The spectral mismatch between the light sources used to generate lighting basis images and the various direct and indirect light sources captured in image-based lighting (IBL) environments has largely been ignored, since the final values output by rendering systems are tristimulus RGB pixel values for displays. Typically, the lighting basis images are first white-balanced with a diagonal color matrix, and then the RGB pixel values of the environment lighting are used to weight the individual color channels of each basis image. This tristimulus scaling approach is prone to producing color errors in the final relit result, because it does not consider spectral rendition differences. In tristimulus imaging systems, RGB pixel values are produced by integrating over all wavelengths the fully spectral modulation of the reflectance spectrum of a material by the incident illumination spectrum and the camera spectral sensitivities. Accordingly, a simple white balance operation, or any linear color channel mixing, usually cannot completely correct for material color appearance mismatches across spectrally-different illumination environments.

As an extreme example, suppose that we wish to simulate the appearance of a blue-green material under monochromatic sodium illumination. During lighting capture, the 589nm sodium emission line will likely fall within the spectral sensitivity of both the camera's red and green pixels and produce a positive response for those pixels, but not the blue, producing a yellow color as expected. During reflectance capture, the blue-green material, lit by white LED light, would produce positive response for the green and blue pixels, but not the red. In the IBRL process, the RGB color of the light is multiplied by the RGB color of the material seen under white LED light, yielding zero for red and blue and a positive value for green, indicating green reflectance. However, this represents a significant color mismatch, since the color

of the monochromatic sodium light reflecting from the blue-green material should be, if anything, the yellow sodium color. Thus, performing IBRL in the RGB domain can in theory lead to significant errors in color rendition.

For more accurate color rendition with IBRL, intuition suggests merging the practiced techniques of reflectance field capture with multispectral imaging methods. If both the image basis set and the environment illumination map provided information beyond RGB pixel values, color rendition could theoretically be improved. Various multispectral imaging techniques could augment the RGB data – for instance, we could use an array of imaging filters placed in front of the camera to gain additional color channels for either case, or, for the image basis set, use a set of differently colored LEDs with multiple narrow-band spectral power distributions to illuminate the subject for each lighting basis or direction.

In this work, we record multispectral reflectance fields using an LED sphere (Fig. 1) comprised of six different LEDs of distinct spectra: white, red, amber, green, cyan, and blue (WRAGCB), spectra in Fig. 2. Once we have acquired a multispectral reflectance field, we can relight the subject according to the multispectral lighting conditions of any environment, presuming that the lighting environment has been rendered or captured in a multispectral manner. In practice, we reproduce a lighting environment using each of the multispectral light sources of the sphere, producing relative intensities of each spectral channel for each lighting direction using the methods of LeGendre et al. [8].

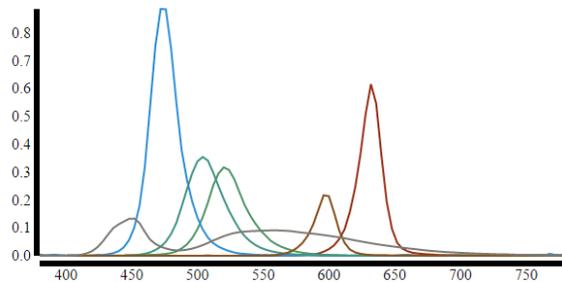


Figure 2. The spectra of the six LED sources used in this work, corresponding to white, red, amber, green, cyan, and blue LEDs.

A main advantage of the typical white light spherical lighting basis image acquisition process is the rapid capture time, which enables visual effects practitioners to acquire such image basis sets for live actors, so that digital doubles can be relit and realistically composited into any scene. Multiplying the number of lighting basis functions or lighting directions by the number of desired spectral channels, while theoretically desirable, would be too slow to practice for live subjects. To avoid lengthy acquisition times, we promote white-LED reflectance functions to multispectral reflectance functions, requiring only an additional n photographs more than for white light reflectance field capture, where n is the number of multispectral channels.

Our process is the following:

1. Acquire white LED reflectance field basis images $R(\theta, \phi, u, v, s = 0, c)$ where (θ, ϕ) is the incident lighting direction, (u, v) is the pixel coordinate in the image, s is the index of the LED color (where 0 indicates white LED light is used), and c is the color channel of the camera.

2. Acquire images of the subject lit by a full sphere of illumination $F(u, v, s, c)$ for each available LED color where (u, v) is the pixel coordinate, s is the index of the LED color where 0 is the white LED and 1.. n are the color LEDs, and c is the index of the color channel of the camera.
3. Promote each pixel (u, v) of each lighting direction basis image (θ, ϕ) to one which includes the response to all LED spectra as $R(\theta, \phi, u, v, s, c) = R(\theta, \phi, u, v, 0, c) * F(u, v, s, c) / F(u, v, 0, c)$.

If multispectral light sources are not available for each direction of the lighting rig, as would be the case with those designed to capture only white light reflectance fields, we also show a method of spectral promotion of reflectance fields basis images that uses only a head light positioned to illuminate the scene from the front. We also demonstrate spectrally promoting the diffuse and specular components of R independently and summing their promoted results. This requires photographing the diffuse and specular components separately, or estimating them from the reflectance functions.

To evaluate multispectral IBRL color rendition, we construct a still life scene with diverse materials of various reflectance spectra, and photograph it illuminated by a mixed illumination real-world environment consisting of diffused incandescent and fluorescent illumination. Using a fully multispectral reflectance field we demonstrate superior color rendition as compared to using the tristimulus-scaling approach. We also demonstrate that multispectral reflectance fields enable the use of monochrome cameras for IBRL. Even with our spectrally-promoted multispectral reflectance field, we show improved color rendition for the still life scene as compared with tristimulus-scaling IBRL, with very few additional images required at capture time.

In summary, the contributions of our paper are the following:

1. To the best of our knowledge, we show the first example of multispectral IBRL.
2. We demonstrate experimentally that using a multispectral reflectance field acquired with a tristimulus (RGB) camera improves IBRL color rendition, as compared with the tristimulus scaling of white light reflectance field basis images.
3. We show that multispectral IBRL enables the uses of monochrome cameras for reflectance field acquisition.
4. We propose a fast method of approximate multispectral reflectance field capture, by “colorizing” the white light reflectance field, and we experimentally assess the accuracy of the approximation.
5. We propose an alternative method of approximate multispectral reflectance field capture, where diffuse and specular components are spectrally-promoted with different models.

Background and Related Work

Image-Based Relighting

Debevec et al. [4] introduced the 4D Reflectance Field $R(\theta, \phi, u, v)$ to denote the image of a subject with pixels (u, v) as lit from any lighting direction (θ, ϕ) . In practice, they acquired such images with an RGB camera by spiraling an incandescent light source around the subject on a customized gantry. They also described the Reflectance Function of a pixel (u, v) as the spherical image of that pixel’s RGB value over all (θ, ϕ) . They then

showed that taking the dot product of the reflectance field with a spherical illumination map recorded as a high dynamic range (HDR), omnidirectional image as in the work of Debevec [3] (denoted as $L(\theta, \phi)$) effectively re-lights the subject using the lighting environment. However, this work restricted the discussion to RGB pixel values only, and thus is unlikely to be able to record or simulate the effect of spectrally complex illumination sources on spectrally complex subjects, e.g. fluorescent light on human skin.

Our current work attempts to efficiently extend this framework to multispectral reflectance fields. We build upon this process to extend the number of channels of these reflectance functions such that we can relight scenes with multispectral lighting environments. To represent a multispectral reflectance field, we will represent the response of the subject to a variety of LEDs of distinct spectra indexed by s , and make explicit that each lighting condition is sensed by the camera's set of spectral response channels (such as RGB) indexed by c . Thus, we will consider $R(\theta, \phi, u, v, s, c)$.

Multispectral Reflectance Capture

Darling et al. [2] explored multispectral image-based lighting for synthetic subjects in a real-time rendering context. They describe a six-channel multispectral environmental illumination capture method and real-time rendering workflow, demonstrating superior color matching results for their six-channel method as compared with XYZ-based rendering or SharpRGB-based rendering, with the target colors defined based on fully-spectral rendering at narrow-bandwidth resolution. Their method of multispectral environmental illumination capture extended the work of Debevec [3] to include capturing HDR photographs of a mirrored sphere as seen by a camera through cyan and yellow filters. The authors solved for the optimal spectra of the six-channel rendering primaries to best match the colors of a spectral materials database under a set of standard illuminants. Importantly, the primaries they define (Fig 4., [2]) closely resemble the LED spectra that we use in our multispectral spherical illumination rig, implying that we should be able to reproduce colors accurately for a wide variety of materials under diverse illumination conditions. The authors of this work do not approach the capture of reflectance functions; their work is limited to the acquisition of multispectral light probes and rendering synthetic objects with real-world illumination.

Park et al. [13] recorded multispectral reflectance images of a subject from a single lighting condition by subsequently illuminating them from a set different LED spectra. They showed that they could make the acquisition more efficient by activating multiple spectral sources (such as red and blue) at the same time, determining the best combination of sources to use. We illuminate our subject with each of the spectrally distinct sources, but only for the full-on lighting condition, and we propagate this spectral reflectance information to each of the individual lighting directions. Ma et al. [9] showed that a high-resolution RGB image could be spectrally promoted to a multispectral image, using a multispectral image acquired with a lower-resolution camera. This bears some similarity to our work, as we operate on the angular illumination domain, and propagate spectral information from a very low resolution image of the sphere (where all lights of each LED spectra are driven to the same brightness) and propagate these spectral relationships to all of the individual lighting directions.

Spherical or dome-shaped LED lighting rigs have been designed and constructed to illuminate a subject from all directions from a variety of spectral channels, for purposes including multispectral material measurement [6, 1, 7]. While intuition suggests improvement in color rendition using images from such a lighting rig for IBRL, to the best of our knowledge no prior work has endeavored to relight a subject using an omnidirectional, multispectral reflectance field.

Multispectral Lighting Reproduction

In our previous work [8], we demonstrated superior studio lighting reproduction using multiple spectral channels as compared with using only RGB LEDs. However, this result is not directly applicable to IBRL, because the problems are not exactly analogous. RGB LED lighting reproduction yields poor color rendition for various materials due to the lack of energy produced across the visible spectrum when using just RGB LEDs – which is a problem mostly avoided in IBRL when using broad-spectrum white light for reflectance field acquisition. However, most broad-spectrum white LEDs, including those we use in our LED sphere, still have a spectral gap between the shorter wavelength blue emitter part of the spectrum and longer wavelength phosphor-converted part. Additionally, the spectral mismatch problem described above is still a concern for IBRL.

Method and Equations

In this section we describe our techniques for acquiring a fully multispectral reflectance field of a subject and producing an IBRL rendered result using such a dataset. We also describe methods for promoting white light reflectance functions to multispectral and for radiometrically and colorimetrically calibrating the various parts of our acquisition system.

Multispectral Reflectance Field Capture

In our multispectral LED sphere, we acquire the appearance of a subject lit by every lighting direction for every available LED of distinct spectra (WRAGCB). Our LED sphere includes 346 light sources surrounding a subject, with 30 light sources spaced around the equator with a separation of 12° , with all lighting directions accounted for excluding a small portion near the bottom of the sphere. Each light source includes one LED of each spectral channel, with the LEDs spaced around a 3 cm diameter circle. The diameter of the LED sphere is 2.7 m, so the angular difference between the LEDs of each spectral channel per light source is small. While we use point lights coming from a dense set of lighting directions on a sphere as our lighting basis, we could also employ any of the basis conditions typically used in IBRL, such as Hadamard patterns or the spherical harmonic basis.

Spectral Promotion of Reflectance Functions

Although our LED sphere allows for the capture of a fully multispectral reflectance field (each LED color for each lighting direction), we would ideally like to minimize the number of images required without sacrificing color rendition or angular resolution. Towards this goal, we evaluate several methods for promoting a white light reflectance field (just the white LEDs for each lighting direction) to a multispectral reflectance field, for both a tristimulus (RGB) and monochrome camera. This is especially notable for the monochrome camera case, since effectively we

need to “colorize” the monochrome images lit from each direction in order to generate a color IBRL image.

Our general technique for spectral promotion is motivated by the fact that a diffuse, Lambertian surface will reflect the same spectrum of light toward the camera – up to a scale factor – no matter which direction it is lit from as long as the illuminant remains the same. Thus, we can measure the spectral reflectance of a point (u, v) on the subject from a reference lighting condition A, and measure only the brightness of the point as lit by a second condition B, and surmise that the spectral reflectance of the subject lit by condition B is the spectral reflectance of the point lit by condition A scaled to have the brightness as when lit by B. We could sensibly choose the reference lighting condition to be a full sphere of even-intensity illumination, since this will minimize (though not eliminate) the appearance of specular reflectance. Thus, if $F(u, v, s, c)$ is the appearance of the subject at each pixel (u, v) under a full sphere of illumination for each LED spectrum s for each camera sensitivity function c , we can write:

$$R(\theta, \phi, u, v, s, c) = R(\theta, \phi, u, v, 0, c) * \frac{F(u, v, s, c)}{F(u, v, 0, c)} \quad (1)$$

However, since many LED sphere systems have only white LEDs, it might not be possible to produce full spheres of differently colored illumination $F(u, v, s, c)$. In this case, we can use a variant of this method of spectral promotion by adding a single multispectral “head light” to illuminate the scene from the front with each of the desired spectral channels, ideally as a ring light around the camera to minimize shadowing. Thus, if $H(u, v, s, c)$ is the appearance of the subject at each pixel (u, v) under a head light of each LED spectrum s , we substitute $H(u, v, s, c)$ for $F(u, v, s, c)$ and $H(u, v, 0, c)$ for $F(u, v, 0, c)$, writing:

$$R(\theta, \phi, u, v, s, c) = R(\theta, \phi, u, v, 0, c) * \frac{H(u, v, s, c)}{H(u, v, 0, c)} \quad (2)$$

Specular reflections from the head light are more likely to affect the estimation of the subject’s diffuse spectral reflectance, although cross-polarizing the head light could conceivably alleviate the issue. Additionally, by lighting the scene from only from the front direction, the interreflections observed in the scene will be specific to the frontal lighting direction, rather than to the more general condition of omnidirectional environmental illumination. Our results shown later indicate that both methods are generally successful, but produce different errors in interreflection regions.

Spectral Promotion with Diffuse and Specular Separation

Spectral promotion with Eq. 1 assumes that the materials in the scene are Lambertian and ignores the effects of specular reflections. For a particular lighting direction, the spectrum of the light reflected off a non-Lambertian object towards the camera could be the color of the light source for dielectrics or a different color for conductors. For each lighting direction, if we can separate the diffuse and specular components of the reflectance, then we could apply a different spectral promotion model to each component and sum the results, potentially producing a closer match to the ground truth multispectral reflectance field.

Separating diffuse and specular components could in theory be done in several ways, including polarization difference imaging as in Ma et al. [10] or Ghosh et al. [5], spherical harmonic

frequency separation [16], chromaticity analysis [15, 4], or a combination of these approaches e.g. Nayer et al. [11]. In this work, we use polarization difference imaging using vertically and horizontally polarized LEDs as in Ghosh et al. [5].

In addition to the unpolarized multispectral LEDs, each light source in our sphere includes a ring of twelve white LEDs, with six polarized vertically and six polarized horizontally, allowing us to compute the specular and diffuse components of each white light reflectance function as in Ghosh et al. [5]. We can promote the diffuse component of the white light reflectance functions with Eq. 1, producing $R_D(\theta, \phi, u, v, s, c)$, where the subscript D denotes the diffuse component. To compute the spectrally-promoted specular component $R_S(\theta, \phi, u, v, s, c)$, where the subscript S denotes the specular component, if $p(s, c)$ is the pixel color of a spectrally-flat object illuminated by each LED spectrum s for the camera sensitivity function c we can write:

$$R_S(\theta, \phi, u, v, s, c) = R_S(\theta, \phi, u, v, 0, c) * \frac{p(s, c)}{p(0, c)} \quad (3)$$

Effectively, with Eq. 3, we tint the white light reflectance function specular component $R_S(\theta, \phi, u, v, 0, c)$ by the color of light source of each LED spectrum s , assuming that the scene materials are dielectrics. This prevents the specular reflections of dielectric materials from becoming tinted by the underlying diffuse color of the material.

We can further improve the spectral promotion of the specular component for the tristimulus (RGB) camera by using chromaticity analysis to separate materials reflecting light of the same color as the incident illumination from those reflecting light of a different color than the incident illumination. As the color of light reflected specularly from conductors closely matches that of the diffuse component of the material, we use Eq. 1 to spectrally promote the parts of white light specular reflectance functions that do not exhibit behavior typical of dielectrics, generating $R_{SC}(\theta, \phi, u, v, s, c)$, with the subscript SC denoting a specular conductor. In practice, we compute both $R_{SC}(\theta, \phi, u, v, s, c)$ and $R_S(\theta, \phi, u, v, s, c)$, and produce a total spectrally-promoted specular reflectance function by linearly interpolating between the images produced by the two models. To control the interpolation, for the white incident light we compute the pixel color ratios $\frac{R_w}{G_w}$ and $\frac{R_w}{B_w}$ from $p(0, c = 1, 2, 3)$. For each pixel of a specular white light reflectance function, we compute the pixel color ratios $\frac{R}{G}$ and $\frac{R}{B}$. If $|\frac{R_w}{G_w} - \frac{R}{G}| = 0$ and $|\frac{R_w}{B_w} - \frac{R}{B}| = 0$, then the color of the specular reflectance function for that pixel is the same as the light source color, and so we spectrally promote this pixel using $R_S(\theta, \phi, u, v, s, c)$. If $|\frac{R_w}{G_w} - \frac{R}{G}| > t$ and $|\frac{R_w}{B_w} - \frac{R}{B}| > t$, where t is a threshold that we empirically set to 0.3, then we use $R_{SC}(\theta, \phi, u, v, s, c)$, interpolating in all other cases. Exploring other methods of separating conductors from dielectrics in specular reflectance functions is of interest for future work.

Multispectral Image-Based Relighting

Once we have a multispectral reflectance field, through exhaustive imaging with the omnidirectional multispectral lighting basis or through one of the spectral promotion techniques, we can relight the object with a multispectral lighting environment. We capture real-world multispectral lighting environments by photographing a mirrored sphere and an arrangement of five “Nano” ColorChecker™ charts from Edmund Optics using traditional

HDR photography with an RGB camera. Using a metameric reflectance matching approach as in our previous work [8], we independently solve for the optimal relative intensity of each spectral channel for each lighting direction, matching color chart appearances for each narrow cone of the lighting environment.

The target pixel values of a color chart for one cone of the lighting environment are represented by pixel values P_{ic} where i is the index of the given color chart patch and c is the camera's c 'th color channel. We record the appearance of the color chart as lit by a full sphere of each of the spectral channels of the LED sphere, and construct a matrix \mathbf{L} where L_{ics} is the averaged pixel value of color chart square i for camera color channel c under LED spectrum s . We consider \mathbf{L} to be the $ic \times s$ matrix whose columns correspond to the LED spectrum s and whose rows unroll the indices i and c to place the RGB pixel values for all chart squares in the same column. To optimally reproduce the target chart appearance, we solve for the LED intensity coefficients α_s that minimize Eq. 4, where m is the number of color chart patches and n is the number of different LED spectra. Rather than driving the LEDs to directly illuminate a subject in a studio environment [8], we use these intensities as scaling factors for the measured (or spectrally-promoted) multispectral reflectance field basis images.

$$\sum_{i=1}^m \sum_{c=1}^3 (P_{ic} - \sum_{s=1}^n \alpha_s L_{ics})^2 = \|\mathbf{P} - \mathbf{L}\alpha\|^2 \quad (4)$$

If a 3×3 color matrix \mathbf{M} is to be applied to the IBRL result for display, and \mathbf{M} unequally weights the color channels, then the quantity to minimize of should be modified, as both the target values \mathbf{P} and the reproduced values $\mathbf{L}\alpha$ will change, and errors for all color channels should be weighted equally in the minimization. Before constructing \mathbf{P} and \mathbf{L} , we apply the color matrix \mathbf{M} to the RGB triples, producing \mathbf{P}_M and \mathbf{L}_M .

$$\operatorname{argmin} (\|\mathbf{P}_M - \mathbf{L}_M\alpha\|^2) \quad (5)$$

When using a monochrome camera, LED intensity coefficients α_s must be computed separately for each color channel, yielding $3s$ degrees of freedom for each lighting direction. Eq. 4 is modified for the monochrome camera case in Eq. 6. N_{is} is the average pixel value of color chart square i under LED spectrum s , and \mathbf{N} is the $i \times s$ matrix whose columns correspond to the LED spectrum s and whose rows correspond to the color chart square i . Eq. 6 should be minimized for each color channel of the target chart, producing in our case three α vectors, which are the lighting primaries for that direction of the lighting environment.

$$\sum_{i=1}^m (P_i - \sum_{s=1}^n \alpha_s N_{is})^2 = \|\mathbf{P} - \mathbf{N}\alpha\|^2 \quad (6)$$

In the monochrome camera case, as the relative intensities of the spectral channels are optimized independently for each color channel of the target chart, applying a 3×3 color channel mixing matrix does not impact the minimization.

Measurement Setup and Calibration

We use a tristimulus (RGB) Ximea MQ042CG-CM machine vision camera with 50mm Fujinon lens and a vertical linear polarizer filter for all measurements, including: the capture of a real-world lighting environment mirrored sphere and color chart photographs, a reference photograph of a subject in this real lighting

environment, the subject's multispectral reflectance field and multispectral full sphere illumination conditions, and its horizontally and vertically polarized white light reflectance fields. This Bayer pattern camera produces linearly encoded 8-bit images of resolution 2048×2048 at frame rates up to 72 frames per second, and its image acquisition can be triggered by an external signal such that the sensor is precisely exposed to each basis lighting condition individually. Since our experimental still life scene contains highly reflective conductors as well as dark regions in shadows, we photographed the lighting environment, the appearance of the still life in the real lighting environment, and all reflectance fields at multiple exposure times, assembling them into HDR images [3], and scaling each to the correct relative radiance. We measured the radiometrically linear range of pixel values for this camera to be from 0.1 to 0.85 on the scale of 0.0 to 1.0. The camera has a non-zero black level, so we subtract an averaged dark current image from each frame for each exposure time prior to demosaicing.

Although the Lumileds Luxeon Rebel ES multispectral LEDs of our LED sphere are binned in a quality control process to ensure that all of the LEDs of the same spectral channel have a similar emission spectrum and overall radiance, we observed LED-to-LED variation in radiance that was greater than measurement-to-measurement or day-to-day variation. Accordingly, we multiply each color channel of the reflectance field basis images by a unique brightness scaling factor per LED, such that each image represents the appearance of the scene as lit by a uniform amount of metameric light from each direction. In practice, calibrating the relative radiance of 346×8 lighting conditions (346 light sources, \times 6 spectral channels, plus the horizontally and vertically polarized LED groups) in a spherical arrangement is a non-trivial task. We compute these scaling factors by placing a diffuse, spectrally-flat 33% reflective gray sphere at the center of the LED sphere and photographing it with the Ximea camera as illuminated by each multispectral and polarized basis lighting condition. For each lighting basis, we then compute the relative light source intensity for each image color channel using photometric stereo [17], assuming the surface normals follow the shape of a sphere. Since the light sources aim in from all directions, we used four cameras placed 90° apart around the equator to photograph the sphere under each basis lighting condition, interpolating between the photometric stereo's light intensity values from each camera according to each light's distance from the camera.

Given our LED arrangement, our full-sphere illumination conditions should produce about $90\times$ more light than a single point light source illuminating the scene from the front. In order to properly expose these conditions with the same exposure bracketing used for the lighting basis images, we dim the LEDs for the full sphere condition using pulse-width modulation (PWM). To achieve the correct relative radiance for the full sphere spectral promotion of reflectance fields, we need to ensure that $\frac{F(u,v,s,c)}{F(u,v,0,c)} = \frac{R(u,v,s,c)}{R(u,v,0,c)}$. For each spectral channel, we sum the scaled lighting basis images of the diffuse gray sphere, comparing the value with PWM dimmed full sphere condition. We observed that summing across the basis images for each lighting direction produced a diffuse sphere image that was on average $1.2\times$ brighter than the photographed full sphere condition for the WRGCB and polarized LEDs, and $1.14\times$ brighter for the amber LED, after adjusting for the PWM, due to the increased current draw required

to power all the light sources simultaneously.

We measured the reflectivity of our mirrored sphere using the technique in Reinhard et al. [14], ch. 9, and found it to differ slightly for each color channel as ($R = 60.7\%$, $G = 61.6\%$, $B = 63.4\%$), so we scale the image of the captured mirrored sphere image of the target lighting environment accordingly to the same relative radiance per color channel.

As our experiments require that we not move the camera, we could not test multispectral IBRL for both a tristimulus and true monochrome camera simultaneously as desired. Instead, after acquiring the reflectance fields and full sphere illumination conditions with our tristimulus camera, we compute artificial monochrome images as $M = 0.986R + 0.368G + 0.578B$, using the pixel values after dark current subtraction and demosaicing, which we measured as the best color channel mix to match the appearances of a color chart illuminated by daylight for the tristimulus camera and the comparable monochrome Ximea camera.

Results

We quantitatively and qualitatively compare the multispectral and tristimulus-scaling IBRL results for photographs of a still life scene in a real lighting environment, using both an RGB and (simulated) monochrome camera. We also compare the IBRL results using several methods of multispectral promotion to results achieved using the multispectral omnidirectional lighting basis.

For these comparisons, we constructed a still life scene containing various materials with diverse reflectance spectra, including colorful paper straws, a red metallic sphere, prism-shaped wax crayons, and twelve fabric swatches arranged as a chart. We also included a “Nano” ColorChecker™ chart from Edmund Optics in the scene, and a paper cup flanked by bright red and green squares of felt to elucidate the rendering of interreflections by each method for reflectance field spectral promotion.

The recorded real-world lighting environment was constructed to include incandescent and fluorescent area lights placed up against the left and right sides of the LED sphere, allowing both the appearance of the still life in the real lighting and the multispectral reflectance field to be captured without moving the still life or the camera. For visualization, we show in Fig. 3 the appearance of the lighting environment reflected off two spheres of different reflectivities (chrome and black acrylic), and the appearance of color charts facing in five different directions.

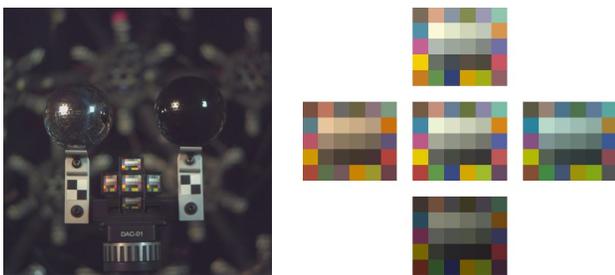


Figure 3. **Left:** The lighting environment observed as reflected in chrome and black acrylic spheres. **Right:** Sampled colors from the five color charts placed in the lighting environment.

For all comparison images we apply one consistent color matrix $M = [1.768, -0.187, -0.329; -0.461, 1.780, -0.525; -0.139,$

$-0.845, 2.882]$ to convert camera RAW values to sRGB for display, but we apply no image-specific color correction. We used M in the lighting environment optimization as described in Eq. 6. We applied no brightness scaling to any images of the still life in this section, although we multiplied the pixel values of the color charts with one overall brightness scaling factor for the qualitative and quantitative color chart matching assessments.

Color Rendition of Multispectral IBRL

We recorded the lighting environment using the multispectral lighting capture technique of [8] and photographed the still life from the front using a Ximea MQ042CG-CM RGB machine vision camera as lit by this environment. We then photographed the still life lit by each each of the six spectral channels of each of the 346 point light sources in the LED sphere to acquire the still life’s *multispectral reflectance field*. One multispectral basis lighting condition example is shown in row C of Fig. 8, as seen by the RGB tristimulus camera. We also simulated how a monochrome camera might have observed the multispectral reflectance field by converting the RGB images to grayscale as previously described.

To synthetically illuminate the scene with the recorded lighting environment based on its recorded multispectral reflectance field, we solved for the spectral channel intensities for each direction of the lighting environment with Eq. 5 for the tristimulus camera and the lighting primaries for each direction of the environment for the monochrome camera with Eq. 6. The lighting basis elements were then weighted by the computed channel intensities and summed. For comparison, we also performed traditional RGB image-based relighting [4] by scaling the RGB channels of the white-balanced white-LED reflectance field basis images according to an RGB light probe image.

We compare the color rendition of each approach in Fig. 4. At first glance, all three approaches (tristimulus-scaling IBRL and multispectral IBRL with both color and monochrome cameras) produce results which are similar to the ground truth lighting. But, on closer inspection, both the (simulated) monochrome camera IBRL and the white light only tristimulus-scaling approach produce color mismatches, as observed in the close-up inset of the colorful straws, and in the color charts below them, both shown in Fig. 4. In these charts, the foreground dots represent the measured values in the IBRL results, and the background squares represent the chart colors measured in the real lighting environment. When the foreground dots visually fade into the background squares, IBRL produced good color rendition. Multispectral IBRL with the color camera visually produces the best color rendition and quantitatively reduces color error for both the fabric and color chart squares for the red and green channels (Fig. 5). We observed with all three IBRL approaches some amount of haze in the dark regions of the images, which we believe to be the result of dark current which was not entirely compensated for. We suspect that this haze is responsible for the lack of improved color rendition for the blue channel for multispectral IBRL, as the haze in the multispectral image produced with the color camera appears blue.

Spectral Promotion of Reflectance Fields

In addition to acquiring the multispectral reflectance fields, we also photographed the scene lit by a full even sphere of illumination for each LED spectrum (Row A of Fig. 8). We generated

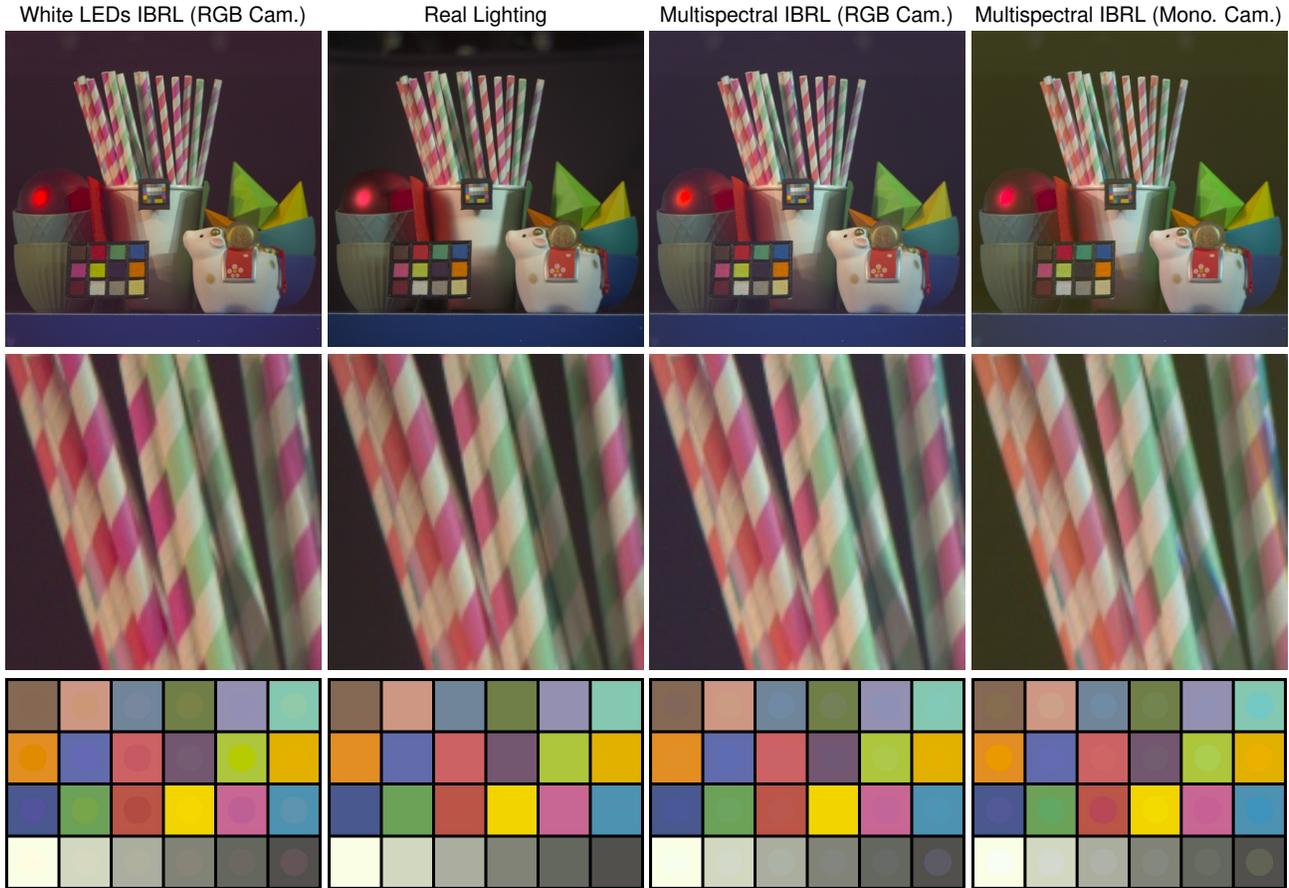


Figure 4. Row 1. Comparison among three different methods of image-based relighting (IBRL) for a still life scene illuminated by a complex multispectral lighting environment and the real photographed appearance of the scene under the actual illumination. Left: IBRL result using white LEDs only for capturing the reflectance field. Second from the left: photograph of the scene lit by the real-world illumination. Second from the right: Multispectral IBRL using six spectral LED channels to capture the multispectral reflectance field, as recorded by a color camera. Right: Multispectral IBRL again using multispectral LEDs, as observed by a monochrome camera. **Row 2.** Close-up insets of the images of Row 1, demonstrating superior color rendition with Multispectral IBRL with a color camera. **Row 3.** After sampling the pixel values from the color charts centered in the scene, we scale the charts to the same overall brightness with a single factor. The foreground dots represent sampled values from the images of the same column, while the background squares represent the target colors from the color chart in the real mixed-illumination lighting environment. Good color rendition is indicated when the foreground dots visually fade into the background squares. (For the real lighting environment color chart, the foreground and background colors would be the same.)

a similar set of images for our simulated monochrome camera. For both the color and monochrome cameras, we promoted the white LED reflectance field to multispectral using Eq. 1, hallucinating the appearance of the still life scene as illuminated by every spectrum of light from every direction. We show an example spectrally-promoted basis image for each spectral channel in Row D of Fig. 8, and, in row E, the absolute value of the difference between the full sphere spectrally-promoted basis image and the ground truth image for each spectral channel for the same lighting direction of the sphere. Using the spectrally-promoted reflectance fields, we then compute the IBRL results using the multispectral relighting equations for color and monochrome cameras, respectively Eq. 5 and Eq. 6. We show the full sphere reflectance field spectral promotion IBRL renderings in the second row of Fig. 10, for comparison with the results computed using the measured multispectral reflectance field (first row of Fig 10). We also show the color differences resulting from the full sphere spectral pro-

motion in the first row of color charts of Fig. 6. The background squares show the sampled colored charts from the multispectral reflectance field IBRL result, and the foreground dots represent the sampled charts from the spectral promotion.

To evaluate the head light spectral promotion, we computed a virtual head light image as the sum of three multispectral reflectance field basis images for the lighting directions closest to our camera. For both the color and simulated monochrome camera, we promoted the white LED reflectance field to multispectral using Eq. 3, again hallucinating our scene's appearance as illuminated by each spectrum of light from each direction. We show an example of such a spectrally-promoted basis image of the reflectance field for each spectrum in Row F of Fig. 8, and the absolute value of the difference between the spectrally-promoted basis image and the actual image for that lighting condition in Row G of Fig. 8. We produce the color and monochrome IBRL results with the head light promoted multispectral reflectance field using

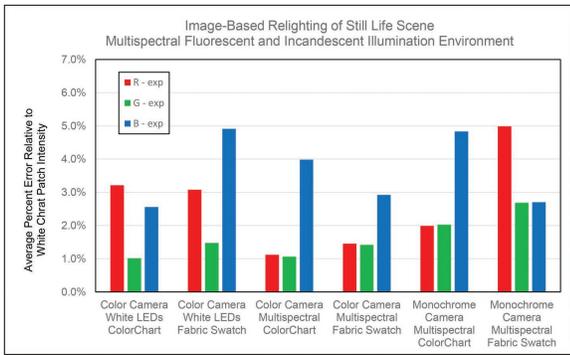


Figure 5. Quantitative experimental error plot for Fig. 4, computing the color error for the three IBRL methods, based on the average squared errors of linear RGB pixel values, with the color matrix M first applied to both the target chart in the real lighting environment and the IBRL results. We report average error from all 30 ColorChecker nano patches, and, then, separately, the twelve fabric swatches from the 3×4 chart in the scene, relative to the white square pixel value of the color chart in the original lighting environment.

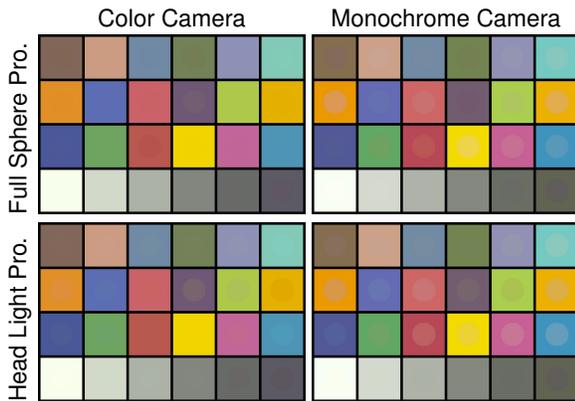


Figure 6. Background squares represent the sampled values from the IBRL results using the full multispectral reflectance field. **Row 1.** Foreground dots represent sampled color chart values from the full sphere spectral promotion IBRL results for each camera. **Row 2.** Foreground dots represent sampled color chart values from head light promotion IBRL results.

Eq. 5 and Eq. 6. We show the head light promoted IBRL renderings in the third row of Fig. 10 and the color differences resulting from this method of spectral promotion in the second row of color charts of Fig. 6, where the background squares again show the sampled colored charts from the multispectral reflectance field IBRL result and the foreground dots represent the sampled color charts from this method of spectral promotion.

In Fig. 7, we report the quantitative error for spectral promotion for the fabric swatches and color chart squares, by comparing the sampled color values from the IBRL result using the different techniques of spectral promotion with the results that could be achieved using the full, measured multispectral reflectance field. The head light approach does introduce color rendition error, and we observed an overall brightness increase for the color squares relative to the white square intensity (Fig. 6), likely owing to the increased amount of light reflected specularly from the not-entirely-diffuse color chart in this lighting configuration.

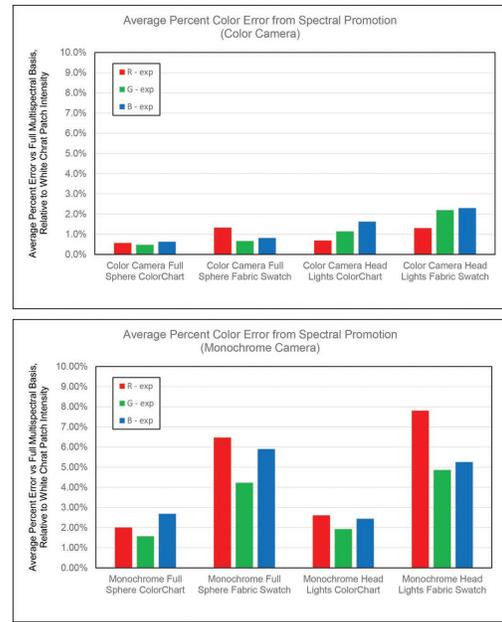


Figure 7. Color error for sampled pixel values from the fabric and color charts, comparing the multispectral IBRL results computed using the spectrally-promoted reflectance fields for the full sphere and head light methods. We evaluate both techniques for color rendition by comparing with the IBRL result using the measured multispectral reflectance fields. We report average error from all 30 ColorChecker nano patches, and, then, separately, the twelve fabric swatches from the 3×4 chart in the scene, relative to the white square pixel value of the color chart rendered using the multispectral IBRL. **Upper:** Color Camera. **Lower:** Monochrome camera.

The average color error for the fabric and color charts rendered using the full sphere spectrally-promoted reflectance field for the color camera is 0.7% (Fig. 7), indicating that full sphere promotion is a viable method that can yield the color rendition improvement of multispectral IBRL, while requiring only a few additional photographs for the lighting basis. As the first row of color chart results of Fig. 6 only compared the full sphere promotion with the color charts rendered by the fully-multispectral IBRL, in Fig. 9 we also compare the full sphere promoted multispectral IBRL result with the sampled color chart pixels from the real environment. We show that even with this approximation step, we achieve superior color rendition as compared to the existing method for this spectrally-complex real-world environment.

In the case of the monochrome camera imagery on the right of Fig. 10, it is apparent that spectral promotion can introduce color errors in the presence of interreflections. In the real image at the top, there are strongly hued shadows on the sides of the central cup from the bounced light from the red and green felt cards to its sides. In the full sphere promotion image in the middle row, these interreflections are somewhat muted, since in the full sphere image the shadows are filled in with more direct white LED light, and thus are less saturated. The situation is even worse for the headlight condition, where the shadows are very desaturated. This results from the headlight illuminating the cup well, but not the red and green felt squares, which are nearly perpendicular to the lighting direction, so that little bounced light is

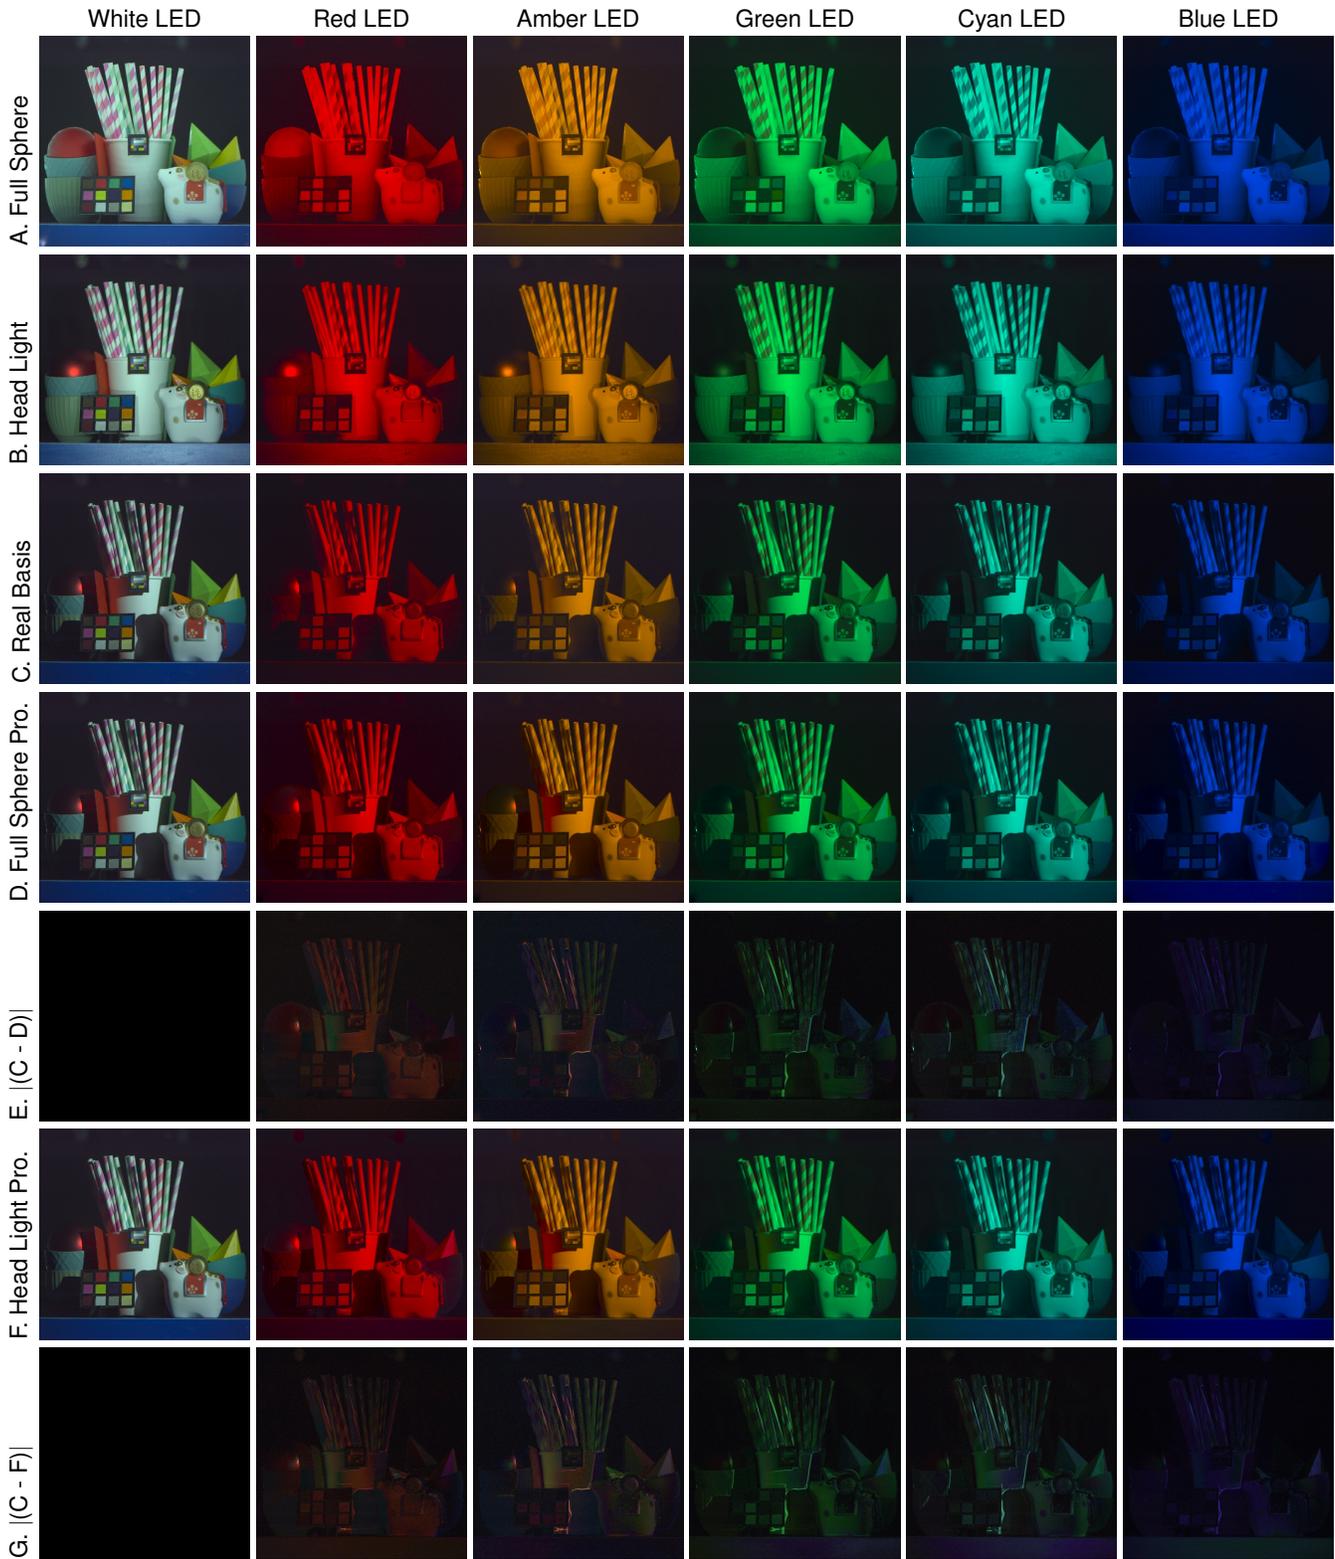


Figure 8. **Row A.** Full, even sphere of illumination of the scene with each spectral channel of the LED sphere. **Row B.** Head light illumination of the scene with each spectral channel. **Row C.** Example measured multispectral reflectance field basis lighting condition images for each spectral channel. **Row D.** Example full sphere white light promoted multispectral reflectance field basis lighting condition images, computed using Eq. 1 (White - actual image; Red, Amber, Green, Cyan, and Blue - Virtual images). **Row E.** The absolute difference of the measured and full sphere promoted multispectral reflectance function example basis images. **Row F.** Example head light white light promoted multispectral reflectance field basis lighting condition images, computed using Eq. 2 (White - actual image; Red, Amber, Green, Cyan, and Blue - Virtual images). **Row G.** The absolute difference of the measured and head light promoted multispectral reflectance function example basis images.

received by the cup. As a result, the shadows are promoted using relatively desaturated lighting estimates and do not match the coloration of the correct shadows at all well.

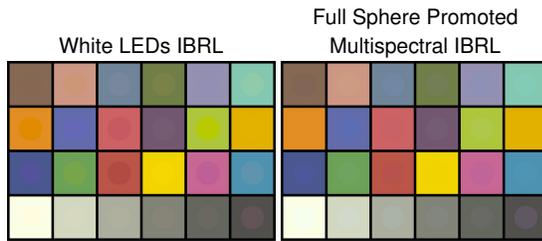


Figure 9. Foreground dots represent sampled color chart values from the IBRL results, while background squares represent the sampled values from the chart in the real lighting environment.

Diffuse/Specular Separation Spectral Promotion

We also acquired the reflectance field of the still life scene as it was illuminated by each of the polarized light sources, such that we could spectrally promote the diffuse and specular components of the reflectance field according to the different proposed models. While the IBRL result (Fig. 12i) did produce color rendition results visually comparable to those of the other methods of spectral promotion, we did not observe an overall benefit for the still life scene in this particular lighting environment, likely because our scene did not contain many dielectric materials, and the light sources of the real lighting environment were relatively diffused and oriented in a way that minimized the appearance of specular reflections relative to the camera position.

Nonetheless, we show in Fig. 12 that promoting the diffuse and specular components of the reflectance fields with different models can reduce color error for the individual basis lighting conditions, where we do observe specular reflections for our scene. Fig. 12g shows the result of combining the spectrally-promoted diffuse, specular conductor, and specular dielectric classified pixels of this lighting basis image separately, where Fig. 12h is the absolute value of the difference between Fig. 12g and the still life appearance when really lit by the amber LED for that lighting direction (Fig. 12j). Fig. 12k shows the result from the spectral promotion using only the full sphere of illumination, with the corresponding error map as compared with Fig. 12j in Fig. 12l. Spectrally promoting the diffuse and specular components differently minimizes the color error particularly for the dielectric, green wax crayon prism to the right in the scene.

Fig. 12b shows the result of the polarization difference imaging for this lighting direction, where we observe only specular reflections. Fig. 12c shows the result of spectrally promoting the specular lighting basis image with Eq. 1, and 12d shows the result of promoting the specular lighting basis image with Eq. 3. Promoting the specular reflection with the diffuse model selects the wrong color for the dielectric specular highlight in Fig. 12c, as it appears red. When using the specular model, the highlight color appears to match that of the real amber LED lighting basis image. Fig. 11 enlarges a region of interest of Fig. 12c and d.

Conclusion

In this work, we have shown that performing image-based relighting in the multispectral domain produces more accurate re-

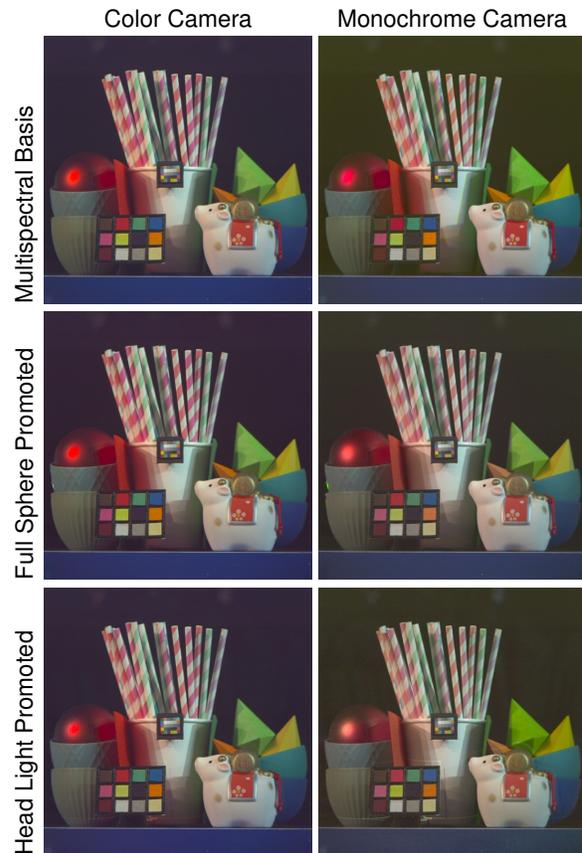


Figure 10. Row 1. IBRL result using the measured multispectral reflectance field, for color and monochrome cameras. Row 2. IBRL result using the full-sphere spectrally-promoted reflectance field, using Eq. 1. Row 3. IBRL result using the head light spectrally-promoted reflectance field, using Eq. 3. The shadows at the sides of the cup are relatively desaturated in the spectrally-promoted images for the monochrome camera.

lighting results than the traditional approach of multiplying lighting and reflectance in the tristimulus domain, at least for a still life scene with a variety of colorful objects. Recording a multispectral reflectance field requires more data, however, so we introduced the process of *spectral promotion* where monochrome or RGB images of an object under a particular direction of white light are promoted into their plausible appearance under a variety of narrow-band spectra of illumination based on the reflectance of the objects as observed by an alternate white-LED lighting condition. We showed that this produces good image-based relighting results in the multispectral domain with only a few more images than traditional reflectance field capture, although errors can arise where there are specular reflections or interreflected light. To address the problem of specular reflections, we showed that if the reflectance field is recorded under two polarization conditions which allow the separation of diffuse and specular reflections, then the reflectance components can be spectrally promoted according to independent color models and more accurate results can be obtained. As a result, we have shown that multispectral image-based relighting can be performed accurately and efficiently, enabling new applications for the simulation of realistic illumination on real-world subjects.

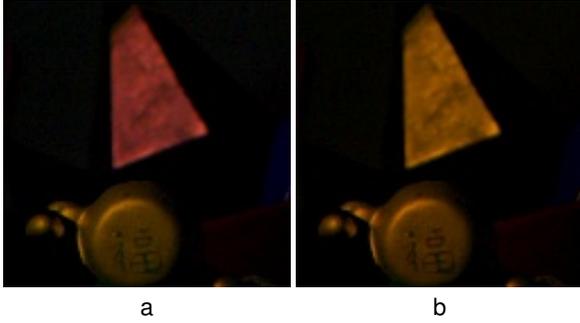


Figure 11. *a.* An inset of Fig. 12c, showing that promoting specular reflections produces color mismatches for dielectric materials. *b.* An inset of Fig. 12d, showing that promoting the specular reflections with Eq. 3 generates the expected appearance, although we did not measure the real specular-only reflectance components when the scene was lit by the amber LED.

Acknowledgments

The authors wish to thank Randy Hill, Kathleen Haase, Christina Trejo, Jay Busch, Andrew Jones, Shanhe Wang, Kathryn Rock, and the Natick Soldier Systems Center for their important support of this work. This project was sponsored by the U.S. Army Research Laboratory (ARL) under contract W911NF-14-D-0005 and in part by a USC Annenberg Ph.D. Fellowship. The content of the information does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred. USC has the potential to receive royalty payments from Google based on possible future sales.

References

- [1] B. Ajdin, M. Finckh, C. Fuchs, J. Hanika, and H. Lensch. Compressive higher-order sparse and low-rank acquisition with a hyperspectral light stage. Technical report, Univ. of Tuebingen, 2012.
- [2] B. A. Darling, J. A. Ferwerda, R. S. Berns, and T. Chen. Real-time multispectral rendering with complex illumination. In *Color and Imaging Conference*, volume 2011, pages 345–351. Society for Imaging Science and Technology, 2011.
- [3] P. Debevec. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '98*, pages 189–198, New York, NY, USA, 1998. ACM.
- [4] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar. Acquiring the reflectance field of a human face. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 145–156. ACM Press/Addison-Wesley Publishing Co., 2000.
- [5] A. Ghosh, G. Fyffe, B. Tunwattanapong, J. Busch, X. Yu, and P. Debevec. Multiview face capture using polarized spherical gradient illumination. *ACM Transactions on Graphics (TOG)*, 30(6):129, 2011.
- [6] J. Gu and C. Liu. Discriminative illumination: Per-pixel classification of raw materials based on optimal projections of spectral brdf. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 797–804, June 2012.
- [7] M. Kitahara, T. Okabe, C. Fuchs, and H. Lensch. Simultaneous estimation of spectral reflectance and normal from a small number of images. In *Proceedings of the 10th International Conference*

- on Computer Vision Theory and Applications, VISAPP 2015*, pages 303–313, 2015.
- [8] C. LeGendre, X. Yu, D. Liu, J. Busch, A. Jones, S. Pattanaik, and P. Debevec. Practical multispectral lighting reproduction. *ACM Transactions on Graphics (TOG)*, 35(4):32, 2016.
- [9] C. Ma, X. Cao, X. Tong, Q. Dai, and S. Lin. Acquisition of high spatial and spectral resolution video with a hybrid camera system. *International Journal of Computer Vision*, 110(2):141–155, 2014.
- [10] W.-C. Ma, T. Hawkins, P. Peers, C.-F. Chabert, M. Weiss, and P. Debevec. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Proceedings of the 18th Eurographics conference on Rendering Techniques*, pages 183–194. Eurographics Association, 2007.
- [11] S. K. Nayar, X.-S. Fang, and T. Boulton. Separation of reflection components using color and polarization. *International Journal of Computer Vision*, 21(3):163–186, 1997.
- [12] J. S. Nimeroff, E. Simoncelli, and J. Dorsey. Efficient re-rendering of naturally illuminated environments. In *Photorealistic Rendering Techniques*, pages 373–388. Springer, 1995.
- [13] J.-I. Park, M.-H. Lee, M. D. Grossberg, and S. K. Nayar. Multispectral imaging using multiplexed illumination. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–8. IEEE, 2007.
- [14] E. Reinhard, G. Ward, S. Pattanaik, and P. Debevec. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2005.
- [15] Y. Sato and K. Ikeuchi. Temporal-color space analysis of reflection. *JOSA A*, 11(11):2990–3002, 1994.
- [16] B. Tunwattanapong, G. Fyffe, P. Graham, J. Busch, X. Yu, A. Ghosh, and P. Debevec. Acquiring reflectance and shape from continuous spherical harmonic illumination. *ACM Transactions on graphics (TOG)*, 32(4):109, 2013.
- [17] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical engineering*, 19(1):191139–191139, 1980.

Author Biographies

Chloe LeGendre received a B.S. in Engineering in 2009 from the University of Pennsylvania and a M.S. in Computer Science from the Stevens Institute of Technology in 2015. From 2011 to 2015, she was an applications technologist in imaging and augmented reality for L’Oreal Research. She is currently pursuing a Ph.D. in Computer Science at the University of Southern California’s Institute for Creative Technologies, advised by Professor Paul Debevec.

Xueming Yu received a B.S. in Electrical Engineering from Shanghai Jiao Tong University in 2005 and an M.S. degree in Computer Science from the University of Southern California in 2010. He developed 3D display technologies at Sony Corporation R&D and the electronics and mechanics for several Light Stage and surface reflectance scanning systems at the USC Institute for Creative Technologies. In 2016 he joined Google VR as a Hardware Engineer.

Paul Debevec received a Ph.D. in Computer Science from the University of California, Berkeley in 1996 under Professor Jitendra Malik specializing in photogrammetry and image-based rendering, with post-doctoral work in High Dynamic Range Imaging, Image-Based Lighting, and Reflectance Field capture. Light Stage systems from his laboratory at the USC Institute for Creative Technologies have been used to digitize models of actors in numerous feature films and were recognized with an Academy Scientific and Engineering Award in 2010. In 2016 he joined Google VR as a Senior Staff Engineer. <http://www.debevec.org/>



Figure 12. *a.* The diffuse component of the reflection of the scene. *b.* The specular reflection of the scene as computed from polarization difference imaging using white LEDs. *c.* The result of spectrally promoting the specular reflection component using Eq. 1, hallucinating the appearance under the amber LED. *d.* The result of spectrally promoting the specular reflection component using Eq. 3. *e.* The result of spectrally promoting the diffuse component using Eq. 1. *f.* The blended result, interpolating between the images of *c* and *d* according to our approach. *g.* The sum of the spectrally-promoted diffuse reflection component and the blended result spectrally-promoted specular component. *h.* The absolute difference between *g* and the real lighting basis image *j*. *i.* The multispectral IBRL result using polarizing difference imaging, with the multispectral reflectance field computed in the same manner as *g*. *j.* The actual appearance of the scene when lit by the amber LED for this direction of the LED dome. *k.* The result of promoting both the diffuse and specular components of the reflection of the scene using 1, equivalent to the sum of *e* and *c*. *l.* The absolute difference between *k* and the real lighting basis image *j*.