Multispectral Lighting Reproduction from RGB Panoramas and Color Charts

Real Photo | LED Sphere Composite | Real Photo | LED Sphere Composite

Figure 1: *The left images show each of two subjects photographed in real outdoor lighting environments. The lighting was captured using panoramic HDR photograpy and color chart observations, allowing the light to be reproduced in an LED sphere with six different LED spectra. The right images show each subject photographed inside the multispectral LED sphere under reproductions of the captured lighting environments, composited into background photos, producing close matches to the originals.*

Abstract

We present a practical framework for reproducing omnidirectional incident illumination conditions with complex spectra using an LED sphere with multispectral LEDs. For lighting acquisition, we augment standard RGB panoramic photography with one or more observations of a color chart. We solve for how to drive the LEDs in each light source to match the observed RGB color of the environment and to best approximate the spectral lighting properties of the scene illuminant. Even when solving for non-negative intensities, we show that accurate illumination matches can be achieved with as few as four or six LED spectra for the entire ColorChecker chart for a wide gamut of incident illumination spectra. A significant benefit of our approach is that it does not require the use of specialized equipment (other than the LED sphere) such as monochromators, spectroradiometers, or explicit knowledge of the LED power spectra, camera spectral response curves, or color chart reflectance spectra. We describe two useful and easy to construct devices for multispectral illumination capture, one for slow measurements of detailed angular spectral detail, and one for fast measurements with coarse spectral detail. We validate the approach by realistically compositing real subjects into acquired lighting environments, showing accurate matches to how the subject would actually look within the environments, even for environments with mixed illumination sources, and demonstrate real-time lighting capture and playback using the technique.

1 Introduction

The way that a person appears – both photometrically and aesthetically – is greatly influenced by how they are lit. And when a subject in a studio is composited into a real or virtual scene, their lighting will either complement or detract from the illusion that they are actually present in the scene. Thus, being able to control and match studio lighting to real-world illumination environments is a useful capability for visual effects, studio photography, product design, and for designing garments and cosmetics.

Lighting reproduction systems as in [Debevec et al. 2002; Hamon et al. 2014] surround the subject with RGB color LEDs and drive them to match the lighting of the scene into which the subject will be composited. The light is recorded as panoramic, high dynamic range images, or rendered omnidirectionally from a global illumination lighting system. While the results can be believable – especially under the stewardship of color correction artists – it is not clear how accurate they are since only RGB colors are used for recording and reproducing the illumination: there is significantly more detail across the visible spectrum than what is being simulated. [Wenger et al. 2003] noted in particular that light reproduced with RGB LEDs can produce unexpected color casts even when each light source mimics the directly observable color of the original illumination. Ideally, a lighting reproduction system could faithfully reproduce the appearance of the subject under any combination of illuminants including incandescent, fluorescent, LED, and daylight, and any filtered or reflected version of such lighting.

Recently, several efforts [Gu and Liu 2012; Ajdin et al. 2012; Kitahara et al. 2015] have produced controllable LED spheres with more than just red, green, and blue LEDs in each light source for purposes such as multispectral material reflectance measurement. These systems add additional colors such as amber and cyan, as well as white LEDs which use phosphors to broaden their emission across the visible spectrum. In this work, we present a practical technique for driving the intensities of such arrangements of LEDs to accurately reproduce the effects of real-world illumination environments with any number of spectrally distinct illuminants in the scene. The practical nature of our approach rests in that we do not require explicit spectroradiometer measurements of the illumination; we require only traditional high dynamic range (HDR) panoramic photography and one or more observations of a color chart reflecting different directions of the illumination in the environment. Furthermore, we drive the LED intensities directly from the color chart and HDR panoramas, with no need to explicitly estimate illuminant spectra, or even to know the reflectance spectra of the color chart samples or the spectral sensitivity functions of the cameras involved. Our straightforward process is:

1. Photograph the color chart under each of the different LEDs
2. Record the illumination using standard panoramic photography plus one or more color chart directions
3. For each LED light source, estimate the appearance of a virtual color chart reflecting its direction of light from the environment
4. Drive the light source LEDs so that they best illuminate the virtual color chart with the estimated appearance

SIGGRAPH 2016 Technical Papers, July 24-28, 2016, Anaheim, CA
ISBN: 978-1-4503-ABCD-E/16/07
DOI: http://doi.acm.org/10.1145/9999997.9999999

Step one is simple, and step four simply uses a nonnegative least squares solver. For step two, we present two assemblies for capturing multispectral lighting environments which trade spectral angular resolution for speed of capture; one assembly acquires unique spectral signatures for each lighting direction; the other permits video rate capture. For step three, we present a straightforward approach to fusing RGB panoramic imagery and directional color chart observations. The result is a relatively simple and visually accurate process for driving multispectral LED sphere lights to reproduce the spectrally complex illumination effects of real-world lighting environments. We demonstrate our approach by recording several lighting environments with natural and synthetic illumination, and reproduce this illumination within an LED sphere with six distinct LED spectra. We show this enhanced lighting reproduction process produces accurate appearance matches for color charts and human subjects and can be extended to real-time multispectral lighting capture and playback.

2 Background and Related Work

In this work, we are interested in obtaining the incident illumination spectra arriving from all directions in a scene toward a subject. Traditional omnidirectional lighting capture techniques [Debevec 1998; Tominaga and Tanaka 2001] capture only tristimulus RGB imagery, often by photographing a mirrored sphere with a radiometrically calibrated camera. More recently, significant work has been done to estimate or capture spectral incident illumination conditions. With measured camera spectral sensitivity functions, [Tominaga and Tanaka 2006] promoted RGB spherical imagery to spectral estimates by projecting onto the first three principal components of a set of illuminant basis spectra, showing successful modeling of daylight and incandescent spectra within the same scene. However, the technique required measuring the camera sensitivity functions and was not demonstrated to solve for more than two illuminant spectra. [Kawakami et al. 2013] promoted RGB sky imagery to spectral estimates by inferring sky turbidity and fitting to a spectral skylight model. [Tominaga and Fukuda 2007] acquired mirror sphere photographs with a monochrome camera and multiple bandpass filters to achieve six spectral bands, and swapped the bandpass filters for a Liquid Crystal Tunable Filter to achieve 31 or 61 spectral bands. Since the filters' bands overlapped, they solved a linear system to best estimate the illuminant spectra. However, the system required specialized equipment to achieve such results.

Capturing natural daylight conditions has received attention of its own. [Stumpfel et al. 2004] used neutral density filters and an especially broad range of high dynamic range exposures with an upward-pointing fisheye lens to accurately record the upper hemisphere of incident illumination, but did not promote the RGB data to spectral measurements. [Debevec et al. 2012] recorded the full dynamic range of natural illumination in a single photograph based on the observed irradiance on diffuse grey strips embedded within a cut-apart mirrored sphere, allowing saturated light source intensities to be reconstructed. [Kider et al. 2014] augmented the fisheye capture technique of [Stumpfel et al. 2004] with a mechanically instrumented spectroradiometer that captured 81 sample spectra over the upper hemisphere. They then used bicubic interpolation over the hemisphere to validate a variety of spectral sky models, but did not explicitly fuse the spectral data with the high-resolution RGB fisheye photography of the sky. HDR image and lighting capture has also been extended to real-time video capture (e.g. [Kang et al. 2003; Unger and Gustavson 2007]) using interleaved exposures, though only for RGB imagery.

Significant work has been done to estimate illuminant spectra and/or spectral camera sensitivity functions (CSF's) using color charts with samples of known or measured reflectance. The MacBeth ColorChecker^{TM} Chart [McCamy et al. 1976] (now sold by X-Rite^{TM}) simulates a variety of natural reflectance spectra, including light and dark human skin, foliage, the sky, certain types of flowers, and a ramp of spectrally flat neutral tones. Its 24 patches have 19 independent reflectance spectra, which, when photographed by an RGB camera yield 57 linearly independent integrals of spectrum of light falling on the chart. [Rump et al. 2011] showed that CSFs could be estimated from an image of such a chart (actually, a more advanced version with additional spectra) under a known illuminant, and [Shi et al. 2014] showed that the illuminant spectrum could be estimated successfully from a photograph of such a chart. Thus, relatively rich spectral information is available from color chart observations, which we leverage in our work.

Reproducing the illumination of a real-world scene inside a studio can be performed as in [Debevec et al. 2002] by surrounding the actor with computer-controlled red-green-blue LED lights and driving them to replicate the color and intensity of the scene's incident illumination from each direction, which has been applied in commercial films such as *Gravity* [Hamon et al. 2014]. [Wenger et al. 2003] showed, however, that using only red, green, and blue LEDs significantly limits color rendition quality of the system, and direct comparisons to the appearance of the subject in the actual illumination environment have been missing. To address this, [Wenger et al. 2003] showed that by measuring the spectral light output of the LEDs, the spectral reflectance of a set of material samples, and the camera's spectral sensitivity functions, the color rendition properties of particular illuminants such as incandescent and fluorescent lights could be simulated reasonably well with RGB LEDs and quite closely with a nine-channel LED light source. However, this work required specialized measurement equipment and limited discussion to individual light sources, not fully spherical lighting environments. Since then, a variety of LED sphere systems have been built with greater spectral illuminant detail, including the 6-spectrum system of [Gu and Liu 2012], the 16-spectrum system of [Ajdin et al. 2012], and the 9-spectrum system of [Kitahara et al. 2015]. In our work, we use a six-channel multispectral LED sphere, but instead of using multiple lighting conditions to perform reflectance measurement, we wish to drive the multispectral LEDs in a single lighting condition which optimally reproduces the appearance a subject would have in a spectrally complex illumination environment as seen by a particular imaging system. And furthermore, we wish to perform this with an easy-to-practice process system for recording the spectral properties of the illumination and driving the LED light sources accordingly, without the need for specialized equipment beyond the LED sphere itself.

3 Method

In this section we will describe our techniques for capturing the incident illuminant and driving the multispectral light sources in order to reproduce the effect of the illumination on a subject as seen by a camera.

3.1 Driving Multispectral Lights with a Color Chart

We first consider the sub-problem of reproducing the appearance of a color chart to a given camera in a particular lighting environment using a multispectral light source. Photographing the chart with an RGB camera produces pixel values P_{ij} where i is the index of the given color chart patch and j is the camera's jth color channel. We now wish to drive the differently colored LEDs at a set of intensities which best reproduce the chart's appearance. In the LED sphere, we photograph the chart lit by each of the different spectra of LEDs k at unit intensity to produce images as seen in Fig. 3). We then construct a matrix L where L_{ijk} is formed by sampling the average pixel value near the center of color chart square i for camera color channel j under LED lighting condition k. To achieve even lighting, we do this using all of the LED sphere light sources of a given spectrum simultaneously, though a single multispectral light could be used. We consider L to be the $ij \times k$ matrix whose columns correspond to the LED spectra k and whos rows unroll the indices i and j to place the RGB pixels for all chart squares in the same column. We then simply need to solve for LED intensity coefficients α_k which minimize:

$$\sum_{i=1}^{m}\sum_{j=1}^{3}(P_{ij} - \sum_{k=1}^{n} L_{ijk}\alpha_k)^2 = ||\mathbf{P} - \mathbf{L}\boldsymbol{\alpha}||^2 \qquad (1)$$

Where m is the number of color chart patches and n is the number of differently colored LEDs. Conveniently, this process does not require knowing the illuminant spectra, the camera spectral sensitivity functions, or even the chart reflectance spectra.

We could easily solve for the LED intensities α_k which minimize Eq. 1 using linear least squares (LLS), but this may lead to negative weights for some of the LED colors, which is not physically realizable. We could, however, simulate such illumination by taking two photographs, one where the positively-weighted LEDs are turned on, and a second where the absolute values of the negatively-weighted LEDs are turned on, and subtract the pixel values of the second image from the first. If our camera can take these two images very quickly, this might be an acceptable approach. However, for greater flexibility and to facilitate motion picture recording, we can also solve for the LED intensity weights using non-negative least squares (NNLS), yielding the optimal solution where the lighting can be reproduced at once and captured in a single photograph.

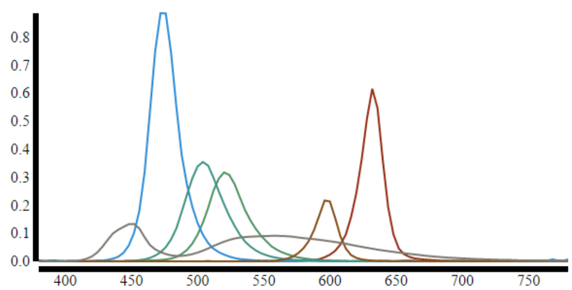

Figure 2: *The illumination spectra of the six LEDs in each light source of the LED sphere.*

To test the technique, we used a Canon 1DX DLSR camera to photograph a Matte ColorChecker Nano chart from Edmund Optics, which includes 12 neutral grayscale squares and 18 color squares with spectra as in [McCamy et al. 1976] under incandescent, fluorescent, and daylight illuminants. We then recorded the color chart's $\mathbf{L}$ matrix under our six LED spectra (Fig. 2) producing the images in Fig. 3. The background squares of the first row of Fig. 4 show the appearance of the color chart[1] under the three illuminants, while the circles inside the squares show the appearance under the nonnegative least squares solve for six LEDs (red, green, blue, cyan, amber, and white, or RGBCAW) for direct comparison. This yields charts which are very similar in appearance, to the point that many circles are difficult to see at all. The second two rows show the results of reproducing the illumination with just four (RGBW) or three (RGB) LED spectra. The RGBW matches are also quite good, but the RGB matches are generally poor, producing oversaturated colors which are easily distinguishable from the original appearance. These differences are reported quantitatively in Fig. 5.

[1]For display, RAW pixel values are converted to sRGB color space.

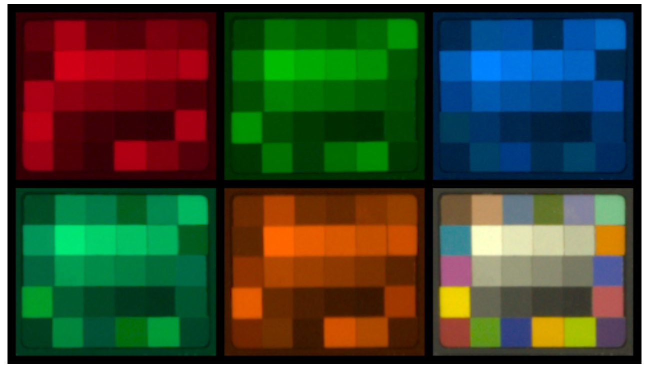

Figure 3: *The color chart illuminated by the red, green, blue, cyan, amber, and white LEDs, allowing measurement of the* $\mathbf{L}$ *matrix.*

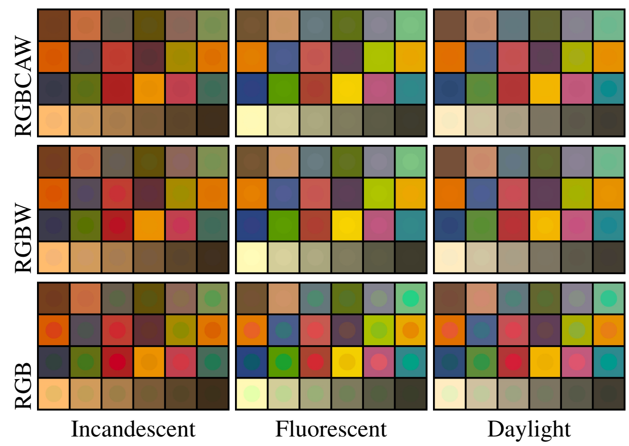

Figure 4: *Comparison color charts for three illuminants where the background of each square is the original chart appearance, and the circles (sometimes invisible) show the color chart under the reproduced illumination in the LED sphere. Six (RGBCAW), four (RGBW), and three (RGB) spectral channels are used.*

Using Different Cameras Although not preferred, cameras with differing spectral sensitivity functions may be used to record the environment lighting and the reproduced illumination; the solving process will still endeavor to find LED intensities which reproduce the chart's appearance. Fig. 6 shows a color chart photographed under several illuminants with a Canon 1DX DSLR, and then as it appears photographed with a Nikon D90 DSLR under reproduced illumination calculated from an $\mathbf{L}$ matrix from the Nikon D90. Despite different sensitivities, the matches are reasonably close, with the most notable discrepancies in the blue and purple squares under fluorescent illumination.

3.2 Recording a Spherical Lighting Environment with a Color Chart

Real lighting environments often feature a variety of illuminant spectra with indirect light from numerous materials modulating the spectra of the illuminants. A color chart, lit by the entire hemisphere of light it points towards, lacks the angular discrimation necessary to record detailed directional effects of such illumination. To

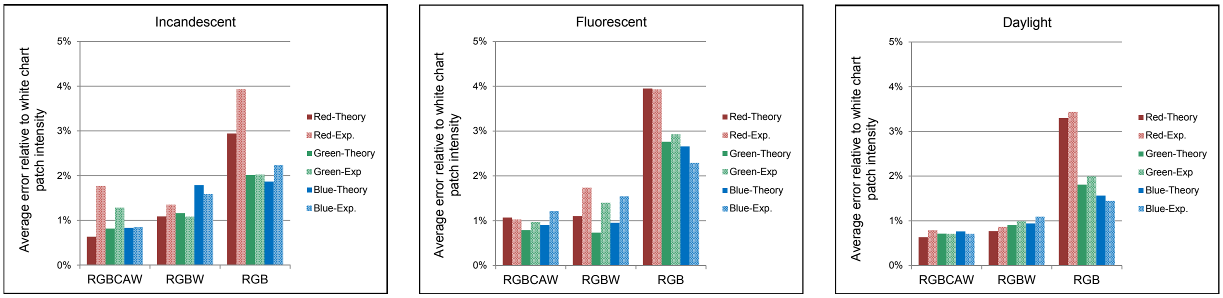

Incandescent Fluorescent Daylight

Figure 5: *Quantitative error plots for Fig. 4 for different illuminants and number of LED spectra used for lighting reproduction, based on the average squared errors of raw pixel values from all 30 ColorChecker Nano patches. Both the theoretical error from the minimization and actual error from reproducing the illumination are reported.*

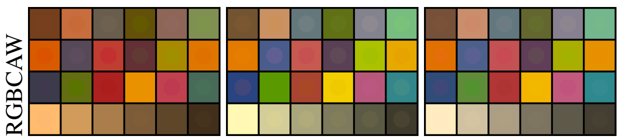

Incandescent Fluorescent Daylight

Figure 6: *Charts photographed with a Canon 1DX camera (background squares) with lighting reproduced in the LED sphere photographed by a Nikon D90 camera (foreground circles).*

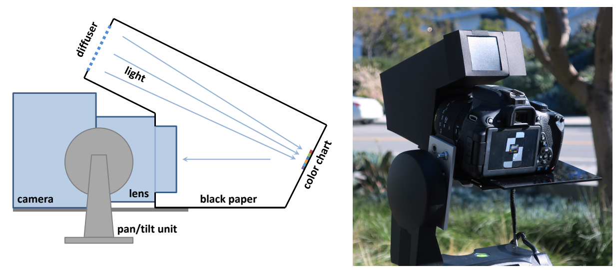

Figure 7: *We use a black box to photograph a small color chart illuminated by only a small section of the environment. The box is placed over the lens of a camera on a pan/tilt rig to capture an omnidirectional multispectral lighting environment.*

address this, we built a black paper box (Fig. 7) to illuminate a color chart by just a narrow cone of light from the environment and secure it in front of a DSLR camera. We chose the width of the cone to roughly match the angular resolution of our LED sphere, which has light sources spaced 12° apart around its equator. We cover the aperture with three sheets of 5-degree light-shaping diffuser from Luminit, Inc. to antialias the incident light. A printable pattern for this color chart box is provided as supplemental material.

To capture a spherical lighting environment, we place a Canon Rebel 3Ti DSLR with a Sigma 50mm macro lens with the color chart box on a GigaPan EPIC 100 pan/tilt rig, and program the rig to record the 360° spherical environment as a set of 30 horizontal by 24 vertical directions, which also roughly matches the angular resolution of the LED sphere. Because the chart becomes brightly illuminated when pointed toward a light source, we set the camera to aperture-priority (Av) mode so that the shutter speed will be automatically chosen to expose the color chart properly at each position. The shutter speeds are recorded within each image's EXIF metadata. To avoid stray light affecting the light meter, we cover the viewfinder during capture. Capturing the 720 images requires approximately one hour.

We stabilize the color chart image sequence using fiducial markers at the sides of the chart, allowing us to reliably average a 20×20 pixel area in the center of each square to sample the color chart's pixel values. We divide the pixel values by the EXIF exposure time for each photograph to produce values proportional to the actual HDR radiance from the chart squares. By choosing an appropriate f/stop and ISO, the exposure times vary from around 1/2 second to 1/1000 second in a typical environment. A full lat-long panorama of charts can be seen in Fig. 8(top). Transposing the data yields a chart of lat-long panoramas, where each section is low-resolution image of the environmental illumination spectrally modulated by the reflectance of each chart square as in Fig. 8(bottom).

3.3 Lighting a Subject with a Color Chart Panorama

We know from Sec. 3.1 how to drive a multispectral LED light source to match the appearance of a color chart, and we have now measured how each section of our lighting environment, at about the resolution of our LED sphere, illuminates a color chart. Thus, for each LED sphere light source, we bilinearly interpolate the four nearest charts to the precise direction of the light to create an interpolated chart P for that light. We then drive the LEDs in the light so that it illuminates the chart as it really appeared lit by the environment. The LED sphere creates the varying LED intensities through pulse-width modulation (PWM), using 12 bits of lighting input to achieve 4K possible intensity values. We automatically scale the overall brightness of the lighting environment so that all of the LED intensity values fit within range, and the lighting environment is reproduced up to this scaling factor.

Fig. 9 shows a subject illuminated by two lighting environments which were spectrally recorded using the panoramic color chart technique of Sec. 3.2, next to results of driving the LED sphere with 6-channel RGBCAW LED intensities solved for from the chart data. It achieves good matches for the skin tones, clothing, and specular highlights.

3.4 Augmenting RGB Panoramas with Color Charts

Recording a lighting environment with the panoramic color chart method takes significantly longer than typical HDRI lighting capture techniques such as photographing a mirrored sphere or acquiring HDR fisheye photographs. To address this, we decribe a simple process to promote a high-resolution RGB HDRI map to multispec-

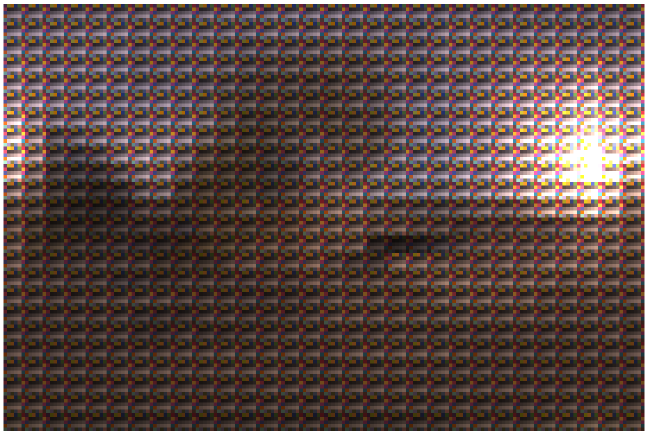

Panorama of sampled directionally lit color charts

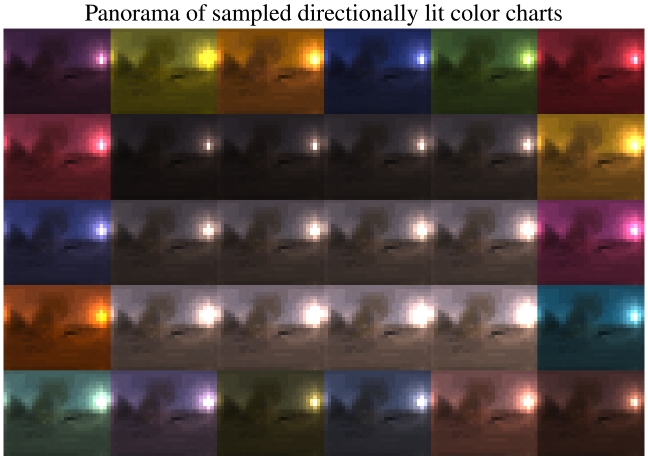

Chart of panoramas reflected by each color sample

Figure 8: *Visualizations of an outdoor multispectral lighting environment panorama made from sampled color charts as in Sec. 3.2.*

(a) Real light (b) LED sphere (c) Real light (d) LED sphere

Figure 9: *(a) A subject in a daylight environment captured as the panorama of color charts in Fig. 8. (b) The subject in the LED sphere with the lighting reproduced as in Sec. 3.3 (c) A subject in an indoor lighting environment with LED, fluorescent, and incandescent light sources also captured as a panorama of color charts. (d) The subject lit in the LED sphere with the indoor lighting.*

tral estimates using a sparse set of color chart observations.

Suppose that one of our LED sphere light sources is responsible for reproducing the illumination from a particular set of pixels T in the RGB HDR panorama, and the average pixel value of this area is Q_j (where j indicates the color channel.) Q's three color channels are not enough information to drive the multispectral light source, and we wish we had seen an entire color chart lit by this part of the environment T. Suppose that we can estimate, from the sparse sampling of color charts available, the appearance of a color chart P_{ij} lit by the same general area of the environment as T, and we consider it to be reflecting an overall illuminant $I(\lambda)$ corresponding to our environment area T. We presume then that:

$$P_{ij} = \int_\lambda I(\lambda) R_i(\lambda) C_j(\lambda)$$

where $R_i(\lambda)$ is the spectral reflectance function of patch i of the color chart, and $C_j(\lambda)$ is the spectral sensitivity function of the j'th color channel of our camera. Our color chart features a white, spectrally flat neutral square which we presume to be the zeroth index patch $R_0(\lambda) = 1$. (In practice, we scale up P_0 to account for the fact that white patch is only 90% reflective.) This patch reflects the illumination falling on the chart as seen by the camera, which should in theory should equal our RGB pixel observation Q. In general, since the illuminant estimate $I(\lambda)$ corresponds to a larger area of the environment than T, Q will not be equal to P_0. For example, if T covers an area of foliage (for example) modulating $I(\lambda)$ by spectral reflectance $S(\lambda)$, and the illuminant I broadly accounts for the incident daylight, we would have:

$$Q = \int_\lambda I(\lambda) S(\lambda) C_j(\lambda)$$

Since we want to estimate the appearance of a color chart P'_{ij} illuminated by the environment area T, we are interested in knowing how the modulated illuminant $I(\lambda)S(\lambda)$ would illuminate the color chart squares $R_i(\lambda)$ as:

$$P'_{ij} = \int_\lambda I(\lambda) S(\lambda) R_i(\lambda) C_j(\lambda)$$

We do not know the spectral reflectance $S(\lambda)$, but we know that environmental reflectance functions are generally smooth, whereas illuminants can be spiky. If we assume that $S(\lambda) \approx \bar{s_j}$ over each camera sensitivity function $C_j(\lambda)$, we have:

$$P'_{ij} = \bar{s_j} \int_\lambda I(\lambda) R_i(\lambda) C_j(\lambda)$$

We can now write $P'_{ij} = \bar{s_j} P_{ij}$ and since $R_0(\lambda) = 1$, we can write:

$$P'_{0j} = \bar{s_j} \int_\lambda I(\lambda) C_j(\lambda) = Q_j$$

so $\bar{s_j} = Q_j / P_{0j}$ and we compute $P'_{ij} = Q_j P_{ij} / P_{0j}$. In effect, we divide the estimated illuminant color chart P by its white square and recolor the whole chart lit by the observed RGB pixel average Q to arrive at the estimate P' for a color chart illuminated by T. This chart is consistent with Q and retains the same relative intensities within each color channel of the estimated illuminant falling on the chart patches.

If our camera spectral response functions were known, then it might be possible to estimate $S(\lambda)$ as more than a constant per color channel to yield a more plausible P'. This is of interest to investigate in future work.

3.5 Fast Multispectral Lighting Environment Capture

Recording even a sparse set of color chart observations using the pan/tilt technique is still relatively slow compared to shooting an RGB HDR panorama. If the scene is expected to comprise principally one illuminant, such as daylight, one could promote the entire HDRI map to multispectral color information using a photograph of a single color chart to comprise every color chart estimate P'. However, for scenes with mixed illumination sources, it would be desirable to record at least some of the angular variation of the illumination spectra.

To this end, we constructed the fast multispectral lighting capture system of Fig. 10(a), which points a DSLR camera at chrome and black 8cm spheres from Dube Juggling Equipment and five Matte

Nano ColorChecker charts from Edmund Optics aimed in different directions. The chrome sphere was measured to be 57.5% reflective using the reflectance measurement technique of [Reinhard et al. 2005] Ch. 9. The black acrylic sphere, included to increase the observable dynamic range of illumination, reflects 4% of the light at normal incidence, increasing significantly toward grazing angles in accordance with Fresnel's equations for a dilectric material. The five color charts face forward and to point $\pm 45^{\circ}$ vertically and horizontally. A rigid aluminum beam secures the DSLR camera and 100mm macro lens 135cm away from the sphere and chart arrangement. Two checker fiducial markers can be used to stabilize dynamic footage if needed.

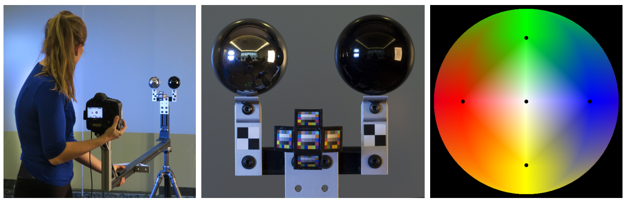

Figure 10: *(a) Photographer using the fast multispectral lighting capture system (b) Image from the lighting capture camera featuring spheres and color charts. (b) Color-coded visualization of interpolating the five charts over the frontal hemisphere of surface normals.*

If the lighting is static, we can record an HDR exposure series and reconstruct the RGB lighting directly from the chrome sphere as in [Debevec 1998]. If the lighting is dynamic and must be recorded in a single shot, we can set the exposure so that the indirect light from the walls, ceiling, or sky exposes acceptably well in the chrome ball, and that the light sources can be seen distinctly in the black ball, and combine these two reflections into an HDR map. We begin by converting each sphere reflection to a latitude-longitude (lat-long) mapping as in [Reinhard et al. 2005] Ch. 9. Since the spheres are 3° apart from each other with respect to the camera, we shift the black sphere lat-long image to rotate it into alignment with the mirrored sphere image. The spheres occlude each other from a modest set of side directions, so we take care to orient the device so that no major sources of illumination fall within these areas.

We correct both the chrome and black sphere maps to 100% reflectivity before combining their images in HDR. For the black ball, this involves dividing by the angularly-varying reflectivity resulting from Fresnel gain. We measured the angularly-varying reflectivity of the two spheres in a dark room, moving a diffuse light box to a range of angles incident upon the spheres, allowing us to fit a Fresnel curve with index of refraction 1.51 to produce the corrective map as in [Stumpfel 2004]. Since the dielectric reflection of the black ball depends significantly on the light's polarization, we avoid using the black sphere image when reconstructing skylight.

In either the static or dynamic case, if there are still bright light sources (such as the sun) which saturate in all of the available exposures, we reconstruct their RGB intensities indirectly from the neutral grey squares of the five color charts using the single-shot light probe technique of [Debevec et al. 2012]. Other dynamic RGB HDR lighting capture techniques could be employed, such as [Unger and Gustavson 2007].

Once the RGB HDR lighting environment is assembled, we need to promote it to a multispectral record of the environment using the five color charts $P_1 \ldots P_5$. For each LED Sphere light, we estimate how a virtual color chart P' would be appear lit by the light it is responsible for reproducing. Since we do not have a chart which points backwards, we postulate a backwards-facing color chart P_6 as the average of the five observed color charts. The six color charts now point toward the vertices of an irregular octahedron. To estimate P', we determine which octahedral face the LED Sphere light aims toward, and compute P' as the barycentric interpolartion of the charts at the face's vertices. Fig. 10(c) visualizes this interpolation map over the forward-facing directions of a diffuse sphere.

We are nearly ready to drive the LED sphere with the captured multispectral illumination. For each light, we determine the average RGB pixel color Q of the HDR lighting environment area corresponding to the light source. We then scale the color channels of our color chart estimate P to form P' consistent with Q as in Sec. 3.4. We then solve for the LED intensities which best reproduce the appearance of color chart P' using the technique of Sec. 3.1.

Figure 11: *(upper left) Unwrapped chrome and black acrylic sphere images from Fig. 10 in latitude-longitude mapping. (upper right) Sampled colors from the five color charts. (bottom) Derived high-resolution maps for driving the red, green, blue, cyan, amber, and white LEDs to match the multispectral illumination.*

3.6 Reproducing Dynamic Multispectral Lighting

The fast multispectral capture system allows us to record dynamic lighting environments. Most digital cinema cameras, such as the RED Epic and ARRI Alexa, record a compressed version of RAW sensor data which can be mapped back into radiometrically linear measurements according to the original camera spectral sensitivity curves. While this feature is not usually present in consumer cameras, the Magic Lantern software add-on (http://www.magiclantern.fm/) allows many Canon DSLR cameras to record RAW video at high definition resolution. We installed the add-on onto a Canon 5D Mark III video camera to record 24fps video of the multispectral lighting capture apparatus, which we could play back at 24fps on the LED sphere as a spherical frame sequence. Fig. 15 shows some results made using this process.

4 Results

In this section, we present comparisons of various lighting reproduction results to photographs of the subject in the original lighting environment. Some effort was made to match the pose and expression from real to reproduced illumination, even though in some cases the shoots were done hours or a day apart. In the LED sphere, a black foamcore board was placed 1m behind the actor to avoid the distracting appearance of the light sources in the photographs. All still photographs were taken in RAW mode on a Canon 1DX DSLR camera, and no image-specific color correction was performed. To create printable images, the raw RGB pixel values were transformed roughly to sRGB color space using the same 3×3 color matrix:

$$C = \begin{bmatrix} 1.88 & -0.879 & 0.0061 \\ -0.228 & 1.58 & -0.330 \\ 0.0393 & -0.696 & 1.63 \end{bmatrix}$$

which boosts color saturation to account for the significant spectral overlap of the sensor filters. Finally, a single brightness scaling was applied equally to all three color channels of each image to account for the brightness variation between the real and reproduced illumination in the LED sphere and the variations in camera shutter speed, f/stop, and ISO. The original Canon 1DX RAW .CR2 images for Fig. 1 and Fig. 12 are provided as supplemental material.

Fig. 12 shows two subjects in indoor, sunlit, and cloudy lighting environments. The indoor environment featured an incandescent soft box light to the subject's left, and spectrally distinct blue-gelled white LED light panels to their right, with fluorescent office lighting from the ceiling. The lighting was recorded using the five-color-chart and reflective sphere capture technique of Sec. 3.5. Then, as quickly as possible, the subject was photographed in the same environment. Later, in the LED sphere, the lighting was reproduced using nonnegative 6-channel RGBCAW lighting, 4-channel RGBW lighting, and 3-channel RGB lighting solves as described in Sec. 3.4. Generally, the matches are visually very close for RGBCAW and RGBW lighting reproduction, whereas colors appear too saturated using RGB lighting. The fact that the nonnegative RGBW lighting reproduction to performs nearly as well as RGBCAW suggests that these four spectra may be sufficient for many lighting reproduction applications. The bottom rows include a sunset lighting condition reproduced with RGBCAW lighting where the light of the setting sun was changing rapidly. We recorded the illumination both before and after taking the pictures of the subject, and averaged the two sets of color charts and RGB panoramas to solve for the lighting reproduction condition. In the bottom row, the subject appears somewhat shinier in the reproduced illumination; this is possibly because the real lighting environment photo was taken closer in time to the application of her makeup.

Fig. 13 compares the results of a linear least squares RGBCAW solve with positive and negative LED intensities to a nonnegative least squares solve for the indoor environment of Fig. 12. The linear least squares solve reproduces the incandescent light with mostly positive illumination with a small amount of negative blue, and the blue-gelled light as positive white and blue LEDs with negative amounts of red, amber, green, and cyan. Both solutions closely approximate the appearance of the subject in the original light, but slight subject motion between the two photos (taken less than a second apart) causes artifacts in the skin texture and the position of the eyelids. The slight theoretical improvement of a least sqaures solve does not seem to outweigh the convenience of single shot photography from a nonnegative solve.

Fig. 14 compares the results of using interpolated illuminant information from all five color charts from the fast lighting capture device to using only spectral information from only the front-facing color chart, using the indoor environment with LED, fluorescent, and incandescent illuminants. The 5-chart example produces slightly richer colors, and can be seen to match the original effect of the lighting in the fourth row of Fig. 12, but the effect is subtle. The color charts in Fig. 14 visualize the error between the actual appearance of the left-facing and right-facing color charts in the lighting capture device compared to their appearance under the reproduced illumination, showing slightly better matches for the 5-chart result. This indicates that having even a single front-facing color chart observed in the scene is helpful for achieving acceptable multispectral lighting reproduction.

Fig. 1 shows each subject in real and reproduced outdoor illumination using RGBCAW lighting, and image compositing to place the subject in the LED sphere on top of the original background. The matte was obtained by taking a second silhouetted photograph a fraction of a second after the first, with the LED sphere lights off, and the black foamcore lit brightly with a pair of flash units. The result are images where it is difficult to tell which image is real and which is the composited reproduced illumination.

Fig. 15 show our technique applied to an image sequence with dynamic lighting, where an actor is rolled through a set with fluorescent, incandescent, and LED light sources with various colored gels applied. The full dynamic range of the illumination was reconstructed using both the chrome and black acrylic spheres, and promoted to multispectral observations using the five color charts as in Sec. 3.5. The derived six-channel RGBCAW lighting was played back in the LED sphere as the actor repeated the performance in front of a separately-lit green screen. Finally, the actor was composited into a clean plate of the set. The real and reproduced lighting scenarios are similar, although differences arise from discrepancies in the pose of the actor and some spill light from the green screen, especially under the hat, and some missing rim lighting which was blocked by the green screen. The matting is intentionally done with an automated process which does not alter the color of the foreground element pixels; better matte lines could be achieved by a professional compositor. An infrared matte as in [Debevec et al. 2002] which allows the LED lights to shine through the backing could improve the matte lines and rim lighting. The limited lighting resolution can be seen in the shadows of the hat on the face, though not very much in the eye reflections. Overall, the lighting is effectively reproduced from the original set to the LED sphere.

5 Discussion

A strength of our technique is that it does not require measurement, knowledge, or estimation of a single spectral illumination or reflectance function; we just need to know how our camera sees a color chart lit by the desired section of the lighting environment and how it sees the color chart lit by the differently colored LEDs. Another reasonable approach to the problem would involve measuring the the LED spectra as well as the spectrum of every lighting direction in the environment, and then approximating the incident illumation spectra as best as possible using linear combinations of the LEDs. But such a technique would require spectral measurement equipment and would likely be cumbersome for capturing the lighting environments. Moreover, the errors made in aproximating the illuminant spectra with the LEDs will not necessarily be minimal with respect to how the camera will see materials lit by these spectra. Since we optimize directly so that the LEDs illuminate the color charts the way they appeared in the original lighting, in theory there is no better solution to be had from the observations available.

6 Future Work

Our work currents limits its consideration to the visible spectrum, largely because most commonly available cameras are designed to have limited sensitivity to infrared and ultraviolet light. Nonetheless, being able to record and reproduce hyperspectral illumination could have applications in medicine, garment manufacturing, and night vision research. Adding hyperspectral sensors and color chart patches would allow the technique to be extended into this domain.

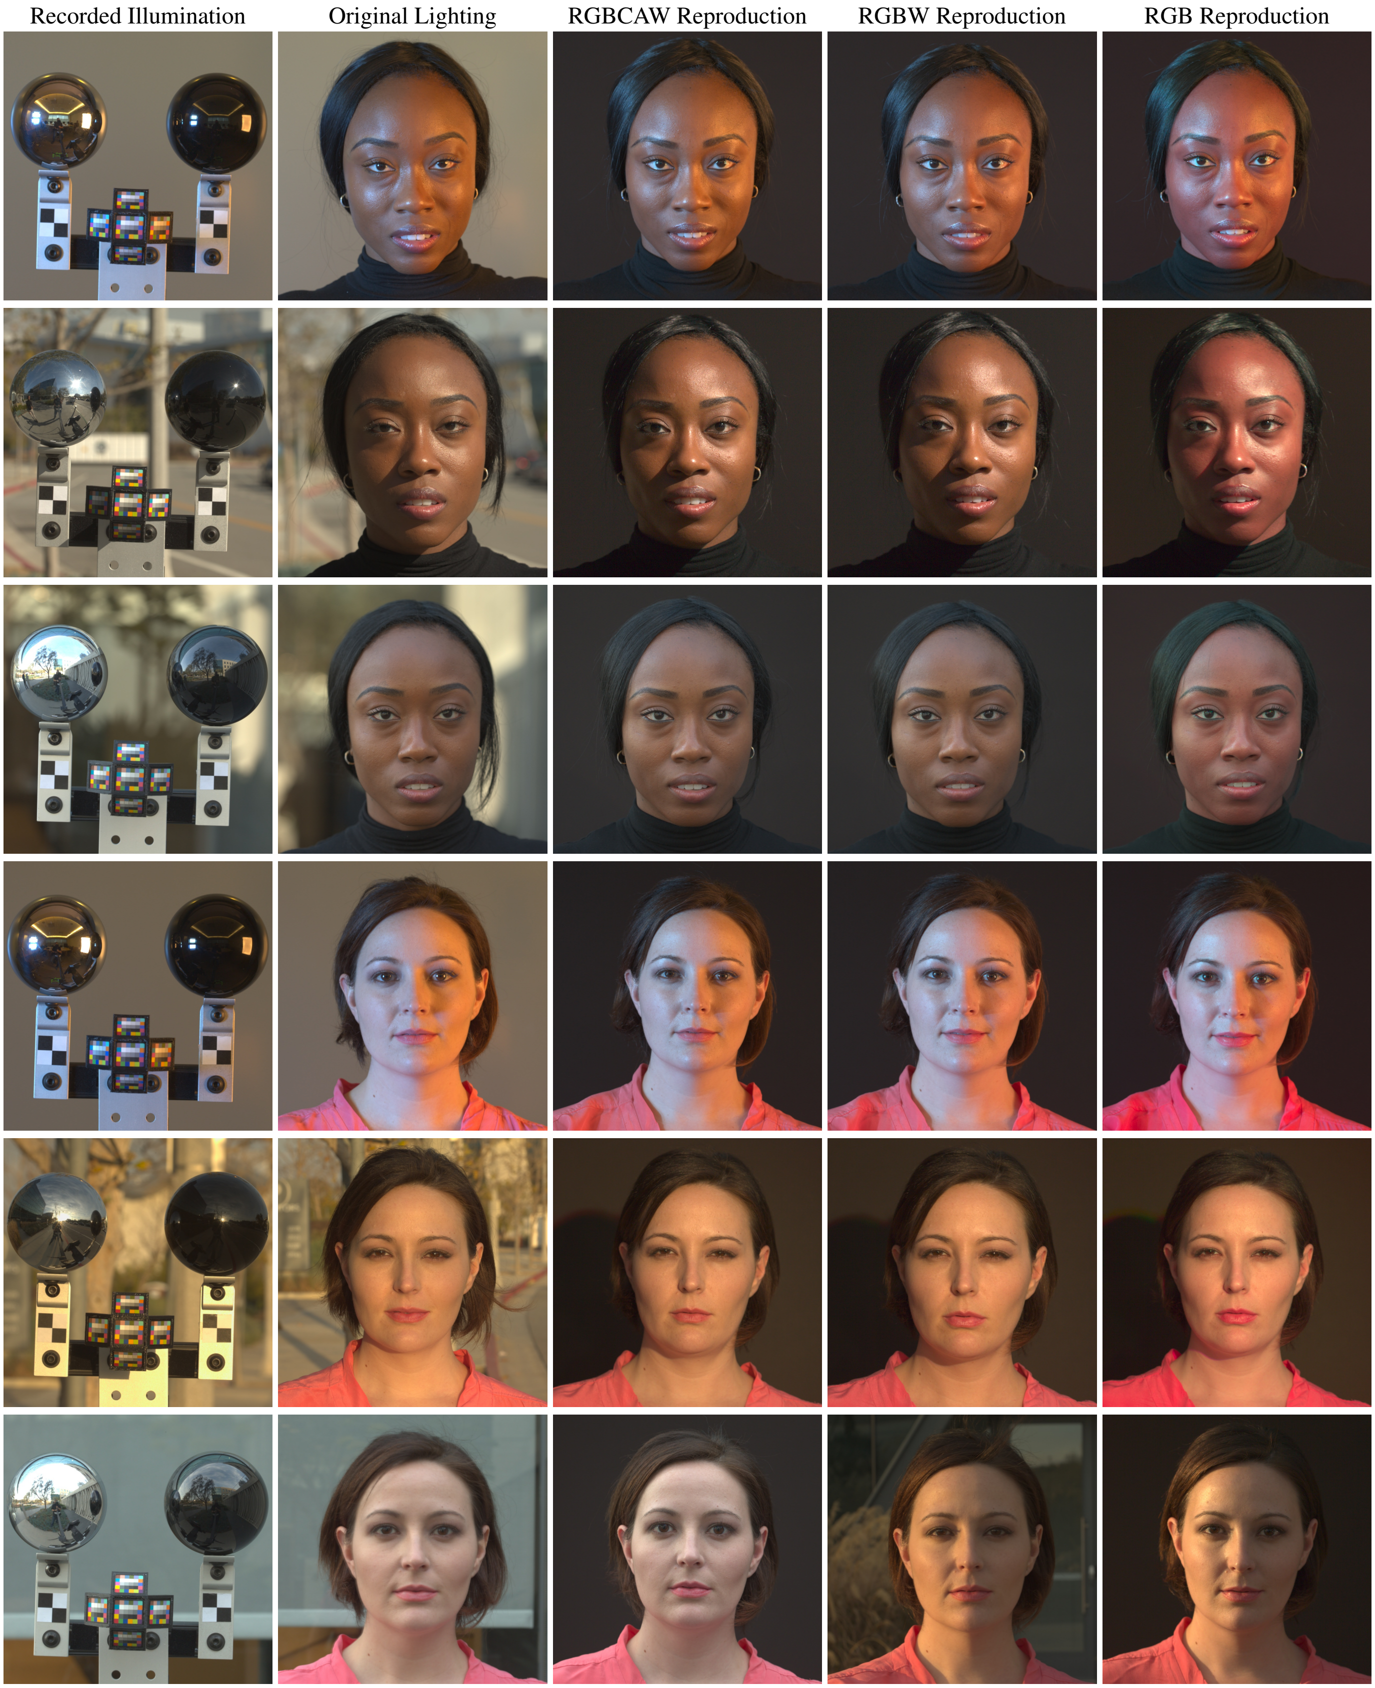

Figure 12: *Two subjects in a variety of lighting environments (left columns) with three different lighting reproduction techniques. Generally, RGBCAW and RGBW reproduction produces accurate results, whereas using just RGB LEDs tends to oversaturate colors. The final two images of the bottom row show the subject and camera rotated 90 degrees to the left in the same sunset lighting of the fifth row, with RGBCAW lighting reproduction.*

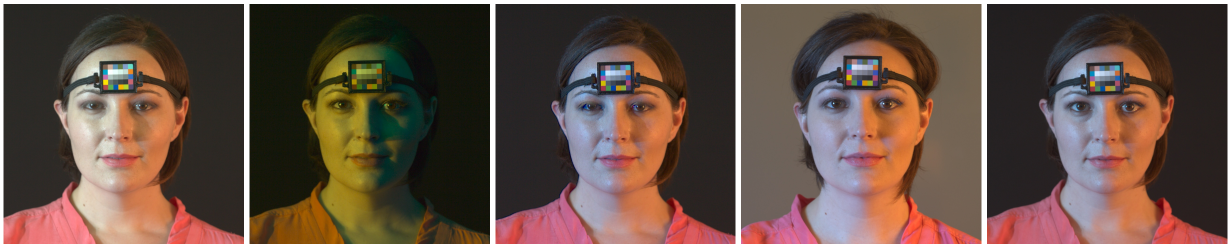

Positive | Negative | Difference | Real light | Nonnegative solve

Figure 13: *Comparing a linear least squares solve with positive and negative LED intensities (which must be photographed in separate images) to a nonnegative least squares solve which can be photographed in a single image, using 6-channel RGBCAW illumination.*

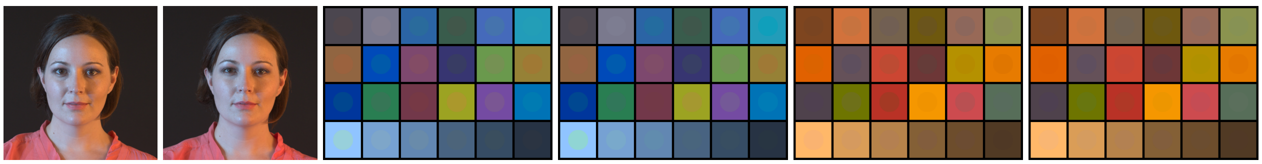

1 chart | 5 charts | Left error, 1 chart | Left error, 5 charts | Right error, 1 chart | Right error, 5 charts

Figure 14: *Comparison of single-chart reconstruction versus five-chart reconstruction.*

We use commonly available ColorChecker chars to reveal multispectral information about the incident illumination, and although we avoid the need to explicitly estimate illumination spectra, this could be done using the technique of [Shi et al. 2014]. Since the ColorChecker has only nineteen distinct spectra, it may fail to record information in certain areas of the spectrum where an item, such as a particular garment, has interesting spectral detail. In this case, a sample of the item could be added to the color chart, and the **L** matrices and lighting capture could proceed with this additional spectral detail taken into consideration. Furthermore, if a particular item's colors are not being reproduced faithfully, its weight in the nonnegative least squares solve could be increased, improving its color rendition at the expense of the other patches.

While our technique successfully tailors its spectral illumination and reflectance matching to the target image sensor, it does not explicitly attempt to replicate the effects of the illumination as seen by the human eye. Empirically, for RGBCAW and RGBW lighting, the subject does appear similarly lit to their original lighting, though it is much dimmer when direct sunlight is being reproduced. Interestingly, 3-channel RGB reproduction looks far worse to the eye than to the camera for which the solution has been optimized. Tailoring the lighting reproduction to the human eye could theoretically be done by measuring the L matrices for the color chart and LEDs in the photometric XYZ color space, which could be realized by pointing a spectroradiometer at each color chart square lit by each LED color to build up an L matrix from the reported XYZ values. It might also be of interest to optimize for multiple observers, such as for both the photographer and their camera.

7 Conclusion

In this paper, we have presented a practical way to reproduce complex, multispectral lighting environments inside an LED sphere with multispecrtal light sources. The process is easy to practice, since it simply adds a small number of color chart observations to traditional HDR lighting capture techniques, and the only calibration required is to observe a color chart under each of the available LED colors in the sphere. The technique produces visually close matches to how the subject actually would look in the real lighting environments, even with as few as four LED spectra available (RGB and white), and can be applied to dynamic scenes. The technique may have useful applications in visual effects production, virtual reality, studio photography, cosmetics testing, and clothing design.

References

Ajdin, B., Finckh, M., Fuchs, C., Hanika, J., and Lensch, H. 2012. Compressive higher-order sparse and low-rank acquisition with a hyperspectral light stage. Tech. rep., University of Tuebingen.

Debevec, P., Wenger, A., Tchou, C., Gardner, A., Waese, J., and Hawkins, T. 2002. A lighting reproduction approach to live-action compositing. In *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques*, ACM, New York, NY, USA, SIGGRAPH '02, 547–556.

Debevec, P., Graham, P., Busch, J., and Bolas, M. 2012. A single-shot light probe. In *ACM SIGGRAPH 2012 Talks*, ACM, New York, NY, USA, SIGGRAPH '12, 10:1–10:1.

Debevec, P. 1998. 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*, ACM, New York, NY, USA, SIGGRAPH '98, 189–198.

Gu, J., and Liu, C. 2012. Discriminative illumination: Perpixel classification of raw materials based on optimal projections of spectral brdf. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, 797–804.

Hamon, P.-L., Harmer, J., Penn, S., and Scapel, N. 2014. Gravity: Motion control and face integration. In *ACM SIGGRAPH 2014 Talks*, ACM, New York, NY, USA, SIGGRAPH '14, 35:1–35:1.

Kang, S. B., Uyttendaele, M., Winder, S., and Szeliski, R. 2003. High dynamic range video. *ACM Trans. Graph. 22*, 3 (July), 319–325.

Real environment | Lighting capture | Lighting reproduction | Final composite

Figure 15: *Frames from an image sequence where a subject moves through an environment, causing dynamic lighting from fluorescent, incandescent, and LED lights with various gels and diffusers. The last column composites the subject onto a clean plate of the set.*

Kawakami, R., Zhao, H., Tan, R. T., and Ikeuchi, K. 2013. Camera spectral sensitivity and white balance estimation from sky images. *Int. J. Comput. Vision 105*, 3 (Dec.), 187–204.

Kider, Jr., J. T., Knowlton, D., Newlin, J., Li, Y. K., and Greenberg, D. P. 2014. A framework for the experimental comparison of solar and skydome illumination. *ACM Trans. Graph. 33*, 6 (Nov.), 180:1–180:12.

Kitahara, M., Okabe, T., Fuchs, C., and Lensch, H. 2015. 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, 303–313.

McCamy, C., Marcus, H., and Davidson, J. 1976. A color-rendition chart. *Journal of Applied Photographic Engineering 2*, 3 (June), 95–99.

Reinhard, E., Ward, G., Pattanaik, S., and Debevec, P. 2005. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting (The Morgan Kaufmann Series in Computer Graphics)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

Rump, M., Zinke, A., and Klein, R. 2011. Practical spectral characterization of trichromatic cameras. *ACM Trans. Graph. 30*, 6 (Dec.), 170:1–170:10.

Shi, J., Yu, H., Huang, X., Chen, Z., and Tai, Y., 2014. Illuminant spectrum estimation using a digital color camera and a color chart.

Stumpfel, J., Tchou, C., Jones, A., Hawkins, T., Wenger, A., and Debevec, P. 2004. Direct hdr capture of the sun and sky. In *Proceedings of the 3rd International Conference on Computer Graphics, Virtual Reality, Visualisation and Interaction in Africa*, ACM, New York, NY, USA, AFRIGRAPH '04, 145–149.

Stumpfel, J. 2004. *HDR Lighting Capture of the Sky and Sun*. Master's thesis, California Institute of Technology, Pasadena, California.

Tominaga, S., and Fukuda, T., 2007. Omnidirectional scene illuminant estimation using a multispectral imaging system.

Tominaga, S., and Tanaka, N. 2001. Measurement of omnidirectional light distribution by a mirrored ball. In *The 9th Color Imaging Conference: Color Science and Engineering: Systems, Technologies, Applications, CIC 2001, Scottsdale, Arizona, USA, November 6-9, 2001*, 22–26.

Tominaga, S., and Tanaka, N. 2006. Omnidirectional scene illuminant estimation using a mirrored ball. *Journal of Imaging Science and Technology 50*, 3 (May), 217–227.

Unger, J., and Gustavson, S., 2007. High-dynamic-range video for photometric measurement of illumination.

Wenger, A., Hawkins, T., and Debevec, P. 2003. Optimizing color matching in a lighting reproduction system for complex subject and illuminant spectra. In *Proceedings of the 14th Eurographics Workshop on Rendering*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, EGRW '03, 249–259.