Multispectral Imaging Using Multiplexed Illumination

Jong-Il Park, Moon-Hyun Lee, Michael D. Grossberg[†], and Shree K. Nayar[‡]

Hanyang University, †City University of New York, ‡Columbia University

{jipark,fly4moon}@hanyang.ac.kr,grossberg@cs.ccny.cuny.edu,nayar@cs.columbia.edu

Abstract

Many vision tasks such as scene segmentation, or the recognition of materials within a scene, become considerably easier when it is possible to measure the spectral reflectance of scene surfaces. In this paper, we present an efficient and robust approach for recovering spectral reflectance in a scene that combines the advantages of using multiple spectral sources and a multispectral camera. We have implemented a system based on this approach using a cluster of light sources with different spectra to illuminate the scene and a conventional RGB camera to acquire images. Rather than sequentially activating the sources, we have developed a novel technique to determine the optimal multiplexing sequence of spectral sources so as to minimize the number of acquired images. We use our recovered spectral measurements to recover the continuous spectral reflectance for each scene point by using a linear model for spectral reflectance. Our imaging system can produce multispectral videos of scenes at 30fps. We demonstrate the effectiveness of our system through extensive evaluation. As a demonstration, we present the results of applying data recovered by our system to material segmentation and spectral relighting.

1. Introduction

One approach to tasks such segmentation, recognition and classification is based on exploiting the variations in appearance due to the material composition of different scene elements. The capability of detecting such variations is important well beyond vision. Fields such as medical imaging, automatic inspection, and remote sensing have developed methods to detect anomalies such as, respectively, skin disease, food contamination, and deforestation using specific devices sensitive to corresponding spectral variations in surface reflectance[20, 16, 2].

Recovery of surface reflectance is impossible exclusively with conventional RGB cameras. The appearance of a captured image depends on both the illumination spectrum and the spectral reflectances of objects in the scene. Even if the illumination spectrum is known, an RGB camera provides only 3 measurements which is insuffcient to recover the spectral reflectance.

A wide variety of methods have been developed for estimating the spectral reflectance of a scene. For a static scene with fixed illumination, the spectral sensitivity of the camera can be varied over time. If the illumination spectrum is known, the multispectral reflectance of the scene can be determined. In the case of a dynamic scene, however, spectral reflectance must be measured with high temporal resolution. Today, there exist no methods for capturing multispectral videos in realtime.

Instead of obtaining the spectral reflectance exclusively with a multispectral camera, the spectrum of the illumination can be modulated temporally, to provide a multispectral light source. From a practical standpoint, there are two critical advantages. First, it is easier to create an illumination source with rapidly changing spectra than a camera with rapidly changing spectral sensitivity. Second, if there are M camera channels and N spectrally distinct illuminations, the number of effective channels is MN. This multiplicative effect dramatically increases the number of independent measurements with a minor increase in system complexity. Varying either the spectral sensitivity of the camera or the illumination typically comes at the cost of lowering the spatial resolution or the frame-rate of the acquired data. Our approach is to view the multispectral imaging problem as one of minimizing the number of measurements needed to obtain spectral reflectance. This is done in two ways. First, an illumination source is constructed using several sources with different spectra. Rather than sequentially activating these sources, we find the optimal way in which they can be multiplexed so as to minimize the number of images acquired while maintaining high signal-to-noise ratio. Second, we use our finite number of spectral measurements to recover a continuous spectral reflectance by using a set of basis functions. The end result is a video-rate multispectral imaging system.

The key technical contributions of our work can be summarized as follows.

Spectral Illumination Multiplexing: For a given set of sources with different spectra (some narrow and others wide) and the number of allowable measurements, we de-

termine the multiplexing sequence of the sources that minimizes least squares error in spectrum estimation.

Model-Based Spectral Reconstruction: We apply a simple, empirical linear model to estimate the full spectral reflectance at each pixel from the small number of multiplexed measurements. Since the model is linear, the reconstruction is efficient and stable.

Multispectral Imaging System: We present a practical, low-cost system that is capable of capturing multispectral videos at 30 fps. The system uses off-the-shelf components, including a commodity RGB video camera and a set of LED light sources.

We begin with a brief review of related work in Section 2. Then, we present our model-based method for spectral reconstruction in Section 3. Section 4 describes our algorithm for finding the multiplexing sequence that minimizes errors in spectral reconstruction. In Section 5, we describe our imaging system and verify its accuracy using surfaces with known spectral reflectances. In Section 6, we present several experiments with real scenes that demonstrate the advantages of using multispectral imaging. Finally, in Section 7 we discuss the limitations of our approach and our plans for addressing the limitations.

2. Related Work

A multispectral video has two spatial dimensions, one spectral dimension and one temporal dimension – it can be viewed as a 3D volume of measurements that varies with time. This 3D volume can be acquired using a wide variety of imaging devices. These devices differ in the mechanisms they use to separate (over space and/or time) the incoming light into its spectral components. This separation is done using prisms, diffraction gratings, tunable filters, or gel filters (see [10] for a survey of available technologies.) Existing systems differ in terms of how they trade off spatial and temporal resolution to obtain multispectral measurements for each point in the field of view.

For static scenes, the temporal dimension can be used to measure the spectrum. For example, a spectrometer can be used to measure the spectrum of a single point, and the entire field of view is scanned over time [29]. Other devices capture a set of monochrome 2D images by changing the spectral sensitivity of the camera over time. A popular way of changing the spectral sensitivity is by using tunable filters [9]. Some hybrid approaches can simultaneously scan a static scene with respect to space and spectrum by modifying a commodity camera [21]. High-cost devices that use complex optics and custom photo-sensors have been developed for remote sensing that can acquire hyperspectral videos of dynamic scenes (for example, see [7].)

All these systems trade off temporal and/or spatial resolution to capture spectral information. A noteworthy approach that avoids this trade-off is the IRODORI system that can capture 6-band HDTV video using an optical splitter, color filters, and two RGB sensors [30]. However, the system does not deal with how to recover the continuous spectrum from the 6 color measurements. It should be noted that our use of spectral multiplexed illumination is complementary to any multispectral camera.

Recovering reflectances from single illumination with ordinary RGB sensors has been proposed [25, 15]. They are based on restrictive assumptions on surface reflectance and thus the achievable performance of reconstruction is limited.

In general, the space of spectral reflectance curves for natural materials can be well-approximated with a low-dimensional linear model [13]. A number of linear models have been proposed based on measured spectral reflectances (see [24] for a review.) A few non-linear models have also been proposed [4, 27]. In our work, we account for the non-linearities of our devices by pre-calibrating them. While recovering the parameters of a linear model for spectral reconstruction, we need to enforce the constraint that spectral reflectance is positive [3, 5, 23]. We enforce this constraint in our reconstruction using quadratic programming.

Multispectral imaging has been used for color analysis, color constancy, object recognition, and the analysis of works of art (for examples, see [1, 17, 6].) There have been a number of works that demonstrate the importance of using spectral reflectance for image-based rendering and relighting [19, 11, 28]. All the above works have relied on multispectral imagers to measure the spectral curves. There has also been some work combining mulitspectral illumination with imaging. A hybrid camera-LED system was developed to match a target's spectral reflectance with a database of reflectances using an image sensor that performs correlations [12, 8]. Color illumination multiplexing using a projector and a color camera was proposed to obtain data for RGB relighting[22]. Vrhel [26] has developed a mouse-like device that uses LED illumination sources and photo-sensors to measure the spectral reflectance of a point. In contrast to these prior works, we use a conventional RGB camera to recover the spectral curves for all points in an image. Furthermore, our system uses multiplexed illumination and a linear spectral reflectance model to produce multispectral video at 30fps.

3. Model-Based Spectral Reconstruction

The concept underlying our active multispectral imaging system is shown in Fig. 1. The scene is illuminated using a sequence of distinct multiplexed illuminations and an image is captured for each illumination using a synchronized RGB camera. The value I_{mn} measured at a pixel in the mth channel, for the nth illumination, is given by

$$I_{mn} = \int s(\lambda)c_m(\lambda)p_n(\lambda)d\lambda, \qquad (1)$$

where $p_n(\lambda)$ is the spectral power distribution of the nth illumination, $s(\lambda)$ is the spectral reflectance of the scene

point, and $c_m(\lambda)$ is the spectral response of the camera's mth color channel. If there are M color channels and N illuminations, then by stacking the N images together, we obtain a multiplexed multispectral image with MN channels.
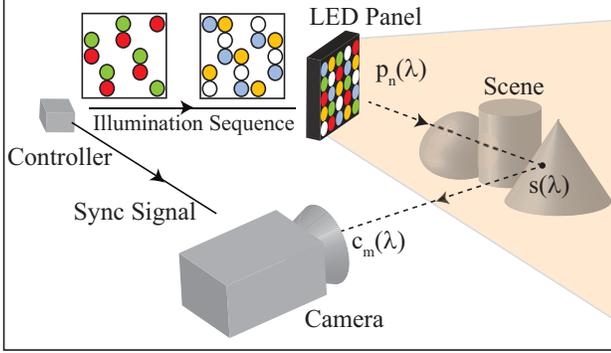

Figure 1. Illustration of our active multispectral imaging system. The scene is lit with a set of distinct illuminations and a synchronized RGB camera captures the corresponding images. By using rapidly changing multiplexed illuminations and processing the acquired video, we obtain a multispectral video of the scene.

Some assumptions are needed to reconstruct the full continuous spectral reflectance $s(\lambda)$ from a finite set of measurements. If the spectral curves are arbitrarily complex, MN must be large. However, in order to capture the spectral reflectance of dynamic scenes (at video rate) with an RGB camera ($M = 3$), the number of multiplexed illuminations must be fairly small, with $N = 2$ or $N = 3$.

Fortunately, the spectral reflectance of most real-world surfaces can be well-approximated using a low-parameter linear model. One such linear model is the set of orthogonal spectral basis functions $b_k(\lambda)$ proposed by Parkkinen *et al.* [18]. These basis functions are eigenvectors of a correlation matrix derived from a database consisting of the measured spectral reflectances of 1257 Munsell color chips. We have found empirically that this model gives fairly accurate spectral reconstructions for a wide range of real-world materials. The model can be written as

$$s(\lambda) = \sum_{k=1}^{K_s} \sigma_k b_k(\lambda) , \qquad (2)$$

where σ_k are scalar coefficients and K_s is the number of parameters of the model. By substituting Eq. 2 in Eq. 1, we get a set of MN equations,

$$I_{mn} = \sum_{k=1}^{K_s} \sigma_k \int b_k(\lambda)c_m(\lambda)p_n(\lambda)d\lambda , \qquad (3)$$

which are linear in the parameters σ_k.

To estimate these parameters from a set of images, we first determine via calibration the spectral response $c_m(\lambda)$ of each of the camera's color channels and the spectrum $p_n(\lambda)$ of each of the multiplexed illuminations. This calibration is done using a spectrometer and known reflectance targets. We can then absorb the known quantities in Eq. 3 into MNK_s coefficients: $f_{mnk} = \int b_k(\lambda)c_m(\lambda)p_n(\lambda)d\lambda$. If we substitute f_{mnk} into Eq. 3 and rewrite the equation in matrix form, we obtain

$$\mathbf{F}\sigma = \mathbf{I} . \qquad (4)$$

When $\mathbf{F^T F}$ is invertible, we can get a least squares solution: $\sigma = (\mathbf{F^T F})^{-1}\mathbf{F^T I}$. However, it is possible that the least squares solution results in a recovered spectral reflectance that is negative for some wavelengths. Since any spectral reflectance function must be positive, the following condition must be satisfied:

$$s(\lambda) = \sum_{k=1}^{K_s} \sigma_k b_k(\lambda) \geq 0, \ \ for\ all\ \lambda . \qquad (5)$$

To this end, we reformulate the problem as a constrained minimization as follows:

$$\sigma^+ = \arg\min_\sigma |\mathbf{F}\sigma - \mathbf{I}|^2 , \ \ \text{subject to } \mathbf{A}\sigma \geq \mathbf{0}, \qquad (6)$$

where $\mathbf{A}_{\ell k} = b_k(\lambda_\ell)$ with $1 \leq \ell \leq L$, and $1 \leq k \leq K_s$. This optimization can be solved using quadratic programming.

A solution to the constrained quadratic minimization in Eq. 6 may not be numerically stable if $\mathbf{F}$ has rank lower than K_s, or if $\mathbf{F}$ has a large condition number. In such cases, we cannot expect a reasonable solution without imposing further constraints. Since real-world spectral reflectances tend to be smooth, we impose a smoothness constraint. We do this by penalizing large values for the second derivative of the spectral reflectance with respect to λ:

$$\min_\sigma \left[|\mathbf{F}\sigma - \mathbf{I}|^2 + \alpha \left| \frac{\partial^2 s(\lambda)}{\partial \lambda^2} \right|^2 \right] , \qquad (7)$$

where α is a smoothness parameter. Then, the optimization can be written as

$$\min_\sigma \left| \tilde{\mathbf{F}}\sigma - \tilde{\mathbf{I}} \right|^2 , \ \ \text{subject to } \mathbf{A}\sigma \geq \mathbf{0}, \qquad (8)$$

where $\tilde{\mathbf{F}} = [\ \mathbf{F^T}\ \alpha\mathbf{P^T}\ ]^\mathbf{T}$ with $\mathbf{P}_{\ell k} = \partial^2 b_k(\lambda_\ell)/\partial \lambda^2$, $1 \leq \ell \leq L$, $1 \leq k \leq K_s$, and $\tilde{\mathbf{I}} = [\mathbf{I^T}\ \mathbf{0}]^\mathbf{T}$. This regularized minimization can also be solved with quadratic programming. In our implementation, we used the `quadprog` routine of Matlab.

Both the numbers of parameters K_s and α can be thought of as different ways to regularize the spectral recovery. K_s conditions the recovery based on the empirical reflectances, while α simply forces the solution to be smooth. By using a combination of both we can robustly recover a reasonable estimate of the full spectrum, even with a small number of channels ($M = 3$) and a small number of images ($N = 2$).

4. Multiplexed Spectral Illumination

Capturing multispectral images at video rate (30 fps) requires a minimum number of distinct illuminations. Moreover, to ensure our measurements have a good signal-to-noise ratio, the irradiance received by the imager must be high. We now describe a method for finding distinct illuminations that satisfy these conditions.

4.1. Finding the Multiplexed Illuminations

To rapidly vary the spectral illumination, we use a compound light source made up of many source elements. Such a light source can be constructed as a cluster of Q types of sources, each type with a distinct spectral curve $l_q(\lambda)$. The illumination used for the nth captured image is a weighted sum of the Q sources. The weights are denoted by d_{nq}, where $0 \leq d_{nq} \leq 1$. In this notation, $d_{11} = 0$ and $d_{11} = 1$ correspond to the sources of the first type being fully turned off and turned on in the first frame, respectively. The spectrum of the compound light source for frame n is

$$p_n(\lambda) = \sum_{q=1}^{Q} d_{nq} l_q(\lambda). \qquad (9)$$

If we substitute the above expression for p_n in Eq. 3, the measured value at a pixel for camera channel m and frame n can be written as

$$
\begin{aligned}
I_{mn} &= \int \sum_{k=1}^{K_s} \sigma_k b_k(\lambda) c_m(\lambda) \sum_{q=1}^{Q} d_{nq} l_q(\lambda) d\lambda, \quad (10) \\
&= \sum_{k=1}^{K_s} \sum_{q=1}^{Q} \sigma_k d_{nq} g_{kmq}, \qquad (11)
\end{aligned}
$$

where $g_{kmq} = \int b_k(\lambda) c_m(\lambda) l_q(\lambda) d\lambda$. Note that g_{kmq} does not depend on the scene or the illuminations used. In particular, for a given basis $\{b_k\}$, a set of camera channels $\{c_m\}$, a set of light source spectra $\{l_q\}$, and a fixed number of frames N, we can find the weights $\mathbf{d} = \{d_{nq}\}$ that minimize the errors in the reconstruction of a representative set of known spectral reflectances. In short, we can find the optimal weights $\mathbf{d}_{opt}$ given by

$$\text{argmin}_{d_1, \cdots, d_N} \sum_{e=1}^{E} \sum_{m=1}^{M} \sum_{n=1}^{N} \left[\sum_{k=1}^{K_s} \sum_{q=1}^{Q} \sigma_{k,e} d_{nq} g_{kmq} - I_{mn,e} \right]^2,$$
(12)

where e is used to denote the representative materials and $I_{mn,e}$ is the ideal image intensity for the material e. The solution will depend on the representative materials used. In our implementation, we used the known reflectance curves of common objects (available from NYU at http://www.cns.nyu.edu/ftp/ltm/SSR/).

When d_{nq} is an indicator variable (i.e., $d_{nq} \in \{0, 1\}$), the sources can only be fully on or fully off. In this case, $\mathbf{d}_{opt}$ can be found using brute force search over all the variables. When the number Q of different types of sources is small, the search can be done in reasonable time.
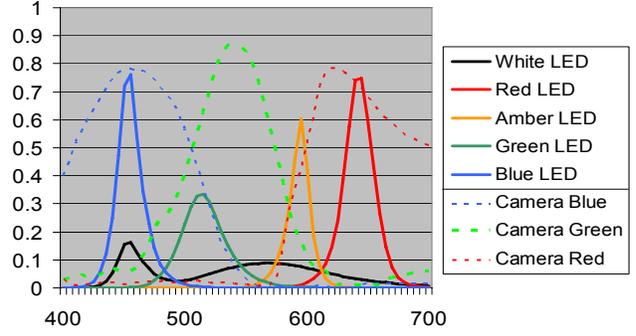

Figure 2. The spectra of the 5 types of LEDs (solid lines) and the spectral responses of the three color channels of the PointGrey Dragonfly Express camera (dashed lines) used in our system.
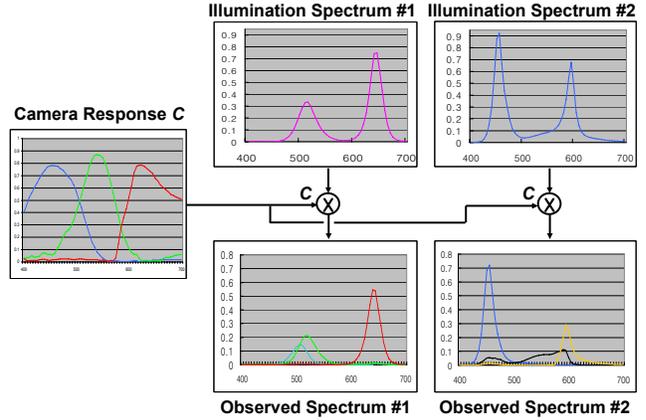

Figure 3. Spectra of an optimal pair of multiplexed illuminations and the corresponding spectra sensed by the color camera.

4.2. Choosing the System Parameters

In our current implementation, the compound light source is a cluster of "white," "red," "amber," "green," and "blue" LEDs, i.e. $Q = 5$. To determine the multiplexed illuminations, we need to specify the camera responses c_m, the basis b_k, and the number of acquired frames N.

The power spectra of the five LEDs as well as the spectral responses of the camera's three channels ($M = 3$) are shown in Fig.2. For spectral reconstruction from the measured image brightnesses, we use the first 8 components of the Parkkinen basis ($K_s = 8$). The two multiplexed illuminations with the lowest reconstruction error (highest rank) for the object spectra were determined using Eq. 12. The spectra of these two illuminations and the corresponding spectra observed by the camera are shown in Fig.3.

The three pairs of multiplexed illuminations that were found to have the highest three ranks are shown in Fig.4. These illuminations were automatically found using the search algorithm and they have the following desirable attributes: (a) **Full Utilization:** The spectral response of each camera channel overlaps with at least one component light source (LED) in each of the two frames. (b) **Independence:** For each camera channel, the spectrum of the illumination is different for the two frames. (c) **Normalization:** The to-

rank	illumination 1	illumination 2
1		
2		
3		

Figure 4. The highest ranking 2-frame multiplexed illuminations (using 5 types of LED sources) found by the optimization algorithm.
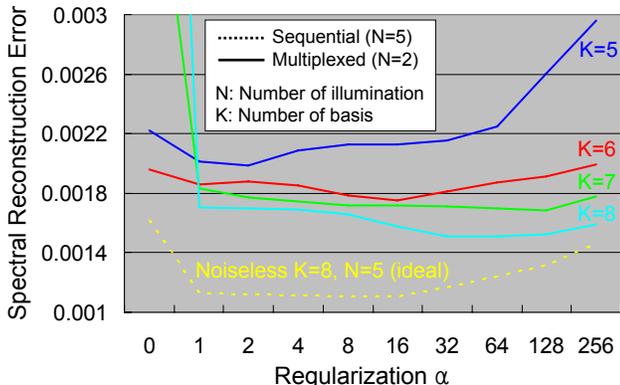

Figure 5. Simulation results that show how the spectral reconstruction error varies as a function of the smoothness parameter α, for different numbers K_s of basis functions. With just two multiplexed illuminations, we can get a spectral reconstruction error that is close to what can be achieved with sequential illumination and noiseless measurements (the lowest curve).

tal power of illumination in a frame is comparable to that obtained when half the sources are on.

In order to understand how the parameters K_s and α affect the performance of the spectral reconstruction algorithm, we have performed extensive simulations using the NYU common object spectra. Some of our results are shown in Fig. 5. Each line represents the RMS error in the recovered object spectra (using an optimal pair ($N = 2$) of illuminations) as a function of α (regularization term), for a given K_s (number of terms of the Parkinnen basis). For all simulations, we added noise using a noise model that was estimated for the PointGrey Dragonfly Express camera[1]. A similar set of simulations were done for $N = 3$ frames. Perhaps surprisingly, the errors did not decrease. In order to realize the benefits of using a larger number of frames, a larger number of sources (with distinct spectra) are required[2].

---

[1] The variance of the Gaussian noise was found to be proportional to the intensity level, similar to the results in [3]. For example, the variances were found to be 0.64 and 3.6 at intensity levels 50 and 200, respectively. The variances were also found to be slightly different for the three color channels.

[2] We have done some preliminary experiments with 8 off-the-shelf LED lamps. The practical difficulty has been that the spectra of the 8 lamps available were neither sufficiently distinct nor bright, to produce superior measurements with all 8 sources versus a subset of 5.
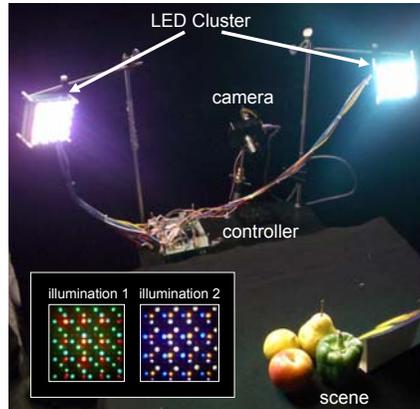

Figure 6. Our system for capturing multispectral video at 30 fps. It consists of two identical clusters of LEDs (each with 5 types of LEDs), a controller, and a 60 fps RGB camera. The two LED clusters emit the same multiplexed illumination spectra for each of the two acquired frames ($N = 2$). The two multiplexed illuminations are shown in the inset image.

We have done many more simulations to evaluate our reconstruction algorithm. Here, we briefly summarize our findings. We found that by carefully choosing the smoothness parameter and the number of basis functions, we could achieve accurate reconstructions of the spectral reflectances. We also found that the error could be further reduced by using more spectrally distinct sources. However, in this case, more frames need to be acquired. This can be done using a high-speed camera, but the resulting reduction in exposure time lowers the SNR. For these reasons, we decided to use just 2 frames and 5 types of light sources. As we will see, this setting is sufficient for accurately measuring spectral reflectances that are reasonably smooth. We believe that such measurements are adequate for a wide variety of application scenarios.

5. Multispectral Imaging System

We chose to use LEDs as the light sources as they are inexpensive, compact, and easy to control. We built two identical clusters, each consisting of 5 kinds of LEDs: "white," "red," "amber," "green," and "blue" (Itswell^{TM}). One of our clusters is shown in Fig.6. The total number of LEDs in the cluster is 12x12=144. Each cluster consists of 16 identical cells where each cell is composed of 9 independently controllable LEDs. The LEDs at the same locations in all cells are simultaneously controlled using a microprocessor (AVR^{TM}). The power spectra of the different types of LEDs were measured (see Fig.2) using a spectroradiometer (Luchem SPR-4001^{TM}).

A CCD camera with external trigger (PointGrey Dragonfly Express) was used for capturing the scenes. This camera works stably at 120fps. To increase image irradiance, we operate the camera at 60fps and acquire multispectral videos at 30fps ($N = 2$). The microprocessor controls the timing of the LED clusters and provides a synchronization

signal to the camera. The radiometric response of each camera channel was calibrated using the method in [14][3].

The accuracy of the spectral estimation was verified by using the Macbeth Color Checker chart. The number of basis functions K_s was 8. The smoothness parameter α was set to 64.0. In Fig.7, we show the known spectra (red solid lines) and the estimated spectra (dotted black lines) for a few of the color chips. Note that the estimated and actual spectra match very well. Similar results were obtained for all the remaining chips on the chart (See Fig.8 for quantitative results). These results provide strong evidence that our system can measure full spectral reflectances with reasonable accuracy.
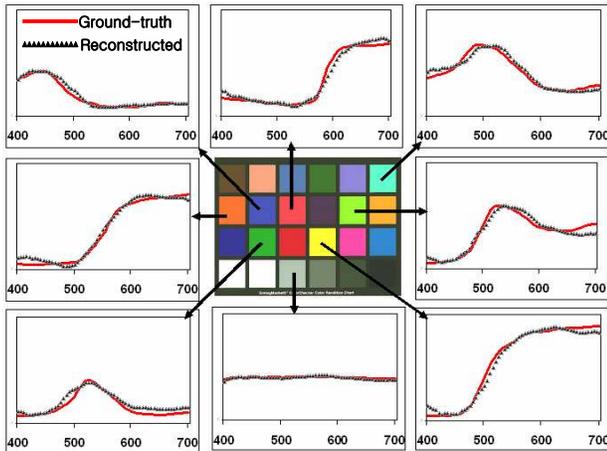

Figure 7. Spectral reflectances of color chips on the Macbeth chart measured using multiplexed illumination. The measured spectra are shown as dotted black lines and the actual spectra (ground truth) are shown as red solid lines.
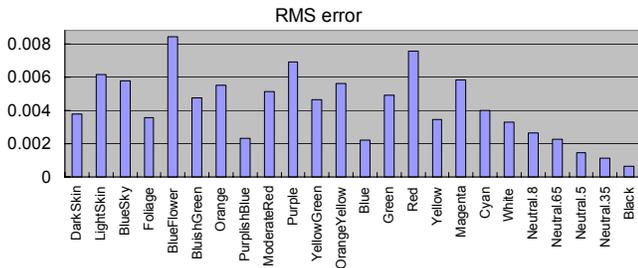

Figure 8. Average reconstruction errors in RMS for Macbeth chart.

6. Applications

Multispectral imaging can provide more measurements at each pixel than is provided by conventional RGB imaging. As a result, many vision applications stand to benefit. We present just two example applications due to limited space.

Object Detection/Segmentation: Fig. 9 shows an example of our system's ability to separate metameric radiances[11]. A star-shaped cosmetic patch has been applied to skin on a women's cheek. Scene points on the cosmetic patch and on the skin are virtually indistinguishable in the RGB images. Instead we consider a 6-component image obtained from two RGB images with the two images lit by optimal multiplexed illumination. From these two images the areas become easily distinguishable in 6-dimensional space. When the spectral reflectance is estimated using our reconstruction algorithm from Eq. 8 separation of the two types of scene point is enhanced as we see in the Fig. 10. We conjecture that the gain of multispectral imaging over the two RGB case comes from the empirical linear model providing a more spectrally balanced metric.

Relighting: Using the spectral reflectance $s(\lambda)$ estimated by our system at each scene point the scene can be spectrally relit accurately. Figs. 11 and 12 show the advantage of multispectral relighting over RGB relighting.

7. Discussion

We conclude by discussing the limitations of our approach and our plans for future work.

Temporal Resolution: Our current system captures multispectral video at 30fps. In principle, the frame-rate can be increased using high-speed cameras. However, as the frame-rate increases, the integration time of the detector decreases requiring more light to maintain the SNR. It is possible that we can obtain our multiplexing sequence from a bright uniform white light source which passes through a tunable filter. One way to address this problem without changing the light source is to use the observation that the spectral reflectances of neighboring scene points tend to be highly correlated. This spatial coherence is commonly exploited for color demosaicing algorithms, and hence could also be applied here to reduce noise at higher frame-rates without significantly compromising spatial resolution. At any frame-rate, rapid enough object motion can result in artifacts at occlusion boundaries. We are currently working address this problem by applying a motion compensation algorithm.

Assumption of Smooth Spectra: Our model-based spectral recovery assumes the reflectance at a point to be reasonably smooth in the spectral domain. However, there are many applications where high spectral resolution is important. To this end, we are exploring the use of a much larger number of LED types. This requires us to have a more efficient way (than brute-force search) of finding the best multiplexed illuminations. Furthermore, in our implementation, the sources are either fully on or fully off. We would like to explore the benefits of finding the optimal brightnesses of the LEDs as well.

Limited Working Volume: Currently, the working volume of our system is a bit larger than the size of a human

---

[3]Since the radiometric responses of the left and right sides of the CCD (Kodak KAI-0340DC) are different (they use separate electronic processing units), we performed the radiometric calibration for each of the sides.
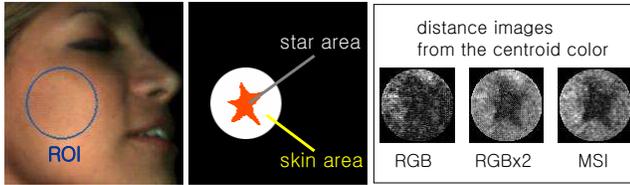
Figure 9. Distinguishing metameric appearances. A metameric cosmetic has been applied to a woman's cheek in a star-shaped area within an ROI. The black and white images show the relative distance to the means of the two classes. Both the multispectral reconstruction and two raw RGB images from multiplexed illumination show far superior distinguishablity than the conventional RGB measurement.
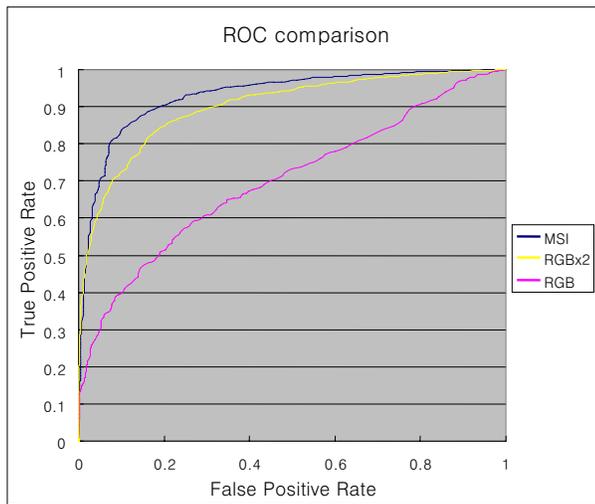

Figure 10. The Receiver Operating Characteristic curves for separation of metameric appearances using our multispectral reconstruction (MSI), measurements from two RGB images using multiplexed illumination (RGB×2), and a conventional RGB measurement (RGB).

head. In order to increase the working volume, we need to build larger panels of LEDs. There are several practical hurdles in achieving this. While large LED panels can be easily constructed, we need to ensure that the radiometric properties of the panel remain uniform within the entire working volume.

References

[1] A. Abrardo, L. Alparone, V. Cappellini, and A. Prosperi. Color constancy from multispectral images. In *Proc. IEEE ICIP*, volume 3, pages 570–574, 1999. 2

[2] L. Anderson, Y. Shimabukuro, R. Defries, and D. Morton. Assessment of deforestation in near real time over the brazilian amazon using multitemporal fraction images derived from terra modis. *IEEE Geoscience and Remote Sensing Letters*, 2(3):315–318, July 2005. 1

[3] K. Barnard and B. Funt. Camera characterization for color research. *Color Research and Applications,*, 27(3):152–163, 2002. 2, 5

[4] J. M. DiCarlo and B. A. Wandell. Spectral estimation theory: beyond linear but before bayesian. *J. Opt. Soc. Am. A,*, 20(7):1796–1806, July 2003. 2

[5] M. S. Drew and G. D. Finlayson. Spectral sharpening with positivity. *J. Opt. Soc. Am. A,*, 17(8):1361–1370, 2000. 2

[6] J. Farrell, J. Cupitt, D. Saunders, and B. A. Wandell. Estimating spectral reflectances of digital artwork. In *Proc. Chiba Conference of Multispectral Imaging*, 1999. 2

[7] D. W. Fletcher-Holmes and A. R. Harbvey. Real-time imaging with a hyperspectral fovea. *J. Opt. A: Pure Appl. Opt.,,* 7:S298–S302, 2005. 2

[8] I. Fryc, S. W. Brown, and Y. Ohno. Spectral matching with an led-based spectrally tunable light source. In *Proc. SPIE Fifth International Conference on Solid State Lighting*, volume 5941, pages 289–297, 2005. 2

[9] N. Gat. Imaging Spectroscopy Using Tunable Filters: A Review. In *Proc. SPIE Wavelet Applications VII*, volume 4056, pages 50–64, 2000. no. 1. 2

[10] A. R. Harvey, J. E. Beale, A. H. Greenaway, T. J. Hanlon, and J. W. Williams. Technology options for hyperspectral imaging. In *Proc. SPIE Imaging Spectrometry VI*, volume 4132, pages 13–24, 2000. 2

[11] G. M. Johnson and M. D. Fairchild. Full-spectral color calculations in realistic image synthesis. *IEEE Computer Graphics and Applications*, 19(4):47 – 53, July 1999. 2, 6

[12] A. Kimachi, H. Ikuta, Y. Fujiwara, M. Masumoto, and H. Matsuyama. Spectral matching imager using amplitude modulation- coded multispectral lightemitting diode illumination. *SPIE Opt. Eng.*, 43(4):975–985, April 2004. 2

[13] L. T. Maloney. Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *J. Opt. Soc. Am. A*, 3(10):1673–1683, 1986. 2

[14] T. Mitsunaga and S. Nayar. Radiometric Self Calibration. In *Proc. IEEE CVPR*, volume 1, pages 374–380, June 1999. 6

[15] P. Morovic and G. D. Finlayson. Metamer-set-based approach to estimating surface reflectance from camera rgb. *J. Opt. Soc. Am. A,*, 23(8):1814–1822, Aug. 2006. 2

[16] J. C. Noordam, W. van den Broek, and L. Buydens. Detection and classification of latent defects and diseases on raw french fries with multispectral imaging. *Journal of the Science of Food and Agriculture*, 85(13):2249–2259, October 2005. 1

[17] Z. Pan, G. Healey, M. Prasad, and B. Tromberg. Face recognition in hyperspectral images. *IEEE Trans. on PAMI,*, 25(12):1552–1560, 2003. 2

[18] J. Parkkinen, J. Hallikainen, and T. Jaaskelainen. Characteristic spectra of munsell colors. *J. Opt. Soc. Am. A,*, 6(2):318–322, 1989. 3

[19] M. S. Peercy. Linear color representations for full spectral rendering. In *Proc. International Conference on Computer Graphics and Interactive Techniques*, pages 191–198, 1993. 2

[20] M. Sambongi, M. Igarashi, T. Obi, M. Yamaguchi, N. Ohyama, M. Kobayashi, Y. Sano, S. Yoshida, and K. Gono. Analysis of spectral reflectance using normalization method from endoscopic spectroscopy system. *Optical Review*, 9(6):238–243, 2002. 1
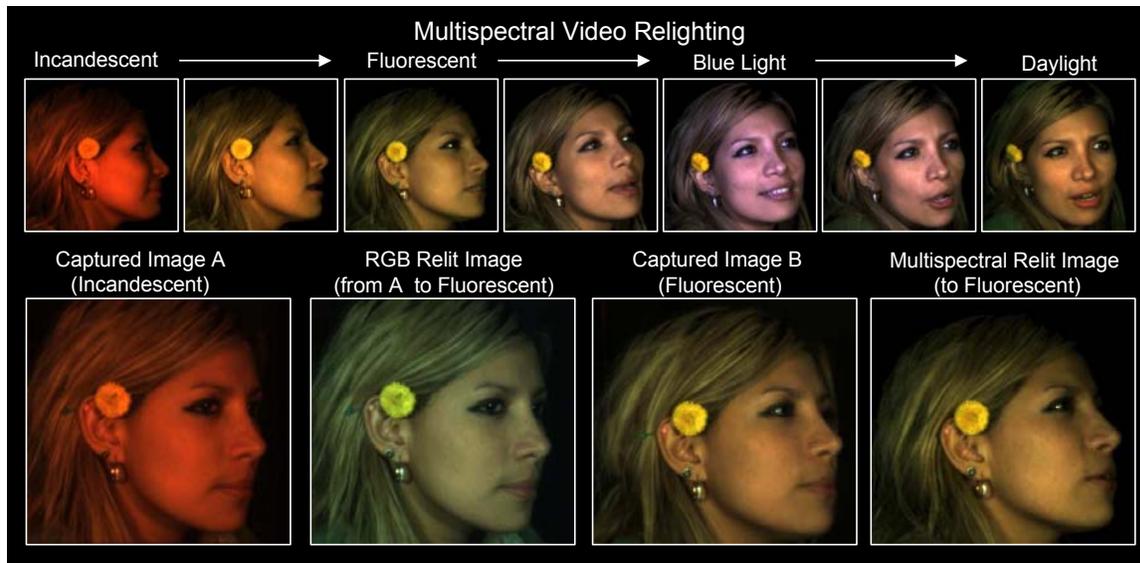
Figure 11. Top Row: The scene shown here was first captured using the proposed system and then spectrally relit to continuously move through four types of illuminations with very different spectra. Bottom Row: The first image (far left) was captured under incandescent illumination with an RGB camera. The second image shows the result of RGB relighting (based on per-channel transformation) the first image using fluorescent illumination. This image has a green tinge and its colors are dull compared to the third image which was captured with an RGB camera under fluorescent illumination (ground truth). The fourth image is the result of relighting with fluorescent illumination a multispectral image captured by our system. This image closely matches the ground truth.
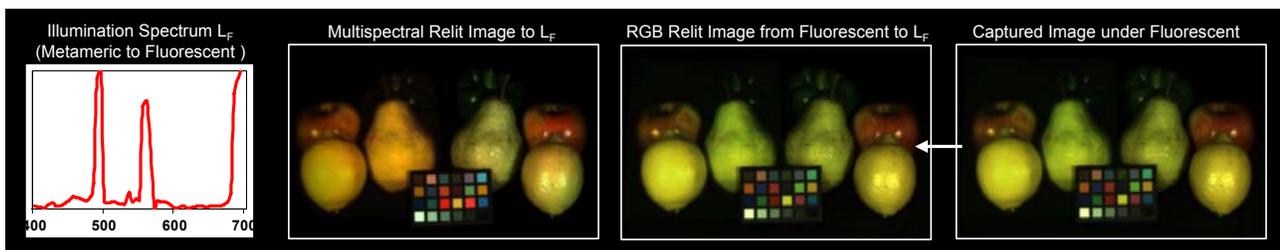

Figure 12. Image comparing RGB and multispectral relighting applied to a scene with three real fruits (on the right) and a printed image of the same fruits (on the left). The printed and real fruits are used as examples of metamers. (left) This target illumination spectrum L_F is constructed to be metameric to fluorescent illumination. (center, left) Image obtained by multispectral relighting with illumination L_F. Note that the real and printed fruits appear different. (center, right) The real and printed fruits remain indistinguishable after RGB relighting with L_F, the real image of the scene (right) captured under fluorescent illumination.

[21] Y. Y. Schechner and S. K. Nayar. Generalized Mosaicing: Wide Field of View Multispectral Imaging. *IEEE Trans. on PAMI,*, 24(10):1334–1348, Oct. 2002. 2

[22] Y. Y. Schechner, S. K. Nayar, and P. N. Belhumeur. A theory of multiplexed illumination. In *Proc. IEEE ICCV*, volume 2, pages 808–815, 2003. 2

[23] G. Sharma, J. Trussell, and M. J. Vrhel. Optimal nonnegative color scanning filters. *IEEE Trans. on IP,*, 7(1):129–133, 1998. 2

[24] D.-Y. Tzeng and R. S. Berns. A review of principal component analysis and its applications to color technology. *Color Research and Application,*, 30(2):84–98, 2005. 2

[25] C. van Trigt. Smoothest reflectance functions. i. definition and main results. *J. Opt. Soc. Am. A,*, 7(10):1891–1904, Oct. 1990. 2

[26] M. J. Vrhel. An led based spectrophotometric instrument. In *Proc. SPIE, Color Imaging: Device-Independent Color,*

Color Hardcopy, and Graphic Arts IV, volume 3648, pages 226–236, Dec. 1998. 2

[27] M. J. Vrhel and H. J. Trussell. Color device calibration: A mathematical formulation. *IEEE Trans. on IP,*, 8(12):1796–1806, Dec. 1999. 2

[28] G. Ward and E. Eydelberg-Vilshin. Picture perfect rgb rendering using spectral pre-filtering and sharp color primaries. In *Thirteenth Eurographics Workshop on Rendering*, pages 117–124, 2002. 2

[29] J. Wellman. Multispectral mapper: Imaging spectroscopy as applied to the mapping of earth resources. In *Proc. SPIE Imaging Spectroscopy*, volume 268, pages 64–73, 1981. 2

[30] M. Yamaguchi, H. Haneishi, H. Fukuda, J. Kishimoto, H. Kanazawa, M. Tsuchida, R. Iwama, and N. Ohyama. High-fidelity video and still-image communication based on spectral information: Natural vision system and its applications. In *Proc. SPIE/IS&T Electronic Imaging*, volume 6062, 2006. 2