

Spectral Imaging Using Basis Lights

Antony Lam¹

antony@nii.ac.jp

Art Subpa-Asa²

art@set.or.th

Imari Sato¹

imarik@nii.ac.jp

Takahiro Okabe³

okabe@ai.kyutech.ac.jp

Yoichi Sato⁴

ysato@iis.u-tokyo.ac.jp

¹ National Institute of Informatics
Tokyo, Japan

² The Stock Exchange of Thailand
Bangkok, Thailand

³ Kyushu Institute of Technology
Fukuoka, Japan

⁴ The University of Tokyo
Tokyo, Japan

Abstract

The spectral reflectance (SR) of objects provides innate information about material properties that have proven useful in applications such as classification, synthetic relighting, and medical imaging to name a few. General approaches to hyperspectral (HS) imaging include brute force capture of narrowband image stacks at consecutive wavelengths, use of conventional cameras with computational techniques for estimation of SR, use of specialized optics or filters, and active lighting. In particular, active lighting methods are fast and accurate. However, to our knowledge, past active lighting methods have not shown their illuminants to be optimal for SR recovery.

In this paper, we propose a new HS imaging method that utilizes optimally derived active illuminants in conjunction with a standard monochrome camera. We begin by showing mathematically, what kinds of illuminants would provide optimal SR recovery using the well-known observation that SR can be compactly represented using basis functions. We then demonstrate that such optimal illuminants can be well approximated in reality and show they can be used for accurate SR imaging of real scenes. In addition, we show that our method experiences little degradation even in the presence of unknown ambient light.

1 Introduction

The spectral reflectance (SR) of a scene provides a wealth of information about the material properties of objects and is useful for various applications such as classification, synthetic relighting, and medical imaging to name a few. As such, there have been many proposed methods for capturing the SR of scenes. These have included methods that perform brute force capture of narrowband images [1], computational techniques with conventional cameras [2], use of specialized optics or filters [3, 4, 5, 6, 7], and active lighting [8, 9, 10, 11].

The most straightforward approach to hyperspectral (HS) imaging is the brute force capture of a stack of narrowband images for all wavelengths but this is time consuming. Also,



Figure 1: Our setup (left) and the MacBeth Color Checker illuminated by “basis lights” (right). We are able to fully capture scene spectral reflectance despite the amount of ambient light shown here.

the limited light in narrowband imaging introduces noise and requires longer exposure times. Computational techniques that use conventional cameras do not require special equipment but may not be applicable in all cases because assumptions about the scene need to be made. There have also been recent developments in snapshot HS imaging systems but they currently suffer from lower spectral resolution. Active lighting methods allow for accurate SR recovery with fast imaging times. To our knowledge though, past active lighting methods have not shown their illuminants to be optimal with respect to SR recovery.

In this paper, we present an active lighting method for HS imaging using a conventional camera and optimal light sources we call “basis lights”. We also show that with only seven images, we can capture HS images with high accuracy. Our main contributions are:

1. We mathematically show that basis lights are optimal for SR recovery using the well-known observation that spectra can be compactly represented by basis functions [1, 2, 3, 4].
2. We show that although basis lights are theoretical, they can be well approximated in reality for accurate SR recovery.
3. We demonstrate that our method works well even with unknown ambient lighting.

As an added benefit, unlike brute force narrowband imaging, all of our images are wide-band which reduces issues with noise and the need for long exposure times.

2 Related Work

Maloney and Wandell used a conventional three-channel camera and were able to computationally recover SR under ambient light [5]. Their method has the advantage that specialized equipment is not needed but the computational method used for SR recovery required assuming constraints on the surfaces and lights. Later, Tominaga proposed a six-channel camera system by using band-pass filters in conjunction with a monochrome camera [6]. However, tasks such as image segmentation needed to be performed to estimate SR.

Balas, *et al.* presented a system for narrowband imaging at multiple wavelengths [7]. Specifically, they showed that a monochromator and CCD camera could be used for accurate capture of spectra in artistic and historic objects without making special assumptions about the scene. However, imaging one scene took about 2 minutes. Furthermore, narrowband imaging typically suffers from limited light which causes issues with noise and requires longer exposure times for better quality images.

Other methods such as DiCarlo, *et al.* used camera flash and a wideband camera to recover SR [8]. However, the use of only one illuminant in the form of a flash limits accuracy. Park, *et al.* [9] and Han, *et al.* [8] used multiple light sources for rapid imaging of SR.

Despite the impressive speeds of their methods, they did not show that their light sources were optimal for the task of SR recovery.¹ Thus accuracy may have been sacrificed.

In Chi, *et al.* [10] a set of available filters were selected to modify a single light source to produce good illuminants for spectral recovery. Their algorithm for selecting a good combination of filters from a very large search space worked well. However, their optimization algorithm requires days to find a good solution with respect to a particular equipment setup (e.g. camera response, light source spectrum) and the solution is not guaranteed to be a global optimum. In our work, our focus is on mathematically showing what optimal light sources for spectral recovery would be and that they can be well approximated in reality for accurate HS imaging.

Our work is also related to Xiong and Funt [18] where an experimental study in good matrix decompositions for representing spectra was conducted. In their work, they found that ICA and PCA provided the best decompositions for representing spectra. However, they argued that since ICA and PCA usually produces basis vectors with negative components, sensors based on these decompositions would be physically implausible. They advocated the use of non-negative matrix factorization as a less than optimal but good decomposition because negative values could be avoided. In our work, we show that decompositions such as PCA satisfy our optimal conditions and that the effects of negative light can be well simulated for HS image capture in real scenes.

It is noteworthy that in recent years, there have also been developments in the area of snapshot HS imaging [4, 5, 6, 7]. These methods can capture the SR of scenes at snapshot speeds but require complex specialized optics. In addition, accuracy and spectral resolution are sacrificed for speed.

3 Ideal Basis Lights

To consider how illuminants can be used to recover SR, let us first focus on a single surface point and how light diffusely reflects from it. Reflectance is modeled by the equation

$$I = \int s(\lambda)l(\lambda)d\lambda \quad (1)$$

where $s(\lambda)$ is the SR of the surface point at wavelength λ and $l(\lambda)$ is the illuminant at wavelength λ . Looking at Eq. (1), we see that s and l are analogous to vectors and similarly, computation of its reflectance is analogous to a dot product. This implies that if illuminant l were a basis vector, the observed reflectance I would be the projection of s onto l . In other words, I would be a basis coefficient.

So if we could project different illuminants l_m where each l_m would be equivalent to a different vector from the same basis, we could observe all the coefficients via reflectance. Having obtained all the coefficients, it would be possible to use the basis to reconstruct the SR s at the surface point.

In addition, we could project such “basis lights” onto an entire surface and then observe the reflectance off all surface points simultaneously which would allow for recovery of a complete HS image. Provided there exists a small number of basis vectors that could span the space of SR distributions, use of such basis vectors as light sources would provide a practical solution for HS image capture.

¹The method in [10] was shown to be optimal in terms of minimizing the number of needed images given their illuminants but not that the illuminants were optimal for accuracy.

We later show that PCA can provide six basis vectors capable of representing SR accurately. However another important question is, would eigenvectors be ideal for applying basis lights to HS image recovery where a camera is used to observe the surface reflectance? We now present an analysis showing that if the basis is orthonormal, we can ideally read the coefficients directly from the camera's pixel values.

We begin by first defining two relations. First, the SR of a surface point can be expressed in terms of a basis vectors as

$$s(\lambda) = \sum_{i=1}^N \sigma_n b_n(\lambda) \quad (2)$$

where $s(\lambda)$ is the SR at wavelength λ , $b_n(\lambda)$ is the n^{th} basis vector at wavelength λ , and σ_n is the coefficient associated with basis vector \vec{b}_n .

Next, we note that the the brightness $I_{p,m}$ of a given pixel p in an image taken using a grayscale camera and under lighting l_m follows the relation

$$I_{p,m} = \int s_p(\lambda) c(\lambda) l_m(\lambda) d\lambda \quad (3)$$

where s_p is the spectral distribution at pixel p in the scene and c is the response of the camera. This is basically the same as Eq. (1) except the camera response is also taken into account.

If we substitute s_p from Eq. (3), with Eq. (2) we get the relation

$$I_{p,m} = \sum_{i=1}^N \sigma_n \int b_n(\lambda) c(\lambda) l_m(\lambda) d\lambda = \sum_{i=1}^N \sigma_n g_{n,m} \quad (4)$$

The intensities $\vec{I}_p = [I_{p,1} \dots I_{p,M}]^T$ of pixel p under all lightings l_m from 1 to M can also be expressed in matrix form as

$$\vec{I}_p = \begin{bmatrix} g_{1,1} & \dots & g_{N,1} \\ \cdot & & \cdot \\ \cdot & \cdot & \cdot \\ \cdot & & \cdot \\ g_{1,M} & \dots & g_{N,M} \end{bmatrix} \begin{bmatrix} \sigma_1 \\ \cdot \\ \cdot \\ \cdot \\ \sigma_N \end{bmatrix} = G \vec{Q} \quad (5)$$

From Eq. (5), we see the basis coefficients \vec{Q} can be solved for from the pixel intensities \vec{I}_p and G . Ideally, if matrix G , has a condition number of 1, \vec{Q} can be solved for reliably.

To see what types of illuminants l_m are needed to make G have a condition number of 1, consider each of the elements of G , $g_{n,m}$. From Eq. (4), each $g_{n,m} = \int b_n(\lambda) c(\lambda) l_m(\lambda) d\lambda$. Now, if we set each illuminant l_m to be $l_m(\lambda) = b_m(\lambda)/c(\lambda)$, $g_{n,m} = \int b_n(\lambda) b_m(\lambda) d\lambda$.

Assuming the basis is orthonormal, if $n = m$ then $g_{n,m} = 1$; otherwise $g_{n,m} = 0$. Thus G becomes the identity matrix which has a condition number of 1. So setting all illuminants $l_m(\lambda) = b_m(\lambda)/c(\lambda)$ where the basis is orthonormal will provide the most effective illuminants for our proposed imaging system.

Mathematically, an orthonormal basis provides an ideal solution to the coefficients and makes the pixel intensities equivalent to the coefficients. Thus PCA is one approach that provides an ideal basis for HS image capture using a camera. Once the eigenvectors are divided by the camera response function they can be used as basis lights l_m . Capturing an image under illuminant l_m then gives the coefficients for basis vector \vec{b}_m in the pixel intensities of the image. With all the coefficients for each basis vector at every pixel, it is then possible to recover the SR of the entire scene. Fig. 2 gives an overview of this process.




Figure 2: Flowchart of our Method

4 Simulating Negative Light and Canceling Unknown Ambient Light

Since eigenvectors typically have negative components, the basis lights would be physically implausible. In reality, an additional illuminant is needed to simulate the effects of negative light. Thus six basis lights require seven illuminants for imaging.

To illustrate how negative light is simulated, let B be a matrix of basis vectors (with negative values) and \vec{b}_n be the n^{th} basis vector. Let v_{\min} be the minimum value in matrix B and $b_n^p(\lambda) = b_n(\lambda) + |v_{\min}|$. Then simulation of reflected negative light from basis vector \vec{b}_n on a surface point can be performed according to the following equation.

$$I_p = \int s_p(\lambda)c(\lambda)l_n^p(\lambda)d\lambda - \int s_p(\lambda)c(\lambda)l_{flat}(\lambda)d\lambda \quad (6)$$

where $s_p(\lambda)$ and $c(\lambda)$ are respectively, the surface point's SR and the camera response at wavelength λ . Illuminant $l_n^p(\lambda) = b_n^p(\lambda)/c(\lambda)$ and illuminant $l_{flat}(\lambda) = |v_{\min}|/c(\lambda)$. Note that l_n^p and l_{flat} would be the illuminants physically projected on the surface. Simulating negative light in this manner also does not affect the orthonormality of the basis lights because the value of I_p is the same as what would be obtained if negative light were physically possible.

The subtraction procedure described above also has the benefit of canceling out ambient light. To see why, let us consider Eq. (6) with ambient light added in. The joint effect of unknown ambient light a and light l_n^p can be modeled as $I_p = \int s_p(\lambda)c(\lambda)(l_n^p(\lambda) + a(\lambda))d\lambda - \int s_p(\lambda)c(\lambda)(l_{flat}(\lambda) + a(\lambda))d\lambda = \int s_p(\lambda)c(\lambda)l_n^p(\lambda)d\lambda - \int s_p(\lambda)c(\lambda)l_{flat}(\lambda)d\lambda$. Thus I_p is the same with or without ambient light present.

Simulating negative light also has advantages in numerical stability for recovering spectra. As mentioned in Sec. 2, Xiong and Funt [18] avoid the issue of negative values by using non-negative matrix factorization (NMF). We can view spectral reflectance as a vector $\vec{S} = B\vec{Q}$ where B is a column matrix of basis vectors and \vec{Q} is the vector of coefficients. If B is a non-negative basis, Xiong and Funt propose sensing the coefficients of \vec{S} using $\vec{Q} = B_T^+ \vec{S}$ where B_T^+ is the pseudoinverse of B with its negative values set to zero.² Their experiments found that B_T^+ was a good approximation to the true pseudoinverse of B . We repeated their experiment using 5,000 spectra from the CAVE Multispectral Database [19] and found that for six basis vectors, PCA produced a basis that represents the 5,000 spectra with an average root-mean-square-error (RMSE) of 0.011 while the NMF approach proposed by Xiong and Funt had an average error of 0.186. In addition, the PCA basis had a condition number of 1.0

²The matrix B can be entirely positive but its pseudoinverse could still have negative values.

while the NMF basis had a condition number of 3.4.³ Thus the use of negative values provides greater numerical stability. We will show in experiments, that simulation of negative light can be used to accurately produce real-life basis lights for HS image capture.⁴

5 Spectral Imaging with Basis Lights

Since we have already detailed how negative light can be simulated in Eq. (6), we will for the sake of clarity, assume negative light can really be projected in the rest for the paper. In our work, the computation of basis lights begins with performing PCA on a large SR dataset. We then select the top basis vectors needed to account for 99.5% of the variance as the basis. The basis lights are then constructed according to Sec. 3. With the basis lights computed, the imaging procedure is as described below and illustrated in Fig. 2.

1. For each basis light l_m from $m = 1$ to N :
 - (a) Illuminate the scene with basis light l_m .
 - (b) Image the scene under basis light l_m .
2. Let $I_{p,m}$ be the p^{th} pixel intensity under basis light l_m . Then for each p in $I_{p,m}$:
 - (a) Treat $\{I_{p,1}, I_{p,2}, \dots, I_{p,N}\}$ as the basis coefficient values for the p^{th} pixel.
 - (b) Use the coefficient values for the p^{th} pixel and basis vectors to reconstruct the SR for pixel p .
3. Output the resultant hyperspectral image.

6 Experiments

6.1 Experimental Setup

For our experiments, we used a Nikon Equalized Light Source (ELS) for generating basis lights and a Sony XC-EI50 digital camera for capturing images. The Nikon ELS is a programmable light source that can produce arbitrary light spectra based on user specifications. We note that our method does not specifically require a programmable light source. In fact, custom filters could be placed in front of a light source to produce basis lights. In this work, we used the ELS because of its convenience in testing different illuminants.

To specify appropriate basis lights, we first used all spectra from the CAVE Multispectral Image Database, performed PCA, and selected the top six eigenvectors to account for 99.5% of the variance in the dataset. Each basis vector was then divided by the response function for our camera and negative light was simulated as described in Sec. 4. The resultant light spectra were then used with the ELS to illuminate scenes for capture with our camera and recovery of scene surface reflectance spectra.

We also demonstrate the effectiveness of our HS capture method by predicting scene appearance under novel lighting. For our experiments in synthetic relighting, we captured ground truth RGB images by placing Edmund RGB filters in front of our camera while illuminating the scene under standard CIE illuminants. We then compared the ground truth RGB images against synthetic relighting of the HS images obtained from our basis lights imaging method.

³This was the NMF basis with the lowest condition number that we found out of 100 runs.

⁴We note that Nayar, *et al.* [10, 11] produced physical eigenfaces using a digital mirror array. However, their physical eigenvectors work in the spatial domain and are based on entirely different physical processes than the spectral domain eigenvectors of our work.




Figure 3: Real Basis Lights (obtained by measuring spectra of illuminants, multiplying by the camera response, and subtracting for negative light) vs. Computed Eigenvectors




Figure 4: RMSE Plots for the 24 spectra on the MacBeth Color Checker

6.2 Dealing with Errors in Real Life Basis Lights

At the end of Sec. 3, we described how ideal basis lights implies the pixel intensities in images captured under these lightings correspond exactly to basis coefficient values, which could be used to recover the spectra of a scene. Provided we can reproduce basis light spectra exactly, the pixel intensities would be exactly the coefficients. Unfortunately, limitations of equipment can introduce small errors in real-life basis lights. This would lead to pixel intensities deviating somewhat from the coefficient values one would expect under ideal conditions.

To compensate for the errors, we first measured the basis lights produced by the ELS using a spectrophotometer. Each measured basis light l_m , was then used to construct matrix G as described in Eq. (5). Then whenever we imaged scenes under our real-life basis lights, we took the resultant set of pixel intensities \vec{I}_p for each surface point and solved for the proper coefficient values as $\vec{Q} = G^{-1}\vec{I}_p$. Once the coefficients \vec{Q} were determined, we used them with the eigenvectors computed from PCA to recover the scene’s spectra.

6.3 Results

We first tested how accurately we could produce basis lights. Recall that a basis light is an eigenvector divided by the camera’s spectral response function. So to visualize how accurately we could produce basis lights, we measured all our real-life basis lights using a spectrophotometer and then multiplied each measured result by the camera response. In Fig. 3, it can be seen that the basis lights when multiplied by the camera response are very close to the computed eigenvectors from PCA. Assuming these basis lights to be ideal would provide good results but we found that slightly better results were obtained using the procedure described in Sec. 6.2 and we report those results.

For our experiments on imaging, we used a MacBeth Color Checker to test the accuracy of our method both with and without ambient light. The error between estimated spectra and ground truth was then computed using the RMSE. To test our method’s performance under ambient light, we turned on our office’s fluorescent lights directly over the scene.

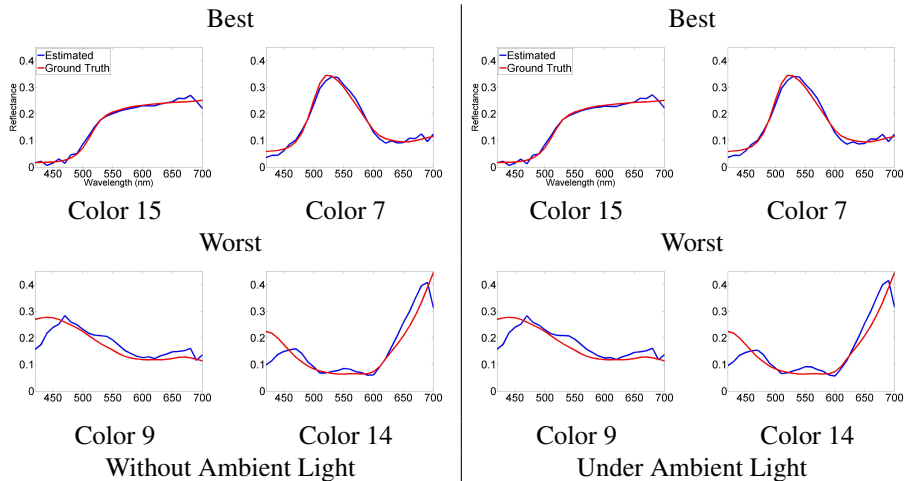


Figure 5: Best (15, 7) and Worst (9, 14) Case Estimated Spectra from MacBeth Color Checker




Figure 6: Real and simulated recovery had similar errors indicating color 14 was not well represented by the choice basis functions. Comparisons for other colors also indicate that the error trends observed largely come from the choice of basis functions and not from our proposed method itself.

Fig. 4 shows RMSE bars between estimated spectra and ground truth for the 24 colors on the MacBeth Color Checker. The average error when there is no ambient light is 0.0251 while with ambient light, the average error is 0.0255. Indeed, imaging under ambient light only increases the error by a negligible amount. To get an idea of what the errors look like visually, we show best and worst case recovered spectra in Fig. 5. The best cases are almost the same as the ground truth. For the worst cases, a major source of the error is due to the basis itself not being able to represent those particular spectra as well. For example, Fig. 6 shows that even in simulation, the recovery of color 14 has similar errors. We found similar results for our other worst cases, indicating that most of the error comes from our choice of basis and not our proposed method itself.

At this point, we would like to comment that it is possible to find a better basis for getting the highest accuracy on the MacBeth Color Checker. In fact, we obtained better results by using eigenvectors computed directly from the MacBeth Color Checker’s spectra. However, in the interest of making our basis lights more generally applicable, we chose to build our basis on the much larger collection of spectra found in the CAVE Database. Regardless of the choice of basis, we can see from the results that real life basis lights do very closely approximate basis vectors and allow for accurate HS imaging.

Results from our relighting tests can be seen in Figs. 7 and 8. At the top rows, ground

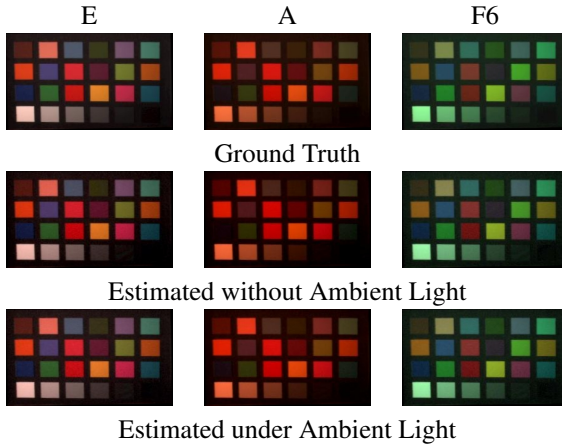


Figure 7: Relighting Results in RGB for the MacBeth Color Checker

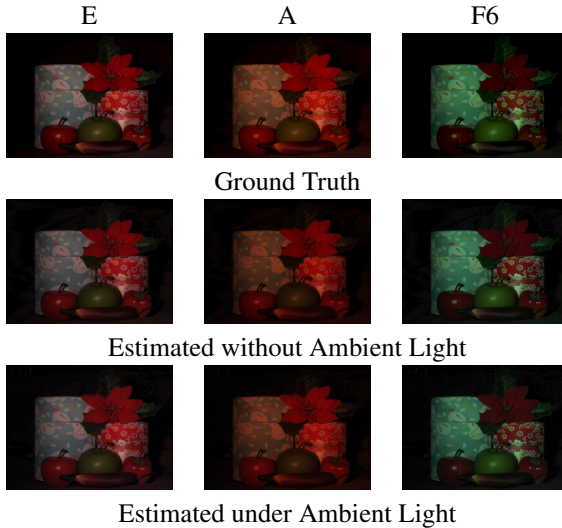


Figure 8: Relighting Results in RGB for a Scene

truth RGB images of the MacBeth Color Checker and a scene with real objects can be seen under CIE illuminants E, A, and F6. The estimated relighting results that follow show that both the no ambient light and ambient light cases are very close to the ground truth. We note that although color 14 (second row, fourth column) had the highest quantitative error, its predicted relightings in Fig. 7 is not far off. Visually, it is difficult to see the difference between the estimated RGB appearance and the ground truth. It is also remarkable that the estimated relightings under ambient light were so close to the no ambient light cases.

7 Conclusion

We proposed an active lighting based hyperspectral imaging method that uses optimal illuminants using the observation that spectra are well represented by basis functions. We mathematically derived the optimal illuminants and then showed how they could be used

effectively in a real setup. We then demonstrated that basis lights could be accurately reproduced in real life and showed their performance quantitatively. A remarkable finding was that imaging under ambient light (fluorescent lights directly over the scene) still yields highly accurate results. A drawback of our method is that accuracy is limited by the choice of basis. In future work, we hope to explore other types of basis functions or even the use of multiple bases that could improve accuracy.

References

- [1] C. Balas, V. Papadakis, N. Papadakis, A. Papadakis, E. Vazgiouraki, and G. Themelis. A novel hyper-spectral imaging apparatus for the non-destructive analysis of objects of artistic and historic value. *J. Cult. Herit.*, 4(1):330–337, 2003.
- [2] Cui Chi, Hyunjin Yoo, and Moshe Ben-Ezra. Multi-spectral imaging by optimized wide band illumination. *IJCV*, 86:140–151, January 2010. ISSN 0920-5691.
- [3] Jeffrey M. DiCarlo, Feng Xiao, and Brian A. Wandell. Illuminating illumination. In *CIC*, pages 27–34. IS&T/SID, 2001.
- [4] Liang Gao, Robert T. Kester, and Tomasz S. Tkaczyk. Compact image slicing spectrometer (ISS) for hyperspectral fluorescence microscopy. *OpEx*, 17:12293–12308, 2009.
- [5] Liang Gao, Robert T Kester, Nathan Hagen, and Tomasz S Tkaczyk. Snapshot image mapping spectrometer (IMS) with high sampling density for hyperspectral microscopy. *OpEx*, 18(14):14330–14344, 2010.
- [6] Alistair Gorman, David William Fletcher-Holmes, and Andrew Robert Harvey. Generalization of the lyot filter and its application to snapshot spectral imaging. *OpEx*, 18: 5602–5608, 2010.
- [7] Alistair Gorman, Gonzalo Muyo, and Andrew Robert Harvey. Snapshot spectral imaging using birefringent interferometry and image replication. In *ORS*. The Optical Society, 2010.
- [8] Shuai Han, Imari Sato, Takahiro Okabe, and Yoichi Sato. Fast spectral reflectance recovery using dlp projector. In *ACCV*, pages 323–335, 2010.
- [9] Te-Won Lee, Thomas Wachtler, and Terrence J. Sejnowski. The spectral independent components of natural scenes. In *First IEEE Intl. Workshop on Biologically Motivated Computer Vision*, BMVC, pages 527–534, London, UK, 2000. Springer-Verlag. ISBN 3-540-67560-4.
- [10] Laurence T. Maloney. Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *JOSA A*, 3(10):1673–1683, Oct 1986.
- [11] Laurence T. Maloney and Brian A. Wandell. Color constancy: a method for recovering surface spectral reflectance. *JOSA A*, 3(1):29–33, 1986.
- [12] David H. Marimont and Brian A. Wandell. Linear models of surface and illuminant spectra. *JOSA A*, 9(11):1905–1913, Nov 1992.

- [13] S. K. Nayar, V. Branzoi, and T. E. Boult. Programmable Imaging: Towards a Flexible Camera. *IJCV*, Oct 2006.
- [14] S.K. Nayar, V. Branzoi, and T. Boult. Programmable Imaging using a Digital Micromirror Array. In *CVPR*, volume I, pages 436–443. IEEE, Jun 2004.
- [15] Jong-Il Park, Moon-Hyun Lee, Michael D. Grossberg, and Shree K. Nayar. Multispectral imaging using multiplexed illumination. In *ICCV*, pages 1–8, oct. 2007.
- [16] J. P. S. Parkkinen, J. Hallikainen, and T. Jaaskelainen. Characteristic spectra of munsell colors. *JOSA A*, 6:318–322, 1989.
- [17] S. Tominaga. Multichannel vision system for estimating surface and illumination functions. *JOSA A*, 13(11):2163–2173, November 1996.
- [18] Weihua Xiong and Brian Funt. Independent component analysis and nonnegative linear model analysis. In *Proc. Illuminant and Reflectance Spectra, AIC Colour*, pages 503–506, 2005.
- [19] F. Yasuma, T. Mitsunaga, D. Iso, and S.K. Nayar. Generalized Assorted Pixel Camera: Post-Capture Control of Resolution, Dynamic Range and Spectrum. Technical report, Columbia University, Nov 2008.