Maximum Entropy Spectral Models for Color Constancy

Sandra Skaff and James J. Clark; Centre for Intelligent Machines, McGill University; Montreal, Canada

Abstract

Earlier work showed that maximum entropy models can be used to represent surface reflectance spectra of Munsell patches. Here, we introduce a new approach to color constancy which is based upon that work. To our knowledge, all color constancy approaches employing spectral models use linear basis function representations for surface and illuminant spectra. This means that a set of basis functions has to be specified in advance in these algorithms. The proposed maximum entropy approach does not require this a priori *information and therefore has a major advantage over other spectral based color constancy approaches. We show that a maximum entropy approach can be used to estimate surface and illuminant spectra given* only *camera sensor responses. We test our approach both in simulation and experiment. We also show that the performance of the proposed approach is similar to the most successful spectral based color constancy approach. This comparison is carried out in simulation in the presence of noise.*

Introduction

Color constancy is the ability of a vision system to compute a measure of a surface's color that is independent of the spectrum of the light incident on a surface. Obtaining such a measure is essential when using color as a cue in machine vision tasks, such as object recognition. This measure can be in 3D vector form (RGB, CMY, YIQ, etc.) or in spectral form. The latter constitutes the surface reflectance spectrum, which is the amount of light reflected off the surface at each wavelength.

Most recent work on color constancy has focused on using 3D color models [14, 15, 16] because the information they provide is sufficient for a large class of problems. However, color constancy using spectral models can be useful when there is a need for a more accurate representation of a color. For example, when classifying vegetation, the difference between the colors of leaves may be sufficiently small to require the extra information contained in the spectral models. Moreover, employing a spectral model for a surface color constitutes a universal representation across different color spaces (RGB, CMY, YIQ, etc.). This can be useful when images of the same scene are taken with cameras of sensor spectral sensitivities belonging to different color spaces. In this case, employing 3D color models of objects is cumbersome in the sense that different color spaces imply different sensor responses for the same object.

Spectral models have been used in a number of color constancy approaches [3, 4, 5, 11, 12, 13]. All these approaches have one major aspect in common. They represent surface and illuminant spectra by linear combinations of spectral basis functions. These basis functions are typically obtained by performing principal components analysis (PCA) on the sets of surface and illuminant spectra. This means that these sets should be available prior to applying the color constancy approach. However, these databases might not be available in advance. Even if they are available, they might not be consistent with the data used in a certain application.

We introduce a new maximum entropy spectral based approach to color constancy which does not require a set of basis functions to be specified in advance. We build upon previous work [1] in which maximum entropy models were successfully used to estimate Munsell patch reflectance spectra given only photoreceptor responses. In [1] the illuminant was assumed to be constant or white, which is not the case in this work. The use of maximum entropy models was inspired by Jaynes, who stated that a physical quantity frequently observed in practice will tend to a value that can be produced in the largest number of ways [2]. In the case of physical processes representing spectra, many surfaces observed in our everyday-life surroundings have spectra that have high entropy, as opposed to monochromatic surfaces which have low entropy spectra [1].

Since we are concerned with color constancy where the illuminant is unknown, we seek a suitable model for the illuminant spectrum. Applying Jaynes' argument, we can say that illuminants observed in our everyday-life surroundings have high entropy spectra, which therefore can be produced in the largest number of ways. This motivates us to represent the illuminant spectra with maximum entropy models as is the case for the surface reflectance spectra.

The paper is organized as follows. First, the proposed maximum entropy approach is explained and derived. Second, the performance of our approach is analyzed both in simulation and experiment. Third, the proposed approach is compared to the best color constancy algorithm employing spectral models. This comparison is performed in simulation in the presence of noise. Finally, the paper is ended with concluding remarks.

Maximum Entropy Spectral Based Color Constancy

A spectral based color constancy algorithm computes surface and illuminant spectra given sensor responses obtained from a camera. The responses can be computed by the following equation:

$$P_k = \sum_{\lambda=1}^{M} R_k(\lambda) s(\lambda) e(\lambda), \quad k = 1,2,...,p, \tag{1}$$

where p is the number of sensor classes, each denoted by k; λ denotes the wavelength, which is taken over the visible range. P_k is the computed sensor response for each sensor class k, each with a spectral sensitivity function $R_k(\lambda)$; $s(\lambda)$ is the surface reflectance spectrum; $e(\lambda)$ is the illuminant spectrum; M is the dimension of these spectra. Usually there are three sensor classes corresponding to each of the long-, medium-, and short-wavelength ranges.

The proposed color constancy approach aims at recovering surface and illuminant spectra by representing them using maxi-

mum entropy models, given only sensor responses. For simplicity, we illustrate the approach with the case of one surface patch in a scene illuminated by a single light source. Note that color constancy problems are addressed for the cases when there are at least two surface patches in the scene, which makes our example a hypothetical case.

In our hypothetical scene, $s(\lambda)$ is the surface spectrum and $e(\lambda)$ is the illuminant spectrum. We represent these spectra by probability distributions in order to compute their entropy. A light spectrum is a collection of photons that can be thought of as a histogram of photons over wavelength. Given a particular photon $photon_0$, the spectral reflectance then represents the probability of a wavelength given $photon_0$. Therefore $s(\lambda)$, for example, can be represented by the corresponding conditional probability distribution $p_s(\lambda|photon_0)$. The same argument can be applied to the illuminant spectrum whose probability distribution representation can be denoted by $p_e(\lambda|photon_0)$. In this paper we adopt the notations $p_s(\lambda)$ and $p_e(\lambda)$ for $p_s(\lambda|photon_0)$ and $p_e(\lambda|photon_0)$ for simplicity knowing that this does not affect the following derivations. Throughout this paper we may refer to the entropy of the probability distribution representation of a spectrum as the entropy of a spectrum for simplicity. $p_s(\lambda)$ and $p_e(\lambda)$ can be obtained from $s(\lambda)$ and $e(\lambda)$ by the following:

$$p_s(\lambda) = s(\lambda)/\sum_{\lambda=1}^{M} s(\lambda), \quad \lambda = 1,...,M, \tag{2a}$$

$$p_e(\lambda) = e(\lambda)/\sum_{\lambda=1}^{M} e(\lambda), \quad \lambda = 1,...,M. \tag{2b}$$

Our goal is to estimate surface and illuminant spectra by representing each by a maximum entropy model. Jaynes showed that given measurements in the form of expectations, which is the case for the sensor responses given by Equation 1, the probability distribution which maximizes the entropy can be computed and is in the form of a product of exponentials [2]:

$$\hat{p}_{s,e}(\lambda) = \frac{1}{Z}\prod_{k=1}^{p}\exp(\alpha_k R_k(\lambda)) = \frac{1}{Z}\exp(\sum_{k=1}^{p}(\alpha_k R_k(\lambda))), \tag{3}$$

where the scale factor Z is given by $\sum_{k=1}^{p}\exp(\sum_{k=1}^{p}(\alpha_k R_k(\lambda)))$. The α_k's are the Lagrange multipliers and they are determined so as to satisfy the constraints given in Equation 1. The probability distribution for which the maximum entropy solution is given is joint over the probability distribution representations of the surface and illuminant spectra and is denoted by $p_{s,e}(\lambda)$. The estimate of $p_{s,e}(\lambda)$ is denoted by $\hat{p}_{s,e}(\lambda)$, which can be obtained by solving for the Lagrange multipliers α_k analytically. However, due to the inherent complexity of an analytical approach, we resort to a numerical one. Therefore, we seek the joint probability distribution $p_{s,e}(\lambda)$ that maximizes the entropy H [10] given by:

$$H = -\sum_{\lambda=1}^{M} p_{s,e}(\lambda)\log p_{s,e}(\lambda). \tag{4}$$

The surface and illuminant spectra are assumed to be independent, and therefore the corresponding probability distributions $p_s(\lambda)$ and $p_e(\lambda)$ are independent. This assumption is valid as the surface and illuminant spectra are characteristic of the surface and light source respectively. Making use of this independence assumption, we can rewrite H [10] as:

$$H = -\sum_{\lambda=1}^{M} p_s(\lambda)\log p_s(\lambda) - \sum_{\lambda=1}^{M} p_e(\lambda)\log p_e(\lambda), \tag{5}$$

where $p_{s,e}(\lambda) = p_s(\lambda)p_e(\lambda)$. Therefore, we can now find $p_s(\lambda)$ and $p_e(\lambda)$ that maximize H, given the constraint in Equation 1.

After illustrating the maximum entropy approach with a one surface patch scene, we move on to a four surface patch scene, which is the case considered in this paper (Figure 1). The probability distributions representing the surface spectra $s_1(\lambda)$, $s_2(\lambda)$, $s_3(\lambda)$, and $s_4(\lambda)$ are denoted by $p_{s_1}(\lambda)$, $p_{s_2}(\lambda)$, $p_{s_3}(\lambda)$, and $p_{s_4}(\lambda)$ respectively and can be obtained by:

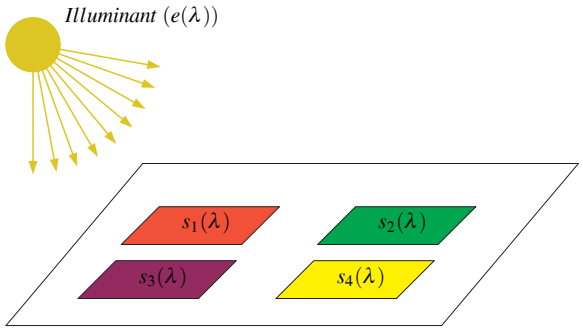

Figure 1. *A scene with four surface patches of reflectance spectra $s_1(\lambda)$, $s_2(\lambda)$, $s_3(\lambda)$, and $s_4(\lambda)$ illuminated by one light source of spectrum $e(\lambda)$.*

$$p_{s_1}(\lambda) = s_1(\lambda)/\sum_{\lambda=1}^{M} s_1(\lambda), \quad \lambda = 1,...,M, \tag{6a}$$

$$p_{s_2}(\lambda) = s_2(\lambda)/\sum_{\lambda=1}^{M} s_2(\lambda), \quad \lambda = 1,...,M, \tag{6b}$$

$$p_{s_3}(\lambda) = s_3(\lambda)/\sum_{\lambda=1}^{M} s_3(\lambda), \quad \lambda = 1,...,M, \tag{6c}$$

$$p_{s_4}(\lambda) = s_4(\lambda)/\sum_{\lambda=1}^{M} s_4(\lambda), \quad \lambda = 1,...,M. \tag{6d}$$

The probability distribution representation for $e(\lambda)$ can be obtained using Equation 2b. We seek the joint probability distribution $p_{s_1,s_2,s_3,s_4,e}(\lambda)$ that maximizes the entropy, denoted by H_4 in Equation 7, given the constraints imposed by the sensor responses (Equation 12).

$$H_4 = -\sum_{\lambda=1}^{M} p_{s_1,s_2,s_3,s_4,e}(\lambda)\log p_{s_1,s_2,s_3,s_4,e}(\lambda). \tag{7}$$

Since we assume that the surface spectra and illuminant spectra are independent, we can write:

$$p_{s_1,s_2,s_3,s_4,e}(\lambda) = p_{s_1,s_2,s_3,s_4}(\lambda)p_e(\lambda). \tag{8}$$

Moreover, we can assume that the surface spectra are independent from each other if the surfaces are drawn randomly from a set of surfaces with a wide range of hues. Since this is the case in this work, the probability distribution representations of these surface spectra are independent:

$$p_{s_1,s_2,s_3,s_4}(\lambda) = p_{s_1}(\lambda)p_{s_2}(\lambda)p_{s_3}(\lambda)p_{s_4}(\lambda). \tag{9}$$

The joint probability distribution can then be written as:

$$p_{s_1,s_2,s_3,s_4,e}(\lambda) = p_{s_1}(\lambda)p_{s_2}(\lambda)p_{s_3}(\lambda)p_{s_4}(\lambda)p_e(\lambda). \quad (10)$$

Therefore, we can express the entropy given in Equation 7 as:

$$\begin{aligned} H_4 \quad = \quad & -\sum_{\lambda=1}^{M} p_{s_1}(\lambda)\log p_{s_1}(\lambda) - \sum_{\lambda=1}^{M} p_{s_2}(\lambda)\log p_{s_2}(\lambda) \\ & -\sum_{\lambda=1}^{M} p_{s_3}(\lambda)\log p_{s_3}(\lambda) - \sum_{\lambda=1}^{M} p_{s_4}(\lambda)\log p_{s_4}(\lambda) \\ & -\sum_{\lambda=1}^{M} p_e(\lambda)\log p_e(\lambda). \end{aligned} \quad (11)$$

The constraints are given by:

$$P_1 - R^T s_1 e = 0, \quad (12a)$$
$$P_2 - R^T s_2 e = 0, \quad (12b)$$
$$P_3 - R^T s_3 e = 0, \quad (12c)$$
$$P_4 - R^T s_4 e = 0, \quad (12d)$$

where P_1, P_2, P_3 and P_4 are the sensor responses of the four surfaces and R is the sensor spectral sensitivities matrix of the camera. Equivalently to seeking the solution that maximizes the entropy H_4, we seek the solution that minimizes the negative of H_4. This means that we seek the solution that minimizes the negative of the sum of the entropies of the probability distribution representations of the spectra s_1, s_2, s_3, s_4, and e subject to the constraints given in Equations 12a, 12b, 12c,and 12d. For this purpose, we use a nonlinear constrained optimization algorithm (*fmincon* in Matlab).

Simulation Results

We test the performance of our approach on *matte* Munsell patches. These are representative of a wide range of hues encountered in printing. The corresponding spectra were measured by Parkkinen *et al.* [6]. The illuminant spectra we use are those of daylight and skylight, measured by Parkkinen and Silftsen [7], and a set of tungsten light spectra. The set of tungsten light spectra is composed of one spectrum at 2800K obtained from the manufacturer and a set of nine spectra with temperatures ranging from 2600K to 2700K and 2900K to 3500K, in steps of 100K, obtained from the IES lighting handbook [8]. For the tungsten light spectrum at 2800K, we use the one obtained from the manufacturer as it corresponds to the light we have in our laboratory. This would result in a correct comparison between the model spectra obtained in simulation and experiment for this particular illuminant.

The surface and illuminant spectra are multiplied to obtain the light falling on the sensor. To obtain the simulated sensor responses for the surface patches used, this spectrum of light is multiplied by the sensor spectral sensitivity curves of a Panasonic WV-CP410 camera (Equation 1 where $p = 3$). These sensitivity curves are obtained from the manufacturer. The wavelength range considered for these spectra is 400 nm to 700 nm, and is discretized into 5 nm bins. This yields a dimension of $M = 61$ in Equation 1 for each spectrum. Noise is added to the responses in the section entitled "Comparisons to Previous Work".

We analyze the performance of our approach in the case of four surface patches in the scene. These scenes are Mondrian as they are composed of several overlapping, matte (lambertian) patches. The light illuminating a Mondrian scene is assumed to be locally constant. This means that the spectral characteristics of the light vary slowly. The artificial scenes are constructed as follows. Twenty surface patch and illuminant pairs are chosen at random. Five 4-surface patch scenes are constructed from each of these pairs by selecting the three remaining patches randomly. Therefore, we construct a total of 100 scenes. The selection of scenes was done in this fashion so as to allow evaluation of the repeatability of spectral estimation in future work.

The model surface and illuminant spectra are computed for each scene. To evaluate the performance of our approach, we normalize each of the model and actual spectra to a maximum of one as we do not intend to recover intensity information. This is a common practice in solving color constancy problems. Then we compute the root mean square (RMS) errors between the normalized actual and model spectra for the surface patches and illuminant in each scene.

We plot the model and actual spectra for one surface patch and the illuminant in each of two scenes, one illuminated by skylight and one illuminated by tungsten light, in Figure 2. We denote the term "RMS error" by RMSE. We would like the reader to note that the model spectra comprise a product of three exponentials, one in each wavelength range (long, medium, and short). This agrees with Equation 3, which implies that the maximum entropy solution of the joint probability distribution of the surface and illuminant spectra is in the form of a product of exponentials. Therefore as one may expect, the daylight spectrum is better estimated by our model than the tungsten light spectrum. However, obtaining an illumination invariant spectral model of a surface is the main objective of our color constancy approach.

Experimental Results

While our emphasis in this paper is on the underlying theory, we present some results on real scenes. We consider scenes with two categories of surface patches: Munsell and construction paper.

Munsell Patches

We captured images of the pages in the Munsell Book of Color [9] with a Panasonic WV-CP410 camera. We obtained a 50x50 pixel sample from each of the Munsell patches using a segmentation algorithm. Since we assume Mondrian scenes where the illumination is locally constant, we average the 3D responses of all the pixels in a segmented patch to obtain one 3D response per patch. We ran the algorithm on the same 100 scenes used in simulation. The only exception was the illumination for which we used a tungsten light bulb of temperature 2800 K. Since we assume flat scenes, where there is no interreflection, we input the average responses corresponding to the patches in the scene into our algorithm. We plot the model and actual spectra for one surface patch and the illuminant from a scene in Figure 3. We can see from the figure that the approach provides a reasonable estimate for the surface patch spectrum shown. However, it does not do as well in estimating the illuminant spectrum. As mentioned previously, this is due to the form of a product of exponentials model which does not constitute a good estimate of the actual spectrum.

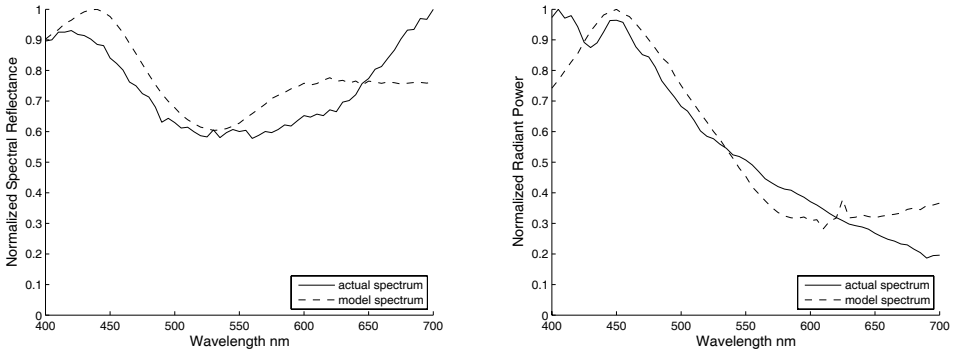

(a) Munsell Patch 1009 (b) Skylight 20

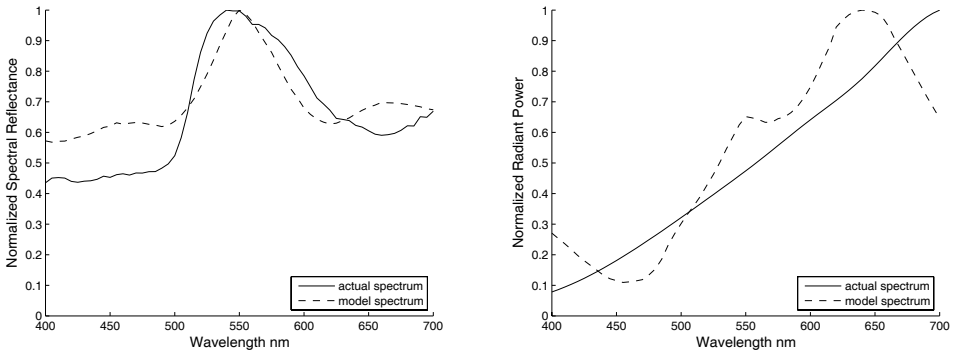

(c) Munsell Patch 455 (d) Tungsten Light at 2900K

Figure 2. *The model and actual spectra obtained in simulation for two scenes with Munsell patches: (a) Munsell patch 1009 (RMSE = 0.0971) illuminated with (b) skylight 20 (RMSE = 0.0892), and (c) Munsell patch 455 (RMSE = 0.1053) illuminated with (d) tungsten light at temperature 2900 K (RMSE = 0.1428).*

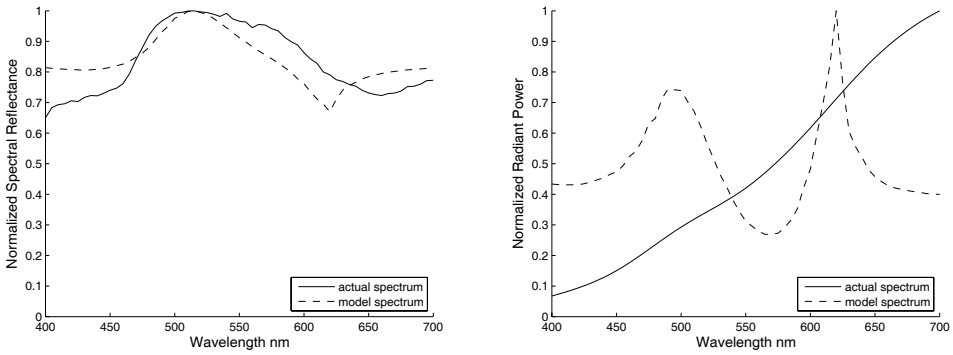

(a) Munsell Patch 510 (b) Tungsten Light at 2800 K

Figure 3. *The model and actual spectra obtained in experiment of (a) Munsell patch 510 (RMSE = 0.0740) and (b) tungsten light at 2800 K (RMSE = 0.3423).*

Construction Paper

We segment four 30x30 pixel samples from an image of construction paper patches shown in Figure 4. We average the camera responses of all the pixels in a patch to obtain one 3D response per patch. We feed these averages into our algorithm to obtain the corresponding spectral estimates. Two of the surface spectral estimates are shown in Figure 5. The actual spectra are measured using a PR-650 spectroradiometer. We can see that the approach provides good spectral estimates of construction paper. This is expected as construction paper spectra have entropies similar to those of the Munsell patches. The average entropy of the spectra of the six construction paper samples is 3.9860, while that of all the Munsell patches is 4.0402.

Discussion

We compare the results obtained in simulation and experiment using the RMS errors between the model and actual spectra as shown in Table 1. On the first row, we show the average of the RMS errors of surface patches in all scenes, followed by that of the illuminants for the simulations case. The second and third

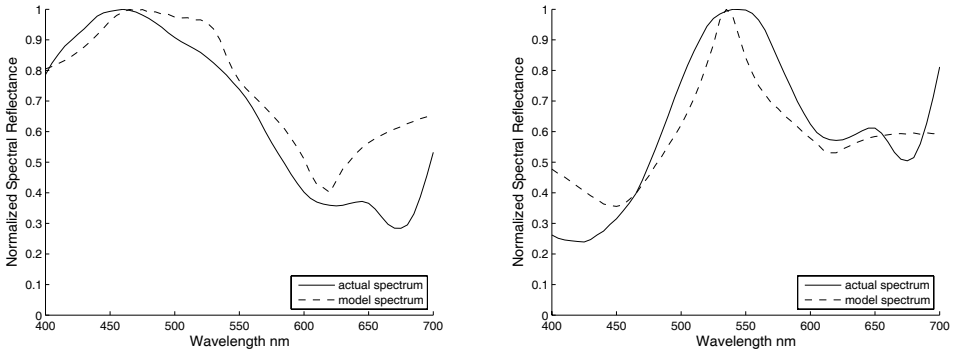

(a) Light Blue (b) Green

Figure 5. *The model and actual spectra of construction paper samples, (a) light blue (RMSE = 0.1339) and (b) green (RMSE = 0.1140).*

rows show equivalent results for the set of simulated scenes using tungsten light and for the set of real scenes respectively. This table shows that the performance of the proposed approach is similar in simulation and experiment for the surface patch spectral estimates in case of tungsten illumination. However, our approach does not do as well in estimating the illuminant spectra of tungsten light for real scenes compared to simulated scenes. Moreover, our approach does better at estimating daylight and skylight spectra than tungsten light spectra.

Table 1. Average of the RMS errors of surface patches and illuminants in all scenes for simulation and experimental data.

Type of Result	Surface Patches	Illuminants
Simulation (all lights)	0.1699	0.1762
Simulation (tungsten)	0.1492	0.2332
Experiment (tungsten)	0.1658	0.4451

Comparisons to Previous Work

As mentioned earlier, all approaches employing spectral models use linear combinations of spectral basis functions to represent surface and illuminant spectra. The objective in these approaches is to compute the weights of these basis functions. Then the surface and illuminant spectra can be recovered.

One of the most popular of these approaches is Maloney and Wandell's [3] which suffers from many limitations. First, in order for a solution to exist, the number of sensor classes has to be greater than the number of basis functions used to model surface spectra. Moreover, in order for the solution to be unique, the number of surfaces in a scene has to be greater than the number of basis functions used to model the illuminant spectrum. D'Zmura and Iverson extended Maloney and Wandell's approach [3] to propose the general linear recovery algorithm [11, 12]. Necessary and sufficient conditions are determined for recovering weights of basis functions of surface and illuminant spectra. Moreover, a model check algorithm is proposed for unique recovery of these weights. In [13], D'Zmura and Iverson extended their approach to use multiple views of the same scene. All these approaches suffer from limitations. For example, in the approach proposed in [13], the number of views has to be less than the number of basis functions used to represent the illuminant. Moreover, the number of basis functions used to represent each of the illuminant and surface spectra has to be greater than the number of photoreceptor types.

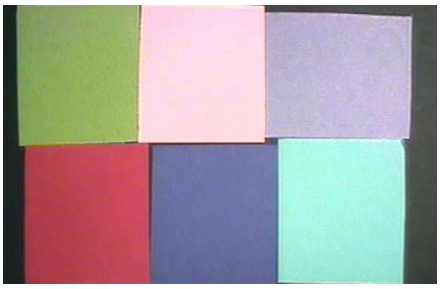

Figure 4. *An image of a few construction paper patches.*

The limitations of these approaches motivated Brainard and Freeman to introduce the Bayesian algorithm for color constancy [4]. This algorithm builds on Maloney and Wandell's approach as it represents each of the surface and illuminant spectra by a linear model. Moreover, it imposes no restrictions on the number of surface and illuminant basis functions, nor on the number of photoreceptor responses. This leads us to choose to compare our proposed maximum entropy approach to the Bayesian approach, which we describe briefly below.

The Bayesian Approach

The Bayesian approach regularizes the problem of computing the weights for the basis functions of the surface and illuminant spectra [4]. The posterior distribution function is computed over the set of surface and illuminant spectral weights given the photoreceptor responses. The solution weights vector is the one that maximizes the posterior if the maximum *a posteriori* (MAP) rule is used to minimize the Bayesian expected loss function.

If we denote the vectors of surface and illuminant spectral model weights by a and b respectively, and the sensor responses by RGB (Red, Green, and Blue), we can write the posterior function $p(a,b|RGB)$ as follows:

$$p(a,b|RGB) = \frac{p(RGB|a,b)\ p(a,b)}{p(RGB)}$$

$$\propto \quad p(RGB|a,b)\ p(a,b)$$
$$= \quad p(RGB|a,b)\ p(a)p(b). \qquad (13)$$

$p(RGB|a,b)$ is the likelihood function which models the relationship between the weights and the sensor responses. $p(RGB)$ represents the probability of the sensor responses and is a normalization term that does not affect the shape of the posterior distribution. $p(a,b)$ represents the prior information on the model weights. The surface and illuminant spectra can be assumed to be independent; therefore, the prior distributions over the corresponding basis function weights are statistically independent ($p(a,b) = p(a)p(b)$). The solution for the surface and illuminant spectra basis function weights is obtained by maximizing the posterior function.

Simulation Results

The simulated sensor responses are computed in the same way explained previously. In this case, however, noise is added to these responses to simulate the real world settings. Possible sources of noise in the real world could be the flickering of the light source while taking the images, dust on the camera lens, and electronic noise in the camera. The added noise is modeled by a normal distribution having a standard deviation of 5% of the entire range of simulated responses for each of the long-, medium-, and short- wavelength range channels. The likelihood is computed using the model predictions of the sensor responses. The surface spectra are represented by eight basis functions while the illuminant spectra are represented by five basis functions. We perform PCA on the set of Munsell patch spectra to obtain the basis functions. The prior distributions are assumed to be Gaussian with means and variances computed from the given databases of surface and illuminant spectra.

We simulate the approach for 100 scenes where the surface patches and illuminants are chosen at random. We assume that there are four surface patches in the scene here as well. For each scene, we compute the average of the RMS errors of all the surface patches and the illuminants as shown in Table 2. The table shows that the performance of our proposed approach is similar to that of the Bayesian approach in estimating surface spectra even though the former requires no *a priori* information. On the other hand, the use of basis functions allows the Bayesian approach to better model the illuminant spectra.

Table 2. Average of the RMS errors of surface patches and illuminants in all scenes for the Maximum Entropy approach and the Bayesian approach. Simulations were carried out in the presence of noise.

Scene Component	Maximum Entropy	Bayesian
Surface Patches	0.3220	0.2827
Illuminants	0.3026	0.2371

Concluding Remarks

Our simulation and experimental results indicate that our maximum entropy approach is successful in estimating surface spectra. Obtaining a spectral model for a surface solves the color constancy problem because this spectrum is illumination invariant. The performance of the proposed approach in estimating surface spectra is similar to that of the best spectral based color constancy approaches which have been reported. However, the proposed approach provides a major advantage in that it requires no *a priori* information contrary to other spectral based approaches.

References

[1] J. J. Clark and S. Skaff, Maximum Entropy Models of Surface Reflectance Spectra, IEEE Instrumentation and Measurement Technology Conference, Ottawa, Canada, pg. 1557 (2005).

[2] E. T. Jaynes, Prior Probabilities, IEEE Transactions on Systems Science and Cybernetics, 4, pg. 227 (1968).

[3] L. T. Maloney and B. A. Wandell, Color Constancy: A Method for Recovering Surface Spectral Reflectance, Journal of the Optical Society of America A, 3, 1, pg. 29 (1986).

[4] D. H. Brainard and W. T. Freeman, Bayesian Color Constancy, Journal of the Optical Society of America A, 14, 7, pg. 1393 (1997).

[5] S. Skaff, T. Arbel and J. J. Clark, Active Bayesian Color Constancy with Non-Uniform Sensors, IEEE International Conference on Pattern Recognition, Quebec City, Canada, pg. 681 (2002).

[6] J. P. S. Parkkinen, J. Hallikainen and T. Jaaskelainen, Characteristic Spectra of Munsell Colors, Journal of the Optical Society of America A, 6, 2, pg. 318 (1989).

[7] J. P. S. Parkkinen and P. Silfsten, Spectra Databases, University of Joensuu, Joensuu, Finland, http://www.it.lut.fi/ip/research/color/database/database.html.

[8] IES lighting handbook (Illuminating Engineering Society of North America, New York, NY, 1981).

[9] Munsell Book of Color-Matte Finish Collection (Munsell Color, Baltimore, MD, 1976).

[10] T. M. Cover and J. A. Thomas, Elements of Information Theory (Wiley & Sons, 1991).

[11] M. D. D'Zmura and G. Iverson, Color Constancy. I. Basic Theory of Two-Stage Linear Recovery of Spectral Descriptions for Lights and Surfaces, Journal of the Optical Society of America A, 10, 10, pg. 2148 (1993).

[12] M. D. D'Zmura and G. Iverson, Color Constancy. II. Results for Two-Stage Linear Recovery of Spectral Descriptions for Lights and Surfaces, Journal of the Optical Society of America A, 10, 10, pg. 2166 (1993).

[13] M. D. D'Zmura and G. Iverson, Color Constancy. III. General Linear Recovery of Spectral Descriptions for Lights and Surfaces, Journal of the Optical Society of America A, 11, 9, pg. 2389 (1994).

[14] G. D. Finlayson, Color in perspective, IEEE Transactions on Pattern Analysis and Machine Intelligence, 18, 10, pg. 1034 (1996).

[15] G. D. Finlayson, S. D. Hordley and P. M. Hubel, Color by Correlation: A Simple, Unifying Framework for Color Constancy, IEEE Transactions on Pattern Analysis and Machine Intelligence, 23, 11, pg. 1209 (2001).

[16] K. Barnard, G. Finlayson and B. Funt, Color Constancy for Scenes with Varying Illumination, Computer Vision and Image Understanding, 65, 2, pg. 311 (1997).

Author Biography

Sandra Skaff received the BS degree in computer science from the American University of Beirut (Beirut, Lebanon) in 2000. She received the MEng degree in electrical engineering from McGill University (Montreal, Canada) in 2002. She has been at McGill's Visual Motor Systems Laboratory since 2000 and is currently pursuing a PhD. Her research interests include spectral models for color perception and color constancy.