

Maximum Entropy Spectral Models for Color Constancy in the Presence of Interreflections

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

Abstract

Interreflections in a scene can be exploited to improve upon surface and illuminant spectral estimation. In this paper, we present a novel maximum entropy approach to spectral color constancy in the presence of interreflections. Previous approaches employ linear model representations of surface and illuminant spectra. Such representations are not always practical as a database of spectra has to be specified in advance in the corresponding algorithms. The proposed approach has a major advantage over previous algorithms in that it requires only camera sensor responses from the mutual illumination or edge region of a folded surface and from the far from the edge regions to estimate surface and illuminant spectra. We demonstrate the feasibility of the approach while assuming a one-bounce, two-zone model of mutual illumination. We test our approach in both simulation and experiment. In the case of one surface patch in a scene, when the color constancy problem has no solution, we are able to obtain promising results.

Introduction

It has been shown that the human visual system makes use of mutual illumination information in color perception [1, 2]. However, there has been little work in the computer vision literature which exploits mutual illumination to improve upon color perception. In this work, we address the problem of spectral color constancy while making use of mutual illumination information. The problem of spectral color constancy lies in computing a surface reflectance spectrum that is independent of the spectrum of light incident on the surface.

Most color constancy approaches with interreflections use spectral models for the surfaces and the illuminant, while employing finite-dimensional linear model representations for these spectra. In [3], Funt *et al.* assumed a one-bounce, two-zone model of interreflections. The one-bounce model takes into account the light reflected off one surface that bounces onto the other. The two zones comprise that in the mutual illumination region, also called the edge, and that far from this region. By taking mutual illumination into account, they effectively added a sensor class to Maloney and Wandell's approach [4]. This is important because one more basis function can then be used to model the surface spectrum, given the restrictions on the number of basis functions in Maloney and Wandell's approach. Drew and Funt [5] extended Funt *et al.*'s [3] approach to account for multiple zones, using the radiosity method from [6]. They also proposed the variational approach [7] that does not assume diffuse-illumination, a two-patches limit, or locations with negligible interreflection as in the previous approaches. It still assumes a one-bounce model of inter-reflection, however. Harder [8] investigated interreflections while assuming known illumination in addition to Lambertian surfaces.

All the mentioned approaches use three basis functions for each of the surface and illuminant spectra. The only exception is in Harder's work [8] where the illuminant spectrum is assumed to be known, which is not the case in this paper. Such a small number of basis functions may not be enough to provide an accurate representation of surface spectra. Moreover, all these approaches require the database of surface and illuminant spectra to be specified in advance in order to obtain the basis functions for the linear model representations. Such a requirement places these approaches at a major disadvantage as databases might not be available in advance. Even if they are available, they might not be consistent with the data in a certain application.

We propose a novel approach that estimates surface and illuminant spectra in the presence of interreflections given only camera sensor responses. This approach is based on the color constancy technique introduced in [9] in which the surface and illuminant spectra are represented by maximum entropy models, and therefore do not require a set of basis functions to be specified in advance. Maximum entropy models were successfully used to estimate Munsell patch reflectance spectra given only photoreceptor responses in [10]. 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 [11]. In the case of physical processes representing spectra, many surfaces observed in our everyday-life surroundings have spectra of high entropy, as opposed to monochromatic surfaces which have low entropy spectra [10]. The illuminant spectra are also represented by maximum entropy models as they are observed in our everyday-life surroundings and therefore can be produced in a large number of ways [9].

In this paper, the proposed approach is explained and derived. To this end, diffuse-illumination and a Mondrian scene composed of matte, Lambertian surfaces are assumed. The light illuminating a Mondrian scene is assumed to be locally constant. This means that the spectral characteristics of the light vary slowly. As for the interreflections, a one-bounce, two-zone model is assumed. Next, the performance of the approach is analyzed in simulation and then in experiment. In particular, in the case of one surface patch in the scene, it is shown how exploiting interreflection information improves upon spectral estimation. Moreover, while, to the knowledge of the authors, none of the previous approaches provided surface or illuminant spectral estimates for real images, such spectra are shown in this paper.

Maximum Entropy Spectral Based Color Constancy with Interreflections

The proposed color constancy approach aims at recovering surface and illuminant spectra in the presence of interreflections given only camera sensor responses. We assume a one-bounce,




Figure 1. A scene with one folded surface patch of reflectance spectrum $s(\lambda)$, illuminated by one light source of spectrum $e(\lambda)$. $I_m(\lambda)$ is the light reflected off the surface in the edge region while $I(\lambda)$ is the light reflected off the surface far from this region.

two-zone model in this paper. The main goal of exploiting interreflection information is to show that in the one surface patch case, when the color constancy problem has no solution, reasonable surface and illuminant spectral estimates can be obtained. Figure 1 shows a folded surface patch in a scene illuminated by a single light source. $s(\lambda)$ is the surface spectrum while $e(\lambda)$ is the illuminant spectrum. Under the assumption of uniform illumination and Lambertian surfaces, the light reflected off the surface in the first zone, the region away from the edge, is:

$$I(\lambda) = s(\lambda)e(\lambda). \quad (1)$$

In the second zone, the mutual illumination region, the light reflected off the surface is:

$$\begin{aligned} I_m(\lambda) &= I(\lambda) + \alpha I(\lambda)s(\lambda) \\ &= s(\lambda)e(\lambda) + \alpha s^2(\lambda)e(\lambda), \end{aligned} \quad (2)$$

where α denotes the proportion of $I(\lambda)$ reflected off one side that bounces onto the edge region of the other side.

Two types of light fall on the camera sensors in this case, $I(\lambda)$ and $I_m(\lambda)$, which after projecting onto the sensor spectral sensitivities $R_k(\lambda)$, yield P_k and P_{km} respectively:

$$P_k = \sum_{\lambda=1}^M R_k(\lambda)s(\lambda)e(\lambda), \quad k = 1, 2, \dots, p; \quad (3a)$$

$$P_{km} = P_k + \alpha \sum_{\lambda=1}^M R_k(\lambda)s^2(\lambda)e(\lambda), \quad k = 1, 2, \dots, p, \quad (3b)$$

where p is the number of sensor classes, each denoted by k ; λ is the wavelength taken over the visible range; 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 surface and illuminant spectra are represented by maximum entropy models. Each of these spectra needs to be represented by a probability density function (pdf) in order to compute

its entropy. The light incident on a surface is a collection of photons, and each photon has a specific wavelength. The pdf representation of the illuminant is therefore the probability density of the wavelength of a photon incident on the surface. Assuming that there is always an incident photon, $e(\lambda)$ can be related to the corresponding pdf $p_e(\lambda)$, which can be obtained by:

$$\begin{aligned} p_e(\lambda) &= \frac{e(\lambda)}{\sum_{\lambda=1}^M e(\lambda)}, \\ &\equiv \frac{e(\lambda)}{I_e}, \quad \lambda = 1, \dots, M, \end{aligned} \quad (4)$$

where $I_e = \sum_{\lambda=1}^M e(\lambda)$ is the intensity of the light emitted by the source. An incident photon can be either absorbed by or reflected off the surface it hits. Denoting the event that a photon is absorbed by the surface by A and the event that a photon is reflected off the surface by R , we can write: $P(A) + P(R) = 1$. The probability of a wavelength given that event R occurred, which means that a photon has been reflected off the surface, is denoted by $p(\lambda|R)$. Bayes Rule gives:

$$p(\lambda|R) = \frac{p(R|\lambda) p(\lambda)}{p(R)}, \quad (5)$$

where $p(\lambda)$ is equivalent to $p_e(\lambda)$. In Equation 5, $p(R|\lambda)$ denotes the likelihood of wavelength λ being observed given a reflected photon. This likelihood is normalized to give the pdf $p_s(\lambda)$:

$$p_s(\lambda) = \frac{p(R|\lambda)}{\sum_{\lambda=1}^M p(R|\lambda)}, \quad \lambda = 1, \dots, M. \quad (6)$$

$p_s(\lambda)$ can serve as the pdf representation of $s(\lambda)$, which is related to the former by the multiplicative factor $\sum_{\lambda=1}^M p(R|\lambda)/p(R)$. This factor can be thought of as the surface albedo relative to the incident light.

Throughout this paper we refer to the entropy of the pdf representation of a spectrum as the entropy of a spectrum for simplicity. For the purposes of the following derivations, we denote the wavelength corresponding to the surface spectrum by λ_s and the wavelength corresponding to the illuminant spectrum by λ_e . Our goal is to estimate the surface and illuminant spectrum. Therefore we seek the joint probability density $p_{s,e}(\lambda_s, \lambda_e)$ that maximizes the entropy H [12] given by:

$$H = - \sum_{\lambda_s=1}^M \sum_{\lambda_e=1}^M p_{s,e}(\lambda_s, \lambda_e) \log p_{s,e}(\lambda_s, \lambda_e). \quad (7)$$

The surface and illuminant spectra are assumed to be independent, and therefore the corresponding probability densities $p_s(\lambda_s)$ and $p_e(\lambda_e)$ 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 [12] as:

$$H = - \sum_{\lambda_s=1}^M p_s(\lambda_s) \log p_s(\lambda_s) - \sum_{\lambda_e=1}^M p_e(\lambda_e) \log p_e(\lambda_e), \quad (8)$$

where $p_{s,e}(\lambda_s, \lambda_e) = p_s(\lambda_s)p_e(\lambda_e)$. Therefore, we can now find $p_s(\lambda_s)$ and $p_e(\lambda_e)$ that maximize H , given the constraints in Equations 3(a) and 3(b). The function `fmincon` from the Matlab Optimization Toolbox is used for this purpose. The α of Equation 2 is also estimated by the optimization as it is not known in real images.

Simulation Results

We test the performance of our approach on two types of patches: *matte* Munsell [13] and construction paper. The former are representative of a wide range of hues encountered in printing. The Munsell spectra were measured by Parkkinen *et al.* [14]. The construction paper spectra were measured with a PR-650 spectroradiometer in our laboratory. The illuminant spectra used are those of daylight and skylight, measured by Parkkinen and Silftsen [15], and a set of tungsten light spectra. This set is composed of ten spectra with temperatures ranging from 2600K to 3500K, in steps of 100K, obtained from the IES lighting handbook [16].

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 3(a) with $p = 3$). To obtain the simulated sensor responses for the patches in the mutual illumination region, the light incident on this region is obtained as expressed in Equation 2. The set of α values considered is 0.1, 0.2, 0.3, as it is assumed that at most 30% of the light reflected off one surface bounces onto the other. This light is then multiplied by the sensor spectral sensitivity curves of the camera (Equation 3(b) where $p = 3$). The sensitivity curves of the camera are obtained from the manufacturer. The wavelength range considered for these spectra is 400 nm to 700 nm, and is discretized into 10 nm bins. This yields a dimension of $M = 31$ in Equations 3(a) and 3(b) for the surface and illuminant spectra.

Two types of artificial scenes, Munsell and construction paper, were constructed as explained below. The model surface and illuminant spectra are computed for a scene for each of the three α 's (0.1, 0.2, 0.3) given the proposed algorithm. These spectra are also computed given the algorithm proposed in [9], where mutual illumination information is ignored, for comparison purposes. To evaluate the performance of our approach, the root mean square (RMS) errors between the normalized actual and model spectra for the surface patches and illuminant in each scene is computed. Each spectrum is normalized to a maximum of one as we do not intend to recover intensity information. This is a common practice in solving color constancy problems. Finally, the average of the RMS errors over all scenes is taken for each of the α values. The performance of the approach is compared for different α 's to the performance of the approach in [9].

The average RMS errors for each of the different α 's for both the surface and the illuminant spectra is depicted in Figures 2 and 3 for each type of scene. The $\alpha = 0$ case corresponds to the approach in [9]. The term RMS error is denoted by RMSE.

Munsell Patches

Eighty scenes of one Munsell patch and one illuminant in each were constructed. The 80 patches were chosen at random from the set of 1269 *matte* Munsell patches [13]. The illuminant for each scene was selected at random from the set of daylight, skylight, and tungsten light spectra.

We plot the actual and model spectra for the surface patch and the illuminant in each of two scenes, one illuminated by daylight and one by tungsten light for $\alpha = 0.2$ in Figure 4. We also plot these spectra when interreflections are ignored. We choose this value for α as it is in the middle of the range of α values considered in this study.

Construction Paper Patches

Eight scenes of one construction paper patch of different hues and one illuminant in each are considered. The illuminant is tungsten light at 2800K in all scenes as this allows for comparison to the experimental results.

We plot the actual and model spectra for the surface patch and the illuminant for a scene with a light blue construction paper patch for $\alpha = 0.2$ in Figure 5.

Discussion

From the bar plots (Figures 2(b) and 3 (b)), we can see that the spectral estimates for the illuminant improve when interreflection information is exploited, for all α 's. For the Munsell case, the average RMS error decreases from 0.2313 ($\alpha = 0$) to 0.1453 ($\alpha = 0.2$). For the construction paper case, the average RMS error decreases from 0.2673 ($\alpha = 0$) to 0.2116 ($\alpha = 0.1$). As for the surface spectra (Figures 2(a) and 3 (a)), the average RMS error decreases from 0.2765 ($\alpha = 0$) to 0.2584 ($\alpha = 0.1$). On average, this is not the case for the Munsell surface spectra in terms of the RMS errors. However, in quite a few cases when the RMS errors are higher, our approach renders a spectral estimate that has a more similar shape to that of the actual one than in the case when interreflections are ignored.

From the spectra shown (Figures 4 and 5), the maximum entropy approach of [9] cannot distinguish between the surface and the illuminant spectra when there is only one surface patch in the scene. By incorporating the sensor responses obtained from the mutual illumination region, better spectral estimates are obtained. For example, the error in the CIE_{xy} chromaticity coordinates between the model and actual spectrum shown in Figure 5(a) is 0.0388 when $\alpha = 0.2$ compared to 0.1263 when $\alpha = 0$, even though the former spectrum seems noisier. Moreover, since the model spectra comprise a product of three exponentials, one in each wavelength range (long, medium, and short), the daylight spectra are better estimated by this approach than the tungsten light spectra.

Experimental Results

No previous spectral color constancy approach with interreflections has reported surface and illuminant spectra obtained from real images. The purpose of the experiments here is to show that we can still obtain reasonable spectral estimates for surfaces and illuminants in real scenes despite the assumptions imposed. We consider scenes made of folded strips of construction paper as explained below. We do not use the Munsell patches in this case as they are 1.6 x 2.3 cm and folding them would not allow for a sensor response far enough from the mutual illumination region. In addition, the Munsell patches have spectra similar to those of the construction paper and therefore similar algorithmic performance would be expected.

We consider eight scenes. Each strip of construction paper was 4 x 21 cm. A strip was folded in the middle to make an opening angle of approximately 120°. Each side of the paper was about 30° to the image plane. The camera was mounted directly above. The light source, tungsten at 2800 K, was about 1 m away from the scene at a 45° angle. We captured images of these scenes with a Panasonic WV-CP410 camera. We segmented two 20x20 pixel samples from each image: one in the mutual illumination region and one far from this region. Since we assume Mondrian

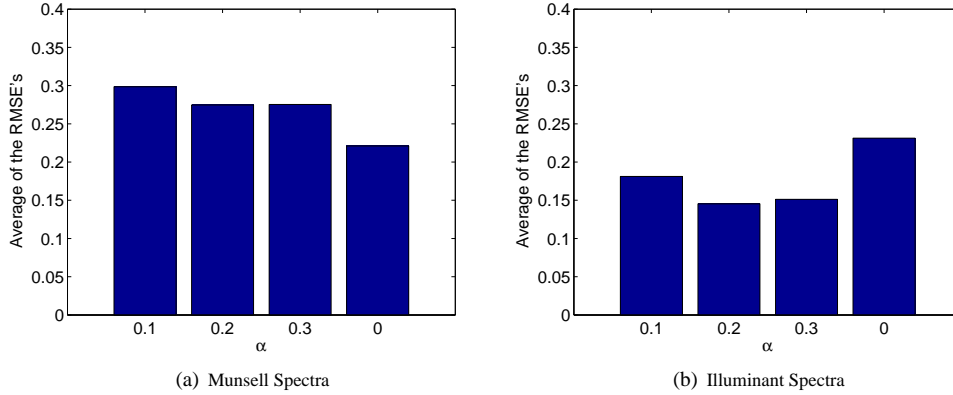


Figure 2. The average of the RMS errors over all synthetic scenes when interreflection information is exploited ($\alpha = 0.1, 0.2, 0.3$) and ignored ($\alpha = 0$), for the (a) Munsell surface patch and (b) illuminant spectra.

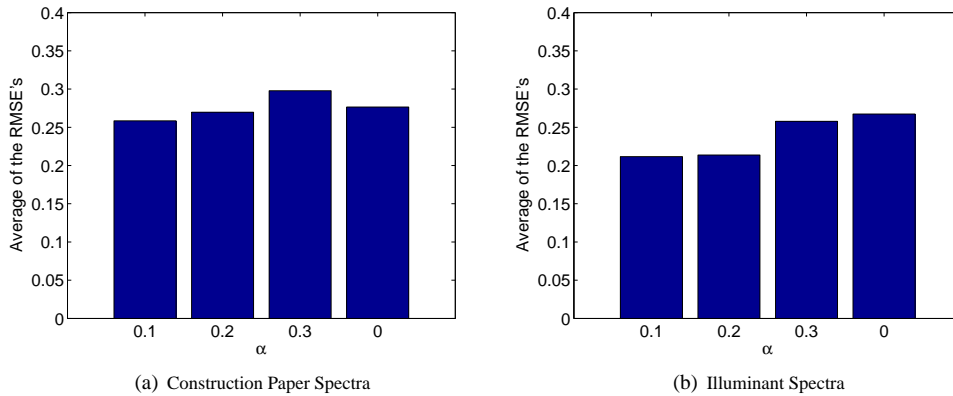


Figure 3. The average of the RMS errors over all synthetic scenes for the case when interreflection information is exploited ($\alpha = 0.1, 0.2, 0.3$) and ignored ($\alpha = 0$), for the (a) construction paper patch and (b) illuminant spectra.

scenes where the illumination is locally constant, we averaged 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 α and spectral estimates.

We compute the average RMS errors for the surface and illuminant spectra over the scenes for which the optimization converged. These were five out of eight. The average RMS errors are slightly higher in the interreflection case than that without interreflection. For the surface spectra, the average error is 0.3006 for $\alpha > 0$ and 0.2597 for $\alpha = 0$. For the illuminant spectra, the average error is 0.4505 for $\alpha > 0$ and 0.4200 for $\alpha = 0$.

We plot the actual and model spectra for the surface and the illuminant for a scene with a pink patch in Figure 6. We also plot these spectra when interreflections are ignored. These plots suggest that surface spectral estimation can be improved with interreflection information in some cases. The same cannot be concluded for the illuminant spectral estimates. However, once a reasonable surface spectrum is obtained, the color constancy problem is solved as this spectrum is illumination invariant. This is very interesting as it suggests that the approach has potential for being used with real images despite the assumptions imposed.

Concluding Remarks

A new maximum entropy approach to solve for surface and illuminant spectra in the presence of interreflections was introduced. The approach provides a major advantage in that it requires no *a priori* information contrary to other spectral based color constancy approaches with interreflections. Our simulation results indicate that exploiting interreflection information generally improves upon spectral estimates. Preliminary experiments did not confirm this hypothesis; however the algorithm performed sufficiently well to suggest that interreflections could be useful with a better control of the experimental setup.

References

- [1] M. G. Bloj, D. Kersten and A. C. Hurlbert, Perception of three-dimensional shape influences colour perception through mutual illumination, *Nature*, 402, pg. 877 (1999).
- [2] M. S. Langer, A model of how interreflections can affect color appearance, *Colour Research and Application*, 26, pg. 218 (2001).
- [3] B. V. Funt, M. S. Drew and J. Ho, Color constancy from mutual reflection, *International Journal of Computer Vision*, 6, pg. 5 (1991).
- [4] 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).
- [5] M. S. Drew and B. V. Funt, Calculating surface reflectance using a

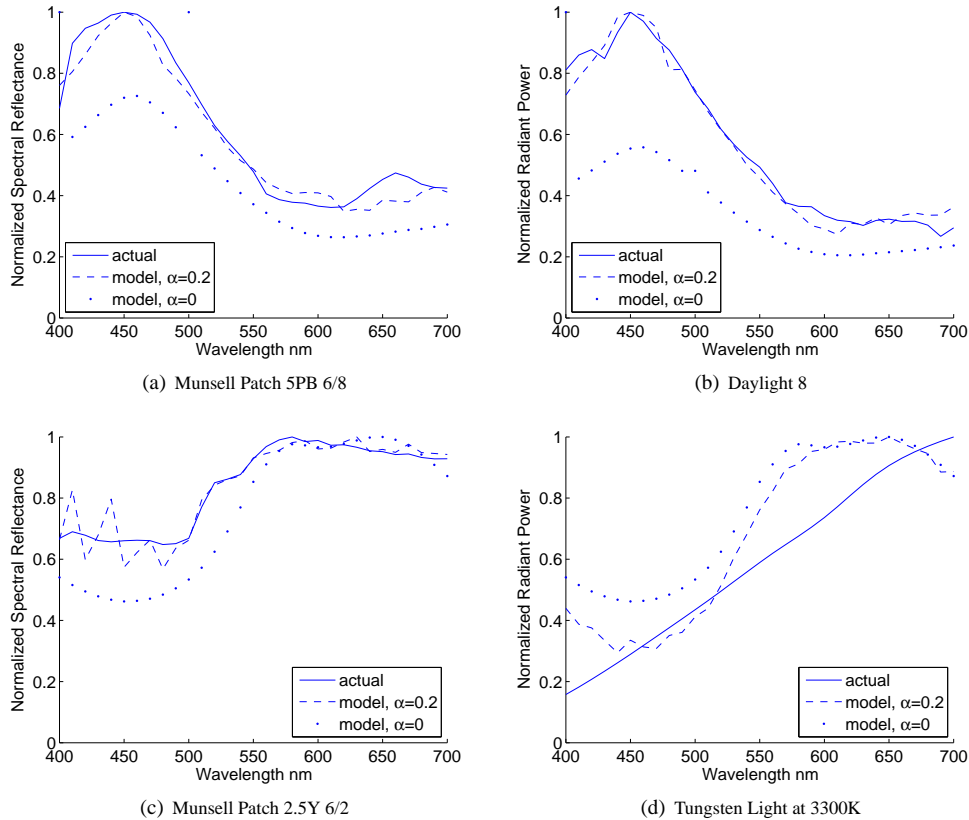


Figure 4. The actual and model spectra, when interreflection information is exploited ($\alpha = 0.2$) and ignored ($\alpha = 0$), obtained in simulation for two scenes with one Munsell patch each: (a) Munsell patch 5PB 6/8 ($RMSE = 0.0486$, $RMSE = 0.1952$) illuminated with (b) daylight 8 ($RMSE = 0.0389$, $RMSE = 0.2475$), and (c) Munsell patch 2.5Y 6/2 ($RMSE = 0.0958$, $RMSE = 0.1813$) illuminated with (d) tungsten light at temperature 3300 K ($RMSE = 0.1680$, $RMSE = 0.2434$).

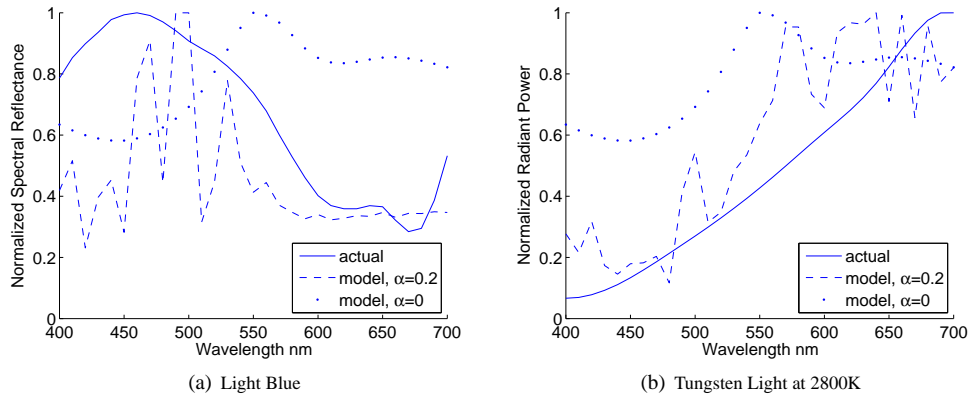


Figure 5. The actual and model spectra, when interreflection information is exploited ($\alpha = 0.2$) and ignored ($\alpha = 0$), obtained in simulation for a scene with a (a) light blue construction paper patch ($RMSE = 0.3106$, $RMSE = 0.3770$) illuminated with (b) tungsten light ($RMSE = 0.2023$, $RMSE = 0.3915$).

single-bounce model of mutual reflection, IEEE International Conference on Computer Vision, Osaka, Japan, pg. 394 (1990).

- [6] C. M. Goral, K. E. Torrance, D. P. Greenberg and B. Battaile, Modeling the interaction of light between diffuse surfaces, *Computer Graphics*, 18, 3, pg. 213 (1984).
- [7] M. S. Drew and B. V. Funt, Variational approach to interreflection in color images, *Journal of the Optical Society of America A*, 9, 8, pg. 1255 (1992).

- [8] I. A. Harder, Spectral analysis of interreflection, Masters Thesis, School of Computer Science, Simon Fraser University, Burnaby, BC (1992).
- [9] S. Skaff and J. J. Clark, Maximum entropy spectral models for color constancy, *Proceedings of the 15th IS&T/SID's Color Imaging Conference*, Albuquerque, NM, pg. 100 (2007).
- [10] J. J. Clark and S. Skaff, Maximum entropy models of surface reflectance spectra, *IEEE Instrumentation and Measurement Technol-*

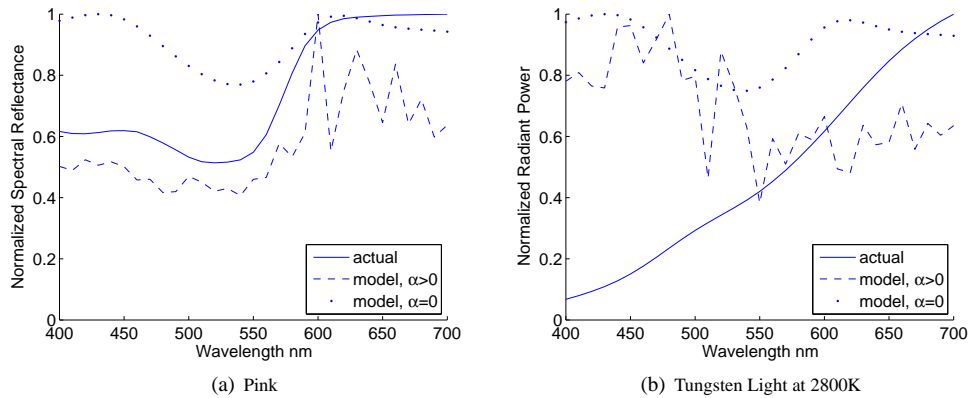


Figure 6. The actual and model spectra, when interreflection information is exploited ($\alpha > 0$) and ignored ($\alpha = 0$), obtained in experiment for a scene with a (a) pink construction paper patch (RMSE = 0.2075, RMSE = 0.2395) illuminated with (b) tungsten light at 2800K (RMSE = 0.4712, RMSE = 0.5307).

ogy Conference, Ottawa, ON, pg. 1557 (2005).

- [11] E. T. Jaynes, Prior probabilities, IEEE Transactions on Systems Science and Cybernetics, 4, pg. 227 (1968).
- [12] T. M. Cover and J. A. Thomas, Elements of information theory (Wiley & Sons, 1991).
- [13] Munsell Book of Color-Matte Finish Collection (Munsell Color, Baltimore, MD, 1976).
- [14] 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).
- [15] 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>.
- [16] IES lighting handbook (Illuminating Engineering Society of North America, New York, NY, 1981).

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.