IMTC 2005 - Instrumentation and Measurement
Technology Conference
Ottawa, Ontario, canada, May 17–19, 2005

Maximum Entropy Models of Surface Reflectance Spectra

J. Clark and S. Skaff

Centre for Intelligent Machines,
McGill University,
3480 University Street, Montreal, Canada
Phone: 1-514-398-2654, Fax: 1-514-398-7348, E–mail: clark@cim.mcgill.ca

Abstract – *We investigate the use of a maximum entropy spectral model for surface spectra from photoreceptor measurements. We compare the accuracy of maximum entropy models againts those provided by the stand linear PCA-based models. We show that the maximum entropy models have similar performance to the linear models, even though they require no* a priori *information.*

Keywords – *color, surface reflectance, spectra, maximum entropy, colorblindness*

I. INTRODUCTION

The measurement of color and the understanding of color perception in humans are areas of vigorous research activities. Color is a subjective quantity, produced by the brain, but is clearly related to physical quantities. Thus there has arisen a dichotomy in the approaches used to investigate color. One approach derives from the apparent 3-dimensionality of color, and works with 3-dimensional representations (e.g. different color spaces such as RGB, HSV, CIE, etc). This approach has the drawback that there is no universal intrinsic color space. Different applications imply different color spaces. For this reason, researchers often concentrate on spectral representations of color, where the high dimensional (potentially infinite dimensional) spectrum of a light-reflecting surface is taken to indicate the color of the surface. One of the primary advantages of a spectral representation is that it provides a common structure that facilitates fusion of information from disparate sources in estimating surface color, as in color constancy algorithms [1,3,6]. The main drawbacks of spectral representations are their high dimensionality and the fact that there are an infinite number of spectra that correspond to a given color (i.e. *metamers*). We can mitigate the difficulties imposed by the high dimensionality of surface spectra by assuming a low dimensional *model* of the spectra, for example, a 3-dimensional model.

A number of different low dimensional spectral models have been used in the color vision and measurement communities, mainly in the context of color constancy algorithms. One of the most popular is the linear model of Maloney and Wandell [3] who proposed representing spectra with a linearly weighted sum of spectral basis functions. The modeled spectra would then be obtained by computing the basis function weights which give the model spectrum that, when projected onto the photoreceptor sensitivity functions, most closely matches the actual photoreceptor measurements. One of the difficulties associated with this approach is that the basis functions must be specified somehow. One common approach is to take the first few principal components of a database of previously measured spectra. Such a database is not always available, however. In particular, it is difficult to imagine how the human brain could incorporate such basis functions, as it is not equipped with a spectrometer which it could use to create a database of spectra. Even if a database is available, it might not be consistent with, or representative of, the data present in a given application. It is also necessary that the spectral sensitivity curves of the photoreceptors be known. It is more reasonable, however, to expect that these can be learned or computed somehow through observation of the photoreceptor response to a wide range of stimuli.

II. PROPOSED APPROACH

To get around the problem of specifying the basis functions used in a linear model we propose to use a maximum entropy spectral model instead. The maximum entropy principle, as presented by Jaynes [2], holds that a physical quantity that is observed in practice will tend to a value that can be produced in the largest number of ways. In the case of the physical processes that specify surface spectra, it is apparent that low entropy spectra, such as monochromatic ones, can only be generated in a small number of ways, whereas more diffuse spectra, with high entropy, can be generated with many different combinations. Therefore, we expect that surfaces we observe in practice will typically be of high entropy. Based on this expectation we propose a low-dimensional spectral model that specifies the spectrum that has maximum entropy subject to the constraints imposed by the measurements provided by the photoreceptors (or camera pixels). The model will have dimensionality equal to the number of constraints provided by the measurements. Jaynes [2] showed that given measurements which are in the form of *expectations*, such as is the case for photoreceptors or camera sensors with spectral sensitivities $R_i(\lambda_j)$,

$P_i = \sum_{j=1}^{M} m(j)R_i(j)$, the maximum entropy solution for m is given by a product of exponentials:

$$\hat{m}(j) = \frac{1}{Z}\prod_{i=1}^{N}\exp(\alpha_i R_i(j)) = \frac{1}{Z}\exp(\sum_{i=1}^{N}\alpha_i R_i(j)) \quad (1)$$

where the scale factor Z is given by $Z = \sum_{j=1}^{M}\exp(\sum_{i=1}^{N}\alpha_i R_i(j))$. The α_i are determined so as to satisfy the constraints $P_i = \sum_{j=1}^{M} m(j)R_i(j)$.

The derivation above assumes that the norm of the spectrum is equal to 1 (as entropy is defined for probability density functions). In practice, spectra can have any value for their norm. We can extend our model to this case, by taking the scale factor $\frac{1}{Z}$ as an additional independent parameter to estimate, rather than fixing it with the expression given above.

The primary purpose of the experiments described in this paper is to compare the performance of the maximum entropy model with a linear principal components model in representing the spectra of naturally occuring surfaces. To this end we ran the models on a large database of colored surface patches with known spectra, the well-known Munsell patches, taken from the Munsell book of color [4]. This is a collection of colored paper squares with a wide range of hues, roughly covering the gamut of colors available from high quality acrylic pigments on white paper. The basis functions used in the linear model were taken as the first 4 principal components of the spectra (quantized into 55 5-nm bins from a wavelength of 430 through 700 nm) of 1240 different Munsell patches. The values of the Munsell spectra were provided by J. Parkkinen *et al* [5]. The photoreceptor sensitivities that were used were taken from data on the human retina by Stockman *et al* [7].

The model spectrum generated by the linear model given photoreceptor measurements P, will be given by:

$$\hat{m}_{lin} = \phi[R^T\phi]^{-1}P \quad (2)$$

where ϕ is an $M \times N$ element matrix whose columns are the N basis functions, with M being the number of discrete spectral samples. R is an $M \times N$ element matrix whose columns are the N photoreceptor spectral sensitivities.

The Maximum Entropy model spectrum was computed with a nonlinear constrained optimization algorithm (*fmincon* in Matlab), minimizing the negative of the entropy subject to the constraints imposed by the measurements. That is,

$$\hat{m}_{me} = \arg\min_{m}[\sum_{i=1}^{i=55} m(i)\log m(i)] \quad (3)$$

subject to

$$P = R^T\hat{m}_{me} \quad (4)$$

An additional normalization constraint of $\sum_{i=1}^{i=55}\hat{m}_{me}(i) = 1$ was added to fix the scale factor. This is equivalent to having a fourth photoreceptor which has the same spectral sensitivity at every wavelength. This could be implemented by the rods in the human visual system.

III. RESULTS

Figures 1 through 3 show the model spectra for three different Munsell patches, which are representative of the models in general. In these figures the solid curve is the measured spectrum of the Munsell patch, the dotted curve is the maximum entropy model and the curve made up of crosses is the linear model. The top plot in each figure shows the spectral models given 3 photoreceptor measurements (corresponding to the responses of the long (L), medium (M) and short (S) wavelength cones in the human retina). The bottom plot in each figure shows the spectral models for the "colorblind" case where only 2 photoreceptors values are available (the M and S receptors).

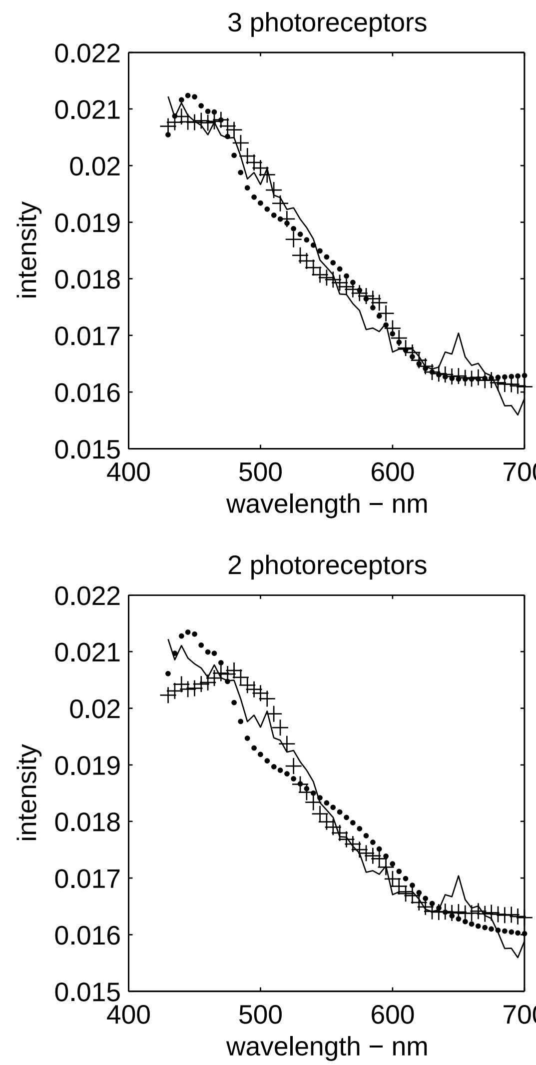

Fig. 1. The spectra predicted by the maximum entropy (dots) and linear (crosses) models as compared with the measured (by Parkkinen *et al* [5]) spectrum (solid) line for Munsell color patch number 10. Three photoreceptors were used for the top models, and two for the bottom models.

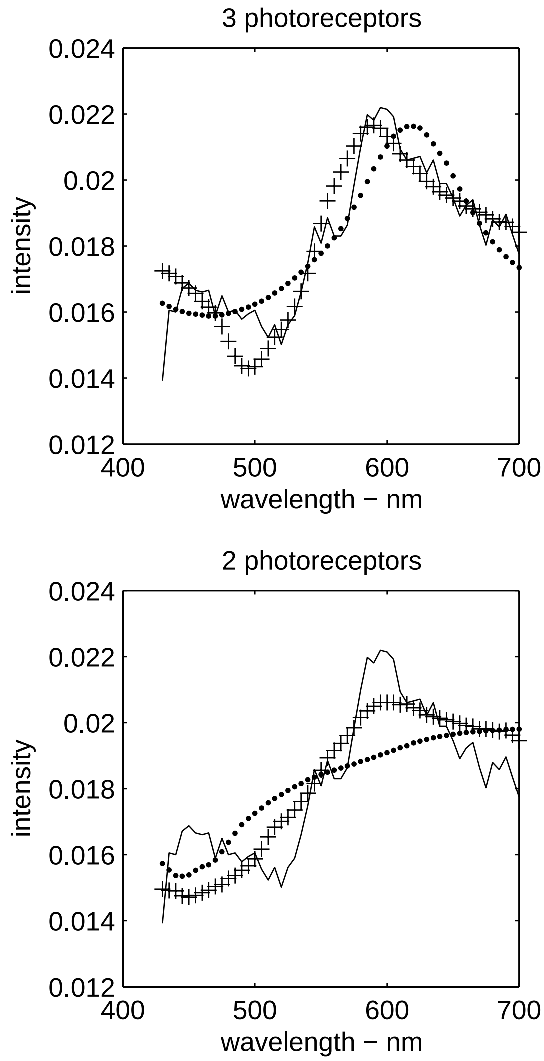

Fig. 2. The spectra predicted by the maximum entropy (dots) and linear (crosses) models as compared with the measured (by Parkkinen *et al* [5]) spectrum (solid) line for Munsell color patch number 200. Three photoreceptors were used for the top models, and two for the bottom models.

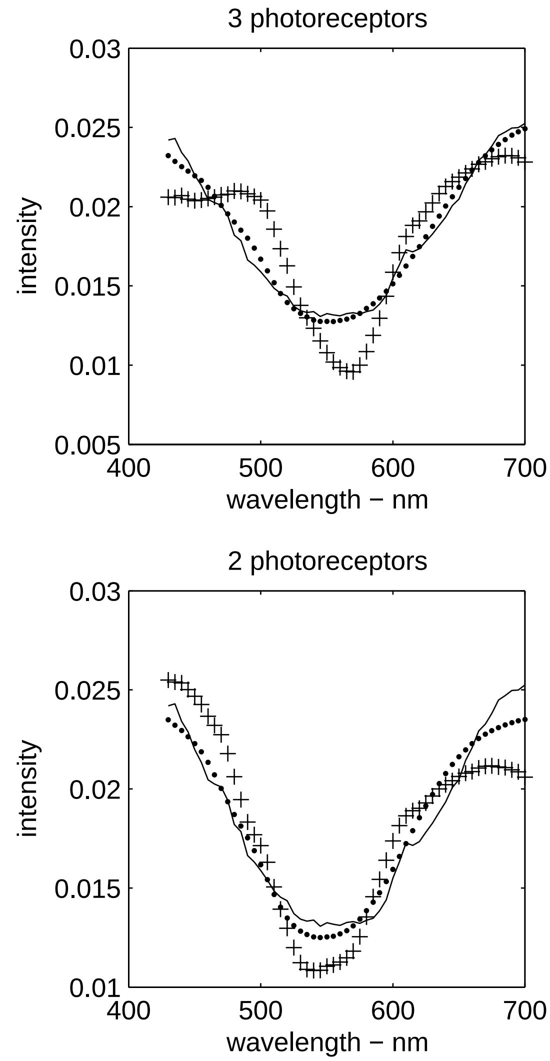

Fig. 3. The spectra predicted by the maximum entropy (dots) and linear (crosses) models as compared with the measured (by Parkkinen *et al* [5]) spectrum (solid) line for Munsell color patch number 777. Three photoreceptors were used for the top models, and two for the bottom models.

The two models give very similar results in approximating the true spectrum. The linear model incorporates much more *a priori* knowledge, however, than does the maximum entropy model, as it uses the entire set of Munsell patch spectra to derive its basis functions. In spite of this, the maximum entropy approach performs quite well. In the 3-photoreceptor case, the average RMS error over all of the Munsell patches is 0.0016 for the maximum entropy model and 0.0016 for the linear model. In the 2-photoreceptor case, the average RMS error over all of the Munsell patches is 0.0024 for the maximum entropy model and 0.0020 for the linear model.

The performance of a model is a reflection of the extent to which the assumptions inherent in the model are satisfied. The linear model assumes that all of the variation in the modeled quantity is captured by the subspace spanned by the basis functions that are used. We observed that, even though the test spectra used were taken from the database used to derive the basis functions, the linear model nonetheless sometimes performs poorly, such as on patch 777.

The maximum entropy model will do well whenever the spectrum being modeled has a high entropy (which most naturally occuring surfaces do). Reducing the number of measurements will incur a greater degree of smoothness on the prediction. Figure 4 shows the histograms of the entropy values for the two models, compared with the entropies of the measured spectra. There are two points in particular to note. The first is that the entropies of the actual physical spectra are generally quite high. Secondly, the linear model has a stronger tail towards lower entropies than the maximum entropy model. In the 3 photoreceptor case, the RMS difference between the actual Munsell patch spectrum entropy and the maximum entropy model spectrum entropy is 0.0058. The corresponding RMS difference for the linear model is 0.022. Thus, the maximum entropy model does a better job in modeling the high entropy nature of the Munsell patch spectra.

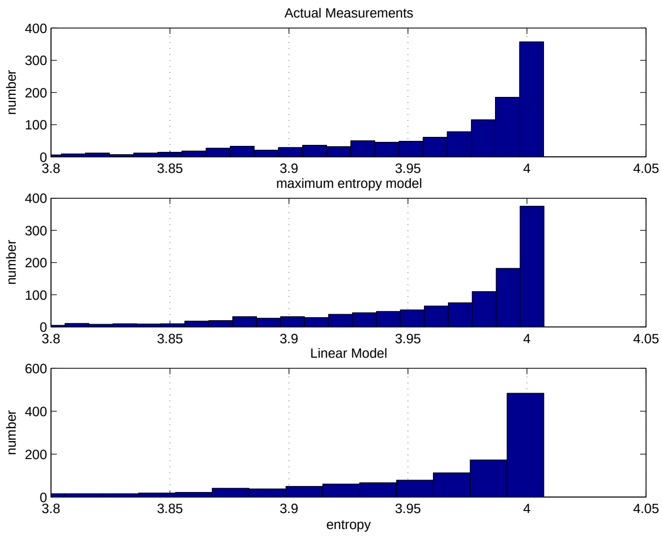

Fig. 4. Histogram of entropy values, for the actual spectral measurements of the Munsell patches (top), for the maximum entropy spectral models (middle) and the linear spectral models (bottom).

IV. CONCLUSIONS

In this paper we have presented a case for the use of maximum entropy models for surface spectra based on limited spectral measurements. The principal advantage of maximum entropy models over linear basis function models is the lack of a need to pre-determine an appropriate set of basis functions. Instead, the only assumption that is made is that the surface spectra most likely to be encountered in practice have a high entropy. Our results on modeling the spectra of the Munsell patch database shows that the maximum entropy model performs nearly as well as a linear model whose basis functions are derived directly from the Munsell database. The linear model performs poorly whenever a particular surface spectrum is not consistent with the overall character of the Munsell database, whereas the maximum entropy model degrades by over-smoothing.

V. REFERENCES

[1] Brainard, D.H. and Freeman, W.T., “Bayesian color constancy”, *Journal of the Optical Society of America, A*, Vol. 14, No. 7, pp. 1393-1411, 1997.

[2] Jaynes, E.T., “Prior Probabilities,’ IEEE Trans. on Systems Science and Cybernetics, SSC-4, pp. 227-241, 1968.

[3] Maloney, L. T. and Wandell B. A., “Color constancy: a method for recovering surface spectral reflectance”, *Journal of the Optical Society of America, A*, Vol. 3, pp. 29-33, 1986.

[4] Munsell Book of Color, Matte Finish Collection, Munsell Color, Baltimore, USA, 1976.

[5] Parkkinen, J. P. S., Hallikainen, J. and Jaaskelainen, T., “Characteristic spectra of Munsell colors”, *Journal of the Optical Society of America, A*, Vol. 6, No. 2, pp. 318-322, 1989.

[6] Skaff, S., Arbel, T., and Clark, J.J., “Active Bayesian Color Constancy with Non-Uniform Sensors”, 2002 International Conference on Pattern Recognition, Quebec City, Canada, pp II-681:684, August 2002

[7] Stockman, A., Sharpe, L. T., and Fach, C. C., “The spectral sensitivity of the human short-wavelength cones”, *Vision Research*, Vol. 39, pp. 2901-2927, 1999.