Modelling Reflectance by Logarithmic Basis Functions

Brian V. Funt
School of Computing Science
Simon Fraser University
Vancouver, B.C.
Canada V5A 1S6
e-mail funt@sfu.ca

Abstract

Non-linear models of reflectance based on exponentiation of linear combinations of 'pseudo-dye' basis functions are compared to linear models of the same dimension. Tests with the new Kodak database of 354 reflectances show that the non-linear model makes a modest gain over its linear counterpart.

1 Introduction

Low-dimensional linear models of spectral reflectance have proved very useful (eg. [?, ?, ?, ?]) because they allow approximations of the reflectance spectrum to be represented by only a few parameters. While 6-8 such parameters may suffice [?] for accurate approximations, in many color applications often only 3 or fewer can actually be used because color imagery provides only 3 knowns to a set of equations in which the linear-model parameters are the unknowns. This paper investigates the extent to which non-linear models might improve the accuracy in approximating reflectance given only a 3-band measurement.

Rodriguez and Stockham [?] describe a non-linear method which extracts the transmittance spectrum of transparency film based on a knowledge of the spectral transmittance functions of the 3 underlying film dyes. Their algorithm calculates the amount of each dye from the response of a 3-band scanner. Since the film obeys Beer's law, the dye amounts and scanner responses are related non-linearly. Their iterative algorithm estimates the dye amounts based on the scanner responses, constructs the transmittance spectrum corresponding to those dye amounts and then generates a set of predicted scanner responses which are then compared with the actual scanner responses in order to adjust the dye estimates.

Their method works well for film; will it also work well for modelling spectral reflectances? This paper presents results showing that in fact low-dimensional, non-linear models of reflectance approximate spectral reflectance somewhat better than linear models of the same dimension. Non-linear models, however, demand significantly more computation time.

2 Non-linear models

Non-linear models of reflectance will work well when reflectances behaves like the absorption of transparency film. For transparency it is more natural to consider transmittance $T(\lambda)$ (the inverse of absorption $A(\lambda)$) which relates the amount of light passing through the object, $I(\lambda)$ to the amount $I_0(\lambda)$ incident upon it:

$$T(\lambda) \;\equiv\; I(\lambda)/I_0(\lambda)$$

The optical density of the object is defined as:

$$D(\lambda) \;\equiv\; -log(T(\lambda)) \;\equiv\; log(A(\lambda))$$

When light passes through 2 or more objects in sequence the transmittances multiply, while conveniently the densities add. Consider light passing through surfaces $|1|, |2|, |3|$ with transmittances and densities as shown below.

$$
\begin{array}{ccccccc}
 & T_1(\lambda) & & T_2(\lambda) & & T_3(\lambda) & \\
I_0(\lambda) \rightarrow & |1| & \rightarrow I_1(\lambda) \rightarrow & |2| & \rightarrow I_2(\lambda) \rightarrow & |3| & \rightarrow I_3(\lambda) \\
 & D_1(\lambda) & & D_2(\lambda) & & D_3(\lambda) &
\end{array}
$$

$$
\begin{aligned}
T_{123}(\lambda) \;&=\; I_3(\lambda)/I_0(\lambda) \\
&= \; (I_1(\lambda)/I_0(\lambda))(I_2(\lambda)/I_1(\lambda))(I_3(\lambda)/I_2(\lambda)) \\
&= T_1(\lambda)T_2(\lambda)T_3(\lambda)
\end{aligned}
$$

$$
\begin{aligned}
D_{123}(\lambda) \;&=\; -log(T_1(\lambda)T_2(\lambda)T_3(\lambda)) \\
&= \; D_1(\lambda) \; + \; D_2(\lambda) \; + \; D_3(\lambda)
\end{aligned}
$$

In the case of 3-layer color negative transparency film, the 3 layers, usually cyan, magenta and yellow are superimposed so that the light passes through them in sequence just as for the case of the 3 objects discussed above. For film, the density spectra of the dyes used in the layers are fixed and measurable. The amount of each dye, however, varies according to how the film is exposed. When the dye amounts ω_i and dye densities D_i $(i = 1..3)$ are known, the resulting film density for film obeying Beer's law becomes completely specified as:

$$D(\lambda) \;=\; \sum_{j=1}^{3} \omega_j D_j(\lambda)$$

The 3 parameters ω_j specify the density linearly and the transmittance non-linearly as:

$$T(\lambda) \;= -exp(\sum_{j=1}^{3} \omega_j D_j(\lambda))$$

2.1 Film Transmittance from RGB

Given RGB measurements from a 3-channel film scanner, Rodriguez and Stockham [?] seek the film's transmittance as represented by the ω_i dye amounts. For a scanner with sensitivities $r(\lambda_i)$, $g(\lambda_i)$ and $b(\lambda_i)$ known at wavelengths $\lambda_1 ... \lambda_M$ and producing measured sensor responses R, G, B, the ω_j are calculated via non-linear least-squares minimization of:

$$
\begin{aligned}
&min\{(R - \sum_{i=1}^{M} -exp(\sum_{j=1}^{3} \omega_j D_j(\lambda_i))r(\lambda_i))^2 \\
&+(G - \sum_{i=1}^{M} -exp(\sum_{j=1}^{3} \omega_j D_j(\lambda_i))g(\lambda_i))^2 \\
&+(B - \sum_{i=1}^{M} -exp(\sum_{j=1}^{3} \omega_j D_j(\lambda_i))b(\lambda_i))^2\}
\end{aligned}
$$

This kind of minimization can be solved by standard numerical techniques. For example, the S-plus language [?] used for all the experiments reported here employs the Gauss-Newton algorithm in its function called nls. Basically, the dye amounts are being chosen so as to produce a transmittance spectrum which, when integrated with the scanner sensitivity functions to produce a predicted scanner response, results in the predicted and actual responses being as similar as possible.

It is helpful to bear in mind an obvious and well-known consequence of the non-linear dependence of transmittance upon dye density, which is that film transmittances will not be modelled exactly by a 3-parameter linear model. Changing a dye's concentration changes not only the intensity of the transmitted light, but its spectrum as well, so the transmittance spectra of the 3 dyes do not form a set of 3 fixed colors for additive mixing the way a CRT's primaries do.

3 Non-Linear Reflectance

Is reflectance analogous to film absorption? In particular, is there a small set of 'pseudo-dyes' for reflectances that might assume the role of the dye layers in color film? As for the case of arguments made in favor of linear models [?], the question is not to be answered in terms of a physical model, but rather in terms of statistical properties.

The basic presumption behind using a non-linear model is that darker colors of roughly the same hue result from increased amounts of pigment. In the case of paints mixed by adding dyes to a white base paint a non-linear model might be expected to work quite well. However, in the case of printing processes where colors are made darker by increasing the percentage of area covered with an ink rather than by the thickness of the ink layer, a linear model would be more appropriate.

3.1 Pseudo-dyes

To provide a non-linear model of reflectance, a set of basis vectors corresponding to the dye density functions must be found. In the linear-model approach, a basis set is calculated as the set of vectors obtained using a characteristic vector analysis [?] of a representative set of reflectance spectra. The discretely sample spectra are treated as vectors and which are used to form the rows of a matrix which then is decomposed via singular value decomposition. The vectors representing spectra generally are normalized to unit magnitude.

If the reflectance basis vectors are $S_j(\lambda), j = 1...N$, then the N-dimensional, linear model approximation to reflectance $S(\lambda)$ with weighting coefficients σ_j is given by

$$S(\lambda) \;=\; \sum_{j=1}^{N}[\sigma_j S_j(\lambda)]$$

The non-linear case hypothesizes that densities combine linearly from which it follows by analogy to the linear case that a linear basis for the density space can be found via the same characteristic vector analysis, but now applied to the set of vectors obtained by taking the *logarithm* of each spectra in the set of representative reflectance spectra. Since the scaling involved in normalizing the spectra would introduce an undesirable translation in logarithm space, they are not normalized in the non-linear analysis.

3.2 Linear and Non-Linear Reflectance Estimates

For the general case of an n-band sensor with spectral sensitivity functions $R_k(\lambda), k = 1..n$ viewing surface reflectance $S(\lambda)$ under illumination $E(\lambda)$, the linear model approximation to the sensor response is:

$$\rho_k \;=\; \sum_{j=1}^{3}\sigma_j \int E(\lambda)S_j(\lambda)R_k(\lambda)d\lambda$$

In matrix form this becomes

$$\rho \;=\; \mathbf{Q}\sigma$$

where $(\mathbf{Q})_{jk} \;=\; \int E(\lambda)S_j(\lambda)R_k(\lambda)d\lambda$. For $\mathbf{Q}$ nonsingular, this can be rewritten for the coefficient vector so long as the dimension of the linear model matches the number of sensor bands.

$$\sigma \;=\; \mathbf{Q^{-1}}\rho$$

Since the reflectance basis functions are orthonormal, the reflectance spectrum reconstructed from the σ coefficient vector will provide the best approximation in the least-squares sense to the original reflectance within the n-dimensional model.

In addition to reconstructing a spectrum from sensor values, it is also helpful to know what the best possible approximating spectrum is for a model of fixed dimension. In the linear case, the best coefficients of the best approximation are simply the coordinates of the spectrum projected onto the basis vectors. In the non-linear case, the coefficients are found by non-linear least-squares minimization of

$$min\left\{\sum_{i=1}^{M}[S(\lambda_i) - exp(\sum_{j=1}^{N}\omega_j D_j(\lambda_i))]\right\}^2$$

In the special case where the sensor sensitivities behave like dirac delta functions, the non-linear case becomes linear.

Instead of doing a non-linear least squares minimization, another alternative would be to do a linear fit of $log(S(\lambda))$ to the pseudo-dye basis vectors, reconstruct the density spectrum and then exponentiate. This, however, would not guarantee the best fitting reconstruction. The fit to the density basis to $log(S(\lambda))$ may be good, but then any errors are amplified by the exponentiation.

Convergence to the global minimum has not been proven for the non-linear case, however, Rodriguez et.al.[?] reported stable performance for their algorithm and similarly in all the tests reported below the non-linear least squares algorithm converged in under 10 iterations to the same answer even when the starting conditions were varied.

4 Results

The non-linear and linear fitting methods were tested on the new database of reflectance spectra provided by Imaging Research Laboratories, Eastman Kodak Company[?]. This data is broken down into 3 subsets: 64 Munsell chip spectra, 120 Du Pont paint chip spectra, and spectra of 170 assorted objects. The data is provided in 2nm and 10nm sampling intervals. For all the calculations describe here, the 10nm data in the range 400-700nm was used.

Tables 1-3 show the results of fitting the complete spectra from the different subsets of the Kodak re-

Dimensions	Non-Linear Mean Error	Linear Mean Error	Percent Mean Increase	Linear Case Unrealizable	Non-Linear Max Error	Linear Max Error	Percent Max Error Increase
3	7.621	11.797	54.796	31	29.332	45.886	56.437
4	4.857	6.949	43.072	26	14.015	32.904	134.777
% Variance Dupont	0.8778247	0.9732508	0.9948155	0.9980541	0.9990478	0.9994800	0.9996748
% Variance Log Dupont	0.8723005	0.9769323	0.9974153	0.9986852	0.9993321	0.9996388	0.9998145

Table 1: Comparison of average and maximum errors for full-spectrum linear and non-linear fits to the 120 reflectances in the 'Du Pont paint chip' subset. Also shown is the percentage of the variance accounted for by the first 7 characteristic vectors.

Dimensions	Non-Linear Mean Error	Linear Mean Error	Percent Mean Increase	Linear Case Unrealizable	Non-Linear Max Error	Linear Max Error	Percent Max Error Increase
3	10.004	10.373	3.689	10 out of 170	40.038	43.882	9.601
4	7.765	6.954	-10.444	10 out of 170	36.948	43.446	17.587
% Variance Objects	0.9285662	0.9749974	0.9913733	0.9948835	0.9971507	0.9984382	0.9990555
% Variance Log Log Objects	0.9487006	0.9820347	0.9953819	0.9977805	0.9986035	0.9992028	0.9995413

Table 2: Comparison of average and maximum errors for full-spectrum linear and non-linear fits to the 170 reflectances in the 'objects' subset. Also shown is the percentage of the variance accounted for by the first 7 characteristic vectors.

flectance database to both linear and non-linear models of 3 and 4 dimensions. Table 4 shows the results for the combined database of all the Kodak reflectances. The errors represent the square root of the sum of squares difference between the exact spectra S_{exact} and approximated spectra S_{approx} normalized relative to the magnitude of S_{exact} according to the formula:

$$error \equiv \frac{|S_{exact} - S_{approx}|}{|S_{exact}|}$$

The percentage increase in error in moving from the non-linear to the linear model is calculated as

$$\frac{(linear\ error) - (nonlinear\ error)}{(nonlinear\ error)} \times 100$$

Tables 1-4 also list the percentage of the total variance accounted for by the first 7 characteristic vectors for the linear and logarithm cases.

Table 5 lists the errors in the spectra recovered from a 3-band sensor. The CIE $\bar{x}, \bar{y}, \bar{z}$ matching functions [?] act as the sensor sensitivities so effectively the spectra are being recovered from their CIE XYZ coordinates under equal-energy white light. Since in both the linear and non-linear cases, the equations are solved for spectra that produce the specified sensor response, the recovered spectra are necessarily metameric to the actual spectra and so will always have ΔE values of zero.

Computation for the non-linear case requires roughly 5 seconds per spectrum estimated from RGB versus only 0.01 seconds for the linear case (S-plus on SUN Sparc-10).

In the linear case, nothing forces the best fitting spectra in a least-squares sense to be physically realizable, so they may well contain negative values at some wavelengths. The tables list the number of spectra having at least one negative value. For the non-linear case, however, physical realizability is guaranteed since exponentiation is always non-negative.

Dimensions	Non-Linear Mean Error	Linear Mean Error	Percent Mean Increase	Linear Case Unrealizable	Non-Linear Max Error	Linear Max Error	Percent Max Error Increase
3	8.563	10.876	27.012	5 out of 64	27.199	27.095	-0.382
4	5.492	6.851	24.745	2 out of 64	23.420	24.872	6.200
% Variance Munsell	0.8487322	0.9569718	0.9915236	0.9964170	0.9984046	0.9989884	0.9994437
% Variance Log Munsell	0.8965447	0.9727840	0.9962366	0.9984542	0.9991402	0.9994633	0.9997248

Table 3: Comparison of average and maximum errors for full-spectrum linear and non-linear fits to the 64 relectances in the 'Munsell chip' subset. Also shown is the percentage of the variance accounted for by the first 7 characteristic vectors.

Dimensions	Non-Linear Mean Error	Linear Mean Error	Percent Mean Increase	Linear Case Unrealizable	Non-Linear Max Error	Linear Max Error	Percent Max Error Increase
3	9.840	11.704	18.943	43 out of 354	41.109	43.720	6.351
4	7.278	8.629	18.563	36 out of 354	34.900	33.257	-4.708
% Variance Combined	0.8881759	0.9691406	0.9922569	0.9960211	0.9977136	0.9985319	0.9991218
% Variance Log Combined	0.9075316	0.9780640	0.9959379	0.9978212	0.9986291	0.9992057	0.9994792

Table 4: Comparison of average and maximum errors for full-spectrum linear and non-linear fits to the combined set (Du Pont, objects, Munsell) of Kodak 354 reflectances. Also shown is the percentage of the variance accounted for by the first 7 characteristic vectors.

5 Conclusion

Overall the non-linear model performs 17.5% better than the linear model in approximating surface reflectance. Unfortunately, it does not do better on the 'objects' subset. A significant advantage of the non-linear method is that it is guaranteed to generate spectral estimates that are positive at all wavelengths. Increased computational time is its prime disadvantage.

6 Acknowledgment

This research was funded by the Natural Sciences and Engineering Research Council of Canada. I wish to thank Eastman Kodak Company, Rochester New York, for providing the reflectance data in electronic form. Thanks also to Graham Finlayson, Michael Brockington and Louis Brassard for their helpful comments.

Reflectances	Non-Linear Mean Error	Linear Mean Error	Percent Mean Increase	Non-linear Max Error	Linear Max Error	Percent Max Error Increase
Dupont	11.682	16.364	40.07	56.952	68.798	20.80
Munsell	11.187	14.154	26.52	31.552	36.953	17.12
Objects	13.655	13.233	-3.08	60.214	74.553	23.81
Combined	13.193	15.496	17.46	53.55	67.453	25.96

Table 5: Percent normalized error between spectra recovered via the linear and non-linear methods given only the spectras' CIE XYZ coordinates.