# Recovering Device Sensitivities with Quadratic Programming

***Graham D. Finlayson, Steven Hordley***
***Colour & Imaging Institute***
***University of Derby***
***Derby, UK***

***Paul M. Hubel***
***H-P Laboratories***
***Hewlett-Packard Company***
***Palo Alto, CA, USA***

## Abstract

To accurately predict the response of colour devices such as digital cameras to spectral stimuli their spectral sensitivities need to be known. In this paper we present a simple, flexible method for characterising such devices, based on a single image of a Macbeth Color Checker Chart.

We begin by showing that device RGBs are linearly related to the Macbeth reflectances and that this linear relationship is defined by the spectral sensitivities of the device. It follows then that it should be possible to solve for these sensitivities by linear regression. However, this simple idea does not work well in practice - recovered sensitivities are very different from the actual device sensitivities.

The simple regression fails because the Macbeth reflectance set has limited dimensionality and so the regression is highly sensitive to image noise. To overcome this problem we incorporate a number of natural constraints: positivity, modality, and band-limitedness into the regression formulation. Each constraint can be written as a linear inequality and so solving for device sensitivities by this constrained regression is a quadratic programming problem. Posing the problem in this form, we can search for the sensors which best fit the data, quickly and efficiently, by trying different combinations of the linear constraints.

The results of performing this constrained regression on a number of colour devices are presented here. In all cases our new technique recovers sensors which are very close to the actual device sensitivities and we are confident that these results will extrapolate to other colour devices.

## 1. Introduction

The use of colour devices such as digital cameras and scanners is becoming increasingly widespread. Good colour correction of images from these devices requires that we are able to obtain colorimetric data from them. That is, we need to be able to relate the device RGB values to XYZ tristimulus values. This can be achieved in a number of ways. For example, given a set of targets we can map from camera RGBs to XYZs using interpolation and look-up tables[12,6,7] A simpler approach is to transform RGB values to XYZ values using a linear transform i.e. a single $3 \times 3$ matrix. Such linear transforms can be determined by making certain assumptions about the spectral correlation statistics of the scene. Previous authors[2] have determined a $3 \times 3$ matrix using only the assumption that all colour signals are equally likely - the so called maximum ignorance assumption. In more recent work[1] they added the further constraint that the white-point be preserved under the mapping. The ease with which this matrix can be calculated and the fact that its visual results compare favourably with other methods based on surface reflectance statistics[5] has led to its proposal[4] as a standard method for transforming from device co-ordinates to tristimulus co-ordinates.

Unfortunately, determining the maximum ignorance linear transform for a device, is not possible unless its spectral sensitivities are known. Knowledge of a device's spectral sensitivity curves is important in other applications too. For example, to obtain good colour reproduction of images from a digital camera the scene illuminant must be known. Practical algorithms to estimate the scene illuminant from the image data have been developed[3] but they can only be applied when the camera sensitivities are known. Since in general manufacturers do not make the spectral characteristics of their sensors available, the user must characterise their own device.

The standard technique for measuring the spectral sensitivities of a colour device requires the user to record the device's response to monochromatic light across the visible spectrum. This is a difficult task since it involves using a monochromator (an expensive piece of physics equipment) or a series of very narrow-band[17,18] interference filters (which too are expensive, and their use is tedious). These difficulties make measuring the spectral sensitivity curves impractical for most users.

In this paper we present an alternative approach to device characterisation, formulating the problem as one of constrained regression. Rather than relying on making accurate measurements of the device sensitivities, our approach requires only, that we can record the response of the device to a number of objects of known surface reflectance (for example a Macbeth Color Checker Chart) under known illumination conditions.

As we will show, the responses of a device (typically a set of RGB values) are linearly related to the spectral surface reflectance data, and furthermore this relationship is defined by the sensor response functions. It follows that we should be able to solve for the sensors by a simple linear regression. However a a regression of this kind leads to very poor estimates of the sensors. This poor performance is due to the limited dimensionality of surface reflectance functions, which leads to sensor estimates which are highly sensitive to noise in the device responses. To overcome this problem we incorporate into the problem formulation a number of natural constraints; sensitivity functions should be positive, have limited modality, and be band-limited. We then solve for the sensor functions by a constrained regression. The fact that all the constraints are linear allows us to pose the regression in a quadratic programming form[8], easily solvable using a package such as MATLAB.

In Section 2 of this paper we state the problem of sensor estimation as a linear regression problem and show that an unconstrained regression leads to poor sensor estimates. We present our method for constraining the regression and its formulation as a quadratic programming problem in Section 3 and relate this new method to previously published work[17,11]. In Section 4 we present the results of using this new method to characterise a number of colour devices. The results show that in all cases the recovered sensors are very close to the actual (measured) sensors. Finally we summarise this work and draw some conclusions from it in Section 5.

## 2. Statement of the Problem

For the purposes of this paper we assume that the output of our device is a three-band digital image. At each pixel, $j$ the value recorded in the $i$th band of the image is given by the equation:

$$r_i^j = \int_\omega E(\lambda) S_j(\lambda) R_i(\lambda) d\lambda, \qquad i = 1,2,3 \qquad (1)$$

where $E(\lambda)$ is the spectral power distribution of the scene illuminant, $S_j(\lambda)$ is the surface reflectance imaged at pixel $j$, $R_i(\lambda)$ is the spectral sensitivity of the $i$th sensor, and the integral is taken over the visible spectrum $\omega$. For practical purposes it is sufficient to approximate the various continuous spectra by their value at a number of discrete sample points; typically 31 sample points are used, between 400nm and 700nm at 10nm intervals. Adopting these approximations the integral is replaced by a summation and Equation 1 becomes:

$$r_i^j = \sum_{k=1}^{31} E(\lambda_k) S_j(\lambda_k) R_i(\lambda_k), \qquad i = 1,2,3 \qquad (2)$$

Or, re-writing Equation 2 in vector form:

$$r_i^j = \underline{c}_j^t \underline{R}_i, \qquad i = 1,2,3 \qquad (3)$$

where $c_{jk} = E(\lambda_k) S_j(\lambda_k)$ and $R_{ik} = R_i(\lambda_k)$. Now suppose we capture an image of a number of objects of known reflectance under known illumination conditions; here we take as our objects the 24 patches of the Macbeth Color Checker Chart which represent a range of naturally occurring reflectance spectra. Averaging the recorded sensor values for each patch results in 24 equations of the form of Equation 3. Placing the 24 sensor values of the $i$th sensor in a vector $\underline{r}_i$, we can combine these 24 equations into a single matrix equation:

$$\underline{r}_i = C \underline{R}_i, \qquad i = 1,2,3 \qquad (4)$$

where the $j$th element of $\underline{r}_i$ is $r_i^j$ and the $j$th row of $C$ is $\underline{c}_j$. Equation 4 represents a system of 24 equations in 31 unknowns - $R(\lambda_k), \quad (k = 1, \ldots 31)$ an under-determined system. Although we cannot find an exact solution to this set of equations we can find a non-unique least-squares approximation to $\underline{R}$, such that this approximation $\hat{\underline{R}}$ minimises the residual squared error:

$$\| C \hat{\underline{R}}_i - \underline{r}_i \|^2 \qquad (5)$$

The bottom graph of Figure 1 shows the result of recovering the sensors (top of Figure 1) of a typical 3-band *rgb* digital camera using a least-squares regression; clearly a very poor estimate of the camera sensitivities. The reason that an unconstrained regression performs so badly is that though the surface spectral reflectance functions are represented by 31 dimensional vectors, the true dimensionality of these functions is much less (studies suggest that the 24 reflectance functions of the Macbeth chart can be accurately represented by as few as 3 basis functions[14]. Moreover adding more reflectances or using a different set of surface reflectance functions is unlikely to improve the situation since it has been demonstrated[9,13] that naturally occurring surface reflectance functions can be approximated by between 6 and 9 basis functions. This reduced dimensionality results in an estimate of $\underline{R}_i$ which is highly sensitive to noise in the measured sensor values.

## 3. Constraining the Problem

The spectral sensitivity curves of the digital camera which are shown in Figure 1 exhibit a number of features which are likely to be common to the majority of device sensors.

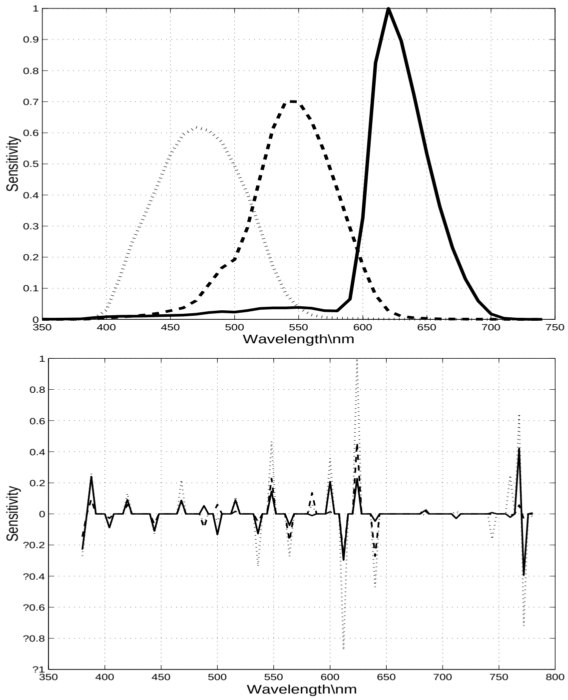


*Figure 1*: *The sensor curves of a typical 3-band digital camera (top) and the sensors obtained by a simple linear regression (bottom).*

First, the curves are everywhere positive; as all sensors will be since a device cannot have a negative response to a stimulus. Second, the curves in Figure 1 are uni-modal — they have a single peak. While we cannot be sure that all sensor curves will have only a single peak it is likely that the number of peaks will be relatively few – allowing curves to be bi-modal or tri-modal will cover the majority of sensor curves. Finally the curves are band-limited; a property shared by all device sensors. We propose that by incorporating these natural constraints into the regression formulation, we can constrain the problem enough to overcome the noise sensitivity problem and thus obtain accurate estimates of the sensors.

A number of other authors have taken a similar approach to sensor estimation. For example, in their work on scanner characterisation, Sharma and Trussel[11] imposed the constraints of positivity and smoothness (by bounding a discrete estimation of the second derivative), and also bounded the maximum, and sum of squares error. Noting that all these constraints are convex, they formulated the regression problem in a set-theoretic form and used the mathematical technique of Projections Onto Convex Sets (POCS) to solve for the sensors. However POCS is an iterative method, and although it is guaranteed to give a solution that satisfies all the constraints (provided the solution set is non-empty) this solution is not optimal. Furthermore the solution returned by POCS is quite sensitive to the choice of the initial iteration point. In later work Barnard[16] used essentially the same constraints to solve for a set of camera sensors. However he reformulated the problem as one of minimising the sum of squares relative error subject to the linear constraints, and solved the regression using a standard numerical solution method.

Importantly, as well as being intuitive, all of the constraints we have introduced here can also be written as simple linear inequalities. This linearity allows us to formulate the regression as a quadratic programming[8] problem; that is we can pose the problem as one of minimising the quadratic objective function:

$$\|C\underline{R}_i - \underline{r}_i\|^2$$

subject to a set of linear constraints:

$$\begin{array}{lcl} a_{1,1}R_i(\lambda_1)+ & \dots & +a_{1,31}R_i(\lambda_{31}) \leq b_1 \\ a_{2,1}R_i(\lambda_1)+ & \dots & +a_{2,31}R_i(\lambda_{31}) \leq b_2 \\ \vdots & \dots & \vdots \\ a_{m,1}R_i(\lambda_1)+ & \dots & +a_{m,31}R_i(\lambda_{31}) \leq b_m \end{array} \tag{6}$$

where the $a_{i,j}, b_i$ will be determined by the nature of the constraint which is being applied. Equally we could choose to minimise the sum of squares relative error:

$$\sum_{j=1}^{24} \left. \frac{C^j\underline{R}_i - r_i^j}{r_i^j} \right) \tag{7}$$

as Barnard[16] suggests. Quadratic programming is exactly the same as linear programming except for the quadratic objective function. Like linear programming there is always a unique global optimum and it is always found. When the regression is posed in this form it can easily be solved using standard mathematical software (MATLAB for example), and moreover the simplicity of the method allows one to search many different combinations of constraints to search for the best estimates of the sensors. We now consider each of the three constraints in more detail, showing that they can be formulated as sets of linear inequalities, and we give some guidelines as to their use.

**Positivity**

The constraint that our sensors are positive is written as :

$$R_i(\lambda_k) \geq 0, \qquad k = 1, \dots, 31, \qquad i = 1, 2, 3 \tag{8}$$

To solve for the sensor curves we must find the $\underline{R}_i$ satisfying Equation 6 where the constraints in this equation are the constraints in Equation 8.

**Modality**

Modality refers to the number of peaks in a sensor curve. For example, all curves in Figure 1 have a single peak, that is they are uni-modal. To see how this constraint can be incorporated into the regression problem consider a curve which has its peak at the $m$th sample point. Uni-modality can then easily be expressed as a set of linear constraints:

$$\begin{array}{ll} R_i(\lambda_{k+1}) \geq R_i(\lambda_k), & k = 1, \ldots, m-1 \\ R_i(\lambda_{k+1}) \leq R_i(\lambda_k), & k = m, \ldots, 31 \end{array} \tag{9}$$

Of course not all sensor curves will be uni-modal but we can easily pose $n$-modality as a similar set of linear constraints; we simply have to choose the sample points where the peaks (and troughs) occur. While in theory we do not know where these peaks will be, in practice we have a reasonably good idea. For example, it is unlikely that the long-wave sensitive mechanism will have peak response in the short-wave region of the visible spectrum. By trying all 'plausible' combinations of peak position, we can solve for the best sensors overall. Because regression in general, and quadratic programming in particular, can be computed quickly the overhead of a combinatorial search is small; trying all plausible combinations is quite feasible.

**Band-limitedness**

By band-limitedness we mean that the sensor curve has a response to light only in a finite range of (low) frequencies. The cone sensors in the human eye have this characteristic and so too do all device sensors. This band-limitedness allows us to represent our sensor curves as linear combinations of a set of band-limited basis functions. There are many possible basis functions we could choose but in this paper we use a standard Fourier basis. That is we represent the camera curves as a linear combination of sine and cosine functions:

$$\underline{R}_i = \sigma_1 B_1 + \sigma_2 B_2 + \ldots + \sigma_l B_l \tag{10}$$

The first few basis functions are:

$$\begin{array}{lll} B_1 = k, & B_2 = sin(x), & B_3 = cos(x), \\ B_4 = sin(2x), & \ldots & \end{array} \tag{11}$$

for constant $k$ and $x = (\lambda - 400)\pi/150$ for $\lambda$ =400nm, $\ldots$, 700nm in 10nm intervals. Equation 4 can now be rewritten:

$$\underline{r}_i = CB\underline{\sigma}_i, \qquad i = 1, 2, 3 \tag{12}$$

where the columns of $B$ are the basis functions $\underline{B}_i$. The problem now becomes to solve for $\underline{\sigma}_i$; the vector of weighting functions for the basis. Equation 11 can be solved by regression and as before we can incorporate the positivity and modality constraints into the regression. The basis function representation limits the high frequency components of the sensor curve in the visible spectrum and so can be thought of as a constraint on the smoothness of the curve. The number and type of basis functions employed controls how smooth the recovered functions will be. In general we have found somewhere between 9 and 15 basis functions to be sufficient.

## 4. Results

We have tested our recovery method on a number of device sensors and have found it to perform well in all cases. We present results for recovering the sensors of three different colour devices here; the sensors of two 3-band digital cameras and a Sharp JX400 colour scanner.

The first stage in the recovery procedure is to obtain the response of the device to a number of objects of known surface reflectance, under a known illuminant. In these experiments we used the 24 patches of the Macbeth Color Checker Chart under standard D65 illumination. For the two cameras the averaged RGB values from an image of these 24 patches form the vectors $\underline{r}_i$ in Equation 4. However, in the case of the scanner, the actual device was unavailable to us, so we generated synthetic RGB values from the measured curves. To these synthetically derived values we added a small amount of random noise (approximately 1% of the signal) to simulate the noise in the measurement process. Spectral reflectance data (the rows of $C$ in Equation 4) can either be measured using a spectrophotometer, or generated synthetically from published[15,10] data. We have found that using published values for the Macbeth Chart[10] and standard D65 illumination[15] gives good results. However, the published data for the Macbeth Chart is only for wavelengths between 400nm and 700nm. In practice many sensors will have a significant response outside this range, so extending the range over which the regression is performed will improve results.

To recover the sensors we then simply have to solve the quadratic programming problem described above. Typically, the number of peaks in our sensors and their location in the visible spectrum will be unknown and we must search for the location which minimises the error. For example, if we assume our sensor to be uni-modal, then it's single peak could fall anywhere in the visible spectrum. In this case we would perform 31 regressions (assuming that we are working with data in the range 400nm to 700nm sampled at 10nm intervals) and get 31 estimated sensors. We choose as our estimate of the sensor, the estimate which minimises the regression error. Similarly, if we allow our sensor to be bi-modal or tri-modal, we must try different combinations for the peaks to find the best match.

Of course, if information as to the location of the sensors' peaks is known this can be used to restrict the search.

Figure 1 shows the sensor curves of our first digital camera. It is clear from this image that the sensors are all positive, uni-modal, and reasonably smooth. Figure 2 shows the results of applying our constrained regression to recover an estimate of these curves. All three of our estimated sensors (dashed lines) match the estimated sensors (solid lines), measured using a monochromator, very closely.

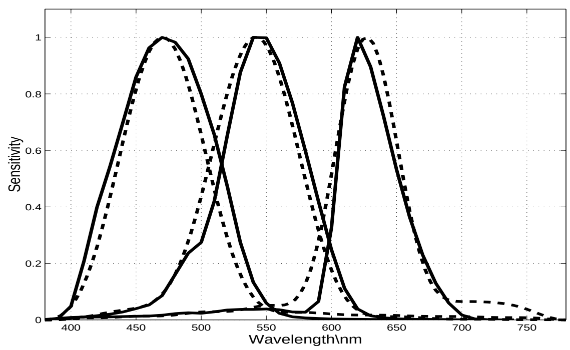


*Figure 2*: *Recovered sensors for a digital camera. The measured sensors (solid lines) are shown together with the recovered sensors (dashed lines). Each sensor is normalised such that its maximum sensitivity is 1.*

The sensor curves of a Kodak DCS460 digital camera, again measured using a monochromator, are shown as the solid lines in Figures 3 and 4. As in the case of the first camera, the red and green sensors are smooth and uni-modal. However, the blue sensor has two clear peaks and attempting to estimate this sensor without allowing bi-modality is likely to result in a poor estimate. Figure 3 shows the measured red and green sensors (solid lines) together with our estimates of them (dashed lines). Again our estimates match the measured data very closely. Figure 4 shows the results of estimating the blue sensor. The measured data (solid line) is shown together with an estimate of the sensor restricted to be uni-modal (dotted line), and an estimate allowed to be bi-modal (dashed line). Clearly the estimates of this sensor is not as good as we would have liked. However, we note that the correspondence between measured blue values and those predicted from the blue sensor (measured with the monochromator) is not good. This discrepancy may be due to noise in the image.
The measured sensors of the final device we tested; the-Sharp JX400 colour scanner are shown as the solid lines in Figure 5. Like the curves of the two cameras these scanner curves are positive and band-limited, and the red and blue curves are uni-modal. The green curve however has a second small peak in the short-wavelength region of

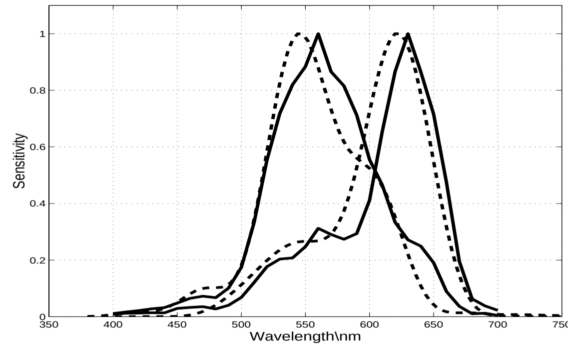


*Figure 3*: *Recovered red and green sensors for the Kodak DCS460 digital camera. The measured sensors (solid lines) are shown together with the recovered sensors (dashed lines). Each sensor is normalised such that its maximum sensitivity is 1.*

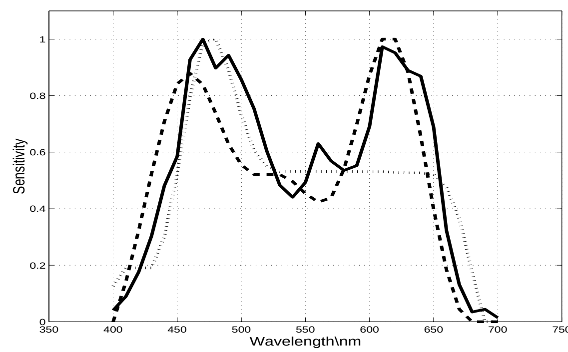


*Figure 4*: *Recovered blue sensor for the Kodak DCS460 digital camera. The measured sensor (solid line) is shown together with an estimate constrained to be uni-modal (dotted line) and an estimate without this constraint (dashed line). Each sensor is normalised such that its maximum sensitivity is 1.*

the visible spectrum. The estimated sensors are shown as dashed lines in Figure 5. Once more our recovered estimated match the measured data very closely, and again the allowing bi-modality results in a slightly better fit for the green sensor than would otherwise be possible.

## 5. Conclusions

Good colour correction of images taken with devices such as digital cameras and scanners requires that the spectral sensitivity curves of these devices be known. In this paper we have shown that these spectral sensitivities define a linear relationship between the measured camera RGBs and the surface reflectance functions of the imaged objects. This linear relationship implies that if we have a set of camera RGBs for objects of known surface reflectance we can solve for the camera sensitivities. We have seen however that solving for the camera sensitivities by simple

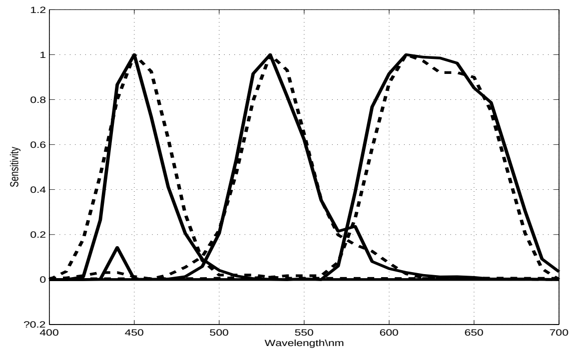


*Figure 5*: *Recovered sensors for a colour scanner. The actual sensors (solid lines) are shown together with the recovered sensors (dashed lines). Each sensor is normalised such that its maximum sensitivity is 1.*

regression results in a very poor estimate of the sensors. This led us to incorporate the linear constraints of positivity, modality, and band-limitedness into the regression by posing the problem in a quadratic programming form.

This results in a simple but very effective method for recovering sensor curves. Results presented in this paper for three devices show that the recovered sensor curves are very close to the actual curves and are likely to be good enough for all practical purposes,

## Acknowledgements

We would like to thank the EPSRC and Hewlett-Packard Incorporated for their support of this work.

## References

[1] G.D. Finlayson and M.S. Drew. White-point preserving color correction. In *Proceedings of the Fifth Color Imaging Conference: Color Standards and Color Measurements*, pages 258–261. IS&T/SID, November 1997.

[2] G.D. Finlayson and M.S. Drew. Constrained least-squares regression in color space. *Journal of Electronic Imaging*, 1997.

[3] Graham D.Finlayson, Paul M. Hubel and Steven Hordley. Color by Correlation. In *Proceedings of the Fifth Color Imaging Conference*, pages 6–11. IS&T/SID, November 1997.

[4] Jack Holm. Issues relating to the transformation of sensor data into standard color space. In *Proceedings of the Fifth Color Imaging Conference*, pages 290–295. IS&T/SID, November 1997.

[5] Paul M. Hubel, Jack Holm, Graham D. Finlayson, and Mark S. Drew. Matrix Calculations for digital photography. In *Proceedings of the Fifth Color Imaging Conference*, pages 105–111. IS&T/SID, November 1997.

[6] P.C. Hung. Color rendition using three-dimensional interpolation. In *Imaging Applications in the Work World*, Vol. 900, pages 111–115. SPIE 1988.

[7] P.C. Hung. Colorimetric calibration for scanners and media. In *Proceedings of the 1991 Electronic Imaging Meeting, San Jose, CA*, February 1991.

[8] Charles L. Lawson and Richard J. Hanson. *Solving Least Squares Problems.* Prentice Hall, 1974.

[9] Lawrence T. Maloney. Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *Journal of the Optical Society of America, A*, 3(10):1673–1683, 1986.

[10] C.S. McCamy, H.Marcus, and J.G. Davidson. A color rendition chart. *Journal App. Photog. Eng.*, pages 95–99, 1976.

[11] Guarav Sharma and H.J. Trussel. Characterization of scanner sensitivity. In *Proc. IS&T/SID Color Imaging Conference: Transforms and Portability of Color*, pages 107–103, November 1993.

[12] S. Suzuki, T. Kusunoki, and M. Mori. Color characteristic design for color scanners. *Applied Optics*, 29:5187–5192, 1990.

[13] M.J. Vrhel, R. Gershon, and L.S. Iwan. Measurement and analysis of object reflectance spectra. *Color Research and Application*, 19(1):4–9, 1994.

[14] Brain A. Wandell. *Foundations of Vision*. Sinauer Associates, 1995.

[15] G.Wyszecki and W.S. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulas.* New York:Wiley, 2nd edition, 1982.

[16] Kobus Barnard. Computational color Constancy: Taking theory into practice. Masters Thesis, Simon Fraser University, 1995.

[17] Poorvi L. Vora, Joyce E. Farell, Jerome D. Tietz, and David Brainard, Digital color cameras - 2 - Spectral response, HP Technical Report, March 1997.

[18] P.M. Hubel, D. Sherman, and J. E. Farrell, A Comparison of Methods of Sensor Spectral Sensitivity Estimation, In Proceedings *IS&T/SID 2nd. Color Imaging Conference: Color Science, Systems and Applications*, pages 45–48, November 1994.