

Robust Color Sensor Response Characterization

Bob Dyas
Motorola Labs
Schaumburg, Illinois

Abstract

Estimation of the spectral characteristics of color image sensors is an important problem. The effective spectral response of a color imaging system may vary from the color filter specification for a variety of reasons. However, color correction and color constancy algorithms require precise characterization of the color filter response in order to attain accurate color reproduction. Also, camera systems need to be qualified and calibrated quickly in a manufacturing environment. CMOS image sensors, which are becoming more popular, complicate the problem, since noise levels on CMOS sensors have yet to equal their CCD counterparts. For these reasons, a fast and reliable characterization of the color sensor response is needed. This paper presents an approach to the problem using numerical regularization techniques. In particular, Tikhonov regularization is used to determine the spectral response of a color image sensor given the digital output codes from an image with known reflectances under a known illuminant.

Introduction

Accurate knowledge of a color image sensor's spectral response is essential in order to guarantee accurate color reproduction. Color correction and color constancy algorithms require precise knowledge of the color filter response in order to attain high image quality. Although the color filter response may be specified, very often the in-situ response varies from the specification. Sensor to sensor and lot to lot variations of color filter response are to be expected, however these variations, even if within specified bounds, may be outside the tolerance required for accurate color reproduction. Furthermore, large deviations from the specified color sensor response may indicate a process or fabrication problem. Therefore, it would be advantageous to have a quick and robust method to identify excessive variations in color response in order to qualify and calibrate color image sensors. Several sources are responsible for variations in color response. These include optical crosstalk, which occurs when light passes through one filter of a color filter array but is actually detected at an adjacent pixel with a different color filter. Electrical crosstalk occurs when electron-hole pairs generated from captured photons in one pixel migrate to an adjacent pixel and are collected there. Additionally, high temperatures are required for the fabrication of solid-state sensors, which can cause degradation in the response of some color filters. Finally, lens induced chromaticity errors can cause variations in the system color response. Techniques have been previously identified for indirectly recovering the color response. These methods have been primarily concerned with relatively low noise applications. CMOS image sensors are becoming more prevalent, however they suffer from poorer noise performance. Any process of recovering the spectral response of CMOS sensors must be robust with respect to noise. This paper will address exactly this problem, namely color sensor response recovery in noisy environments. The response recovery problem requires the linear inversion of a discrete ill-posed problem. These types of problems have been studied extensively in other areas of science and engineering. A number of numerical regularization methods have been developed to attack these problems. This paper will present the results of applying one of these techniques, namely Tikhonov regularization, to the sensor recovery problem. The subsequent sections will present a formal statement of the problem, followed by a review of previous work in the area. The next section will discuss the technique of regularization as applied to this problem. This will be followed by a discussion of experimental results and finally the conclusions drawn from them.

Problem Statement

Most color imaging systems consist of three or perhaps four different spectral filters. This paper will assume, without loss of generality, three spectral bands. The signal from the ith band for a given pixel j is:

$$b_i^j = \int E(\omega) R_j(\omega) C_i(\omega) d\omega \qquad (1)$$

Here $E(\omega)$ is the spectral power distribution of the illuminant, $R_i(\omega)$ is the reflectance of the target at pixel j and $C_i(\omega)$ is the spectral response of the ith color filter. The limits of integration are over the wavelength range of interest, usually the visible spectrum. Measurements of E and R are generally provided on a discrete grid spanning this range. The test target will consist of a collection of m reflectances sampled at n points. For this paper a MacBeth ColorChecker chart will serve as the reflectance target so m=24. The illuminant is assumed to have a smooth spectral response and a sampling grid of $n<m$ is used. This implies

that measurements of E and R sampled at a higher resolution must be integrated to coincide with this grid. Consequently, an error free discretization of (1) can be written in matrix form as:

$$\mathbf{B} = \mathbf{A}\,\mathbf{C} \tag{2}$$

The illuminant and target reflectance have been combined into a single $m \times n$ matrix, $\mathbf{A} = \mathbf{R}\,\mathbf{E}$, where $\mathbf{E}$ is a diagonal $n \times n$ matrix containing the measurements of $E(\omega)$ and $\mathbf{R}$ is an $m \times n$ matrix containing the measurements of $R(\omega)$. The digital output codes from the camera are in the $m \times 3$ matrix $\mathbf{B}$. The $m \times 3$ matrix $\mathbf{C}$ represents the spectral responses of the three color channels. Noise is introduced into the values of $\mathbf{B}$ due to the physics of the light detection process. Quantization of the sensed light also introduces a noise component. Additionally, offsets of the digital codes can contribute to the noise if not explicitly corrected. There is also an implicit scaling to convert from physical units to the digital output codes of the sensor. Consequently the actual output of the sensor is:

$$\mathbf{b}_i = (\mathbf{A}\,\mathbf{c}_i + \mathbf{e}_i)\,g \tag{3}$$

where $\mathbf{e}_i$ is the noise in the ith color channel and g is the implicit system gain. Since the computation of the spectral responses will be done on a per channel basis, the matrices $\mathbf{B}$ and $\mathbf{C}$ have been replaced with individual vectors, $\mathbf{b}_i$ and $\mathbf{c}_i$, for each of the color channels. The sensor is assumed to be linear. If not the data must be linearized using, for example, methods mentioned in [4].

Previous Work

Several methods have been previously suggested for characterization of color sensors. In [1], the technique of Projections Onto Convex Sets (POCS) was applied to the problem. Constraints on the spectral response $\mathbf{C}$, form a number of convex sets that are searched using POCS to arrive at a solution that satisfies all the constraints. These constraints include non-negativity, bounds on the residual norm and a semi-norm based on the second derivative of the estimation. The latter two constraints require parameter selection to bound the errors. Besides the problem of parameter selection, POCS does not necessarily converge to an optimal solution, only one that satisfies all the constraints. This difficulty was addressed in [3] where the same constraints were cleverly reformulated into a bounded and constrained least squares problem. However, this method also requires parameter selection to bound the errors. In [5], a Wiener estimation technique was used, but this also requires parameter selection in the form of estimates of the noise statistics. Sequential quadratic programming was used in [6], with non-negativity constraints and a smoothness constraint achieved through limiting the calculated spectral response to be constructed from a limited number of low frequency basis functions. The expected modality of the spectral response must also be specified. A common theme with these methods is the need for parameters to constrain the optimization. Selection of these parameters is not always obvious and may require significant experimentation in order to determine them. Furthermore, if a different sensor needs to be evaluated the characteristics are likely to change requiring further experimentation. A single parameter, easily searched, estimation technique was mentioned in [1] as the Principle Eigenvector Method and in [2] as the rank-deficient pseudo-inverse method. Both of these refer to the same method of computing the response using a truncated Singular Value Decomposition (SVD). The SVD of a $m \times n$, $n \leq m$, matrix $\mathbf{A}$ is a factorization such that:

$$\mathbf{A} = \mathbf{U}\,\mathbf{S}\,\mathbf{V}^T \tag{4}$$

where $\mathbf{U}$ is a $m \times m$ orthonormal matrix, $\mathbf{V}$ is a $n \times n$ orthonormal matrix and $\mathbf{S}$ is a $m \times n$ diagonal matrix consisting of the n singular values, σ_n, of $\mathbf{A}$ sorted from largest to smallest. The Moore-Penrose pseudo-inverse[10] of $\mathbf{A}$, $\mathbf{A}^+$ is defined as:

$$\mathbf{A}^+ = \mathbf{V}\,\mathbf{S}^+\,\mathbf{U}^T \tag{5}$$

where $\mathbf{S}^+$ is a diagonal $n \times m$ matrix consisting of the reciprocals of the singular values of $\mathbf{A}$. The pseudo-inverse can be used to compute the least squares solution to (2):

$$\mathbf{C}_{ls} = \mathbf{A}^+\,\mathbf{B} \tag{6}$$

Since the singular values may span several orders of magnitude (and do, for the case of the MacBeth chart reflectances), the small singular values, when inverted, will greatly amplify any errors in $\mathbf{B}$. The result is a solution dominated by these amplified errors. In order to avoid this problem, the SVD is truncated by using a reduced number, $r < n$, of the largest singular values in the computation of $\mathbf{C}$. The truncated pseudo-inverse becomes:

$$\mathbf{A}_r^+ = \mathbf{V}_r\,\mathbf{S}_r^+\,\mathbf{U}_r^T \tag{7}$$

where $\mathbf{V}_r$ is $n \times r$, $\mathbf{U}_r$ is $m \times r$ and $\mathbf{S}_r^+$ is $r \times r$. $\mathbf{A}_r^+$ is then used to compute a new solution:

$$\mathbf{C}_r = \mathbf{A}_r^+\,\mathbf{B} \tag{8}$$

As reported in [1] and [2], the results from (8) are better than those obtained from (6), but the method is not robust with respect to larger noise levels in $\mathbf{B}$.

Regularization

Regularization is the process of replacing the original problem, for example (6), with a modified or regularized problem with a stable solution less sensitive to perturbations and close to the desired solution. The truncated SVD is an attempt to regularize the color sensor recovery problem. This is the preferred approach for numerically rank deficient problems, however this problem is, strictly speaking, ill-posed. The distinction is that numerically rank deficient problems have a well-determined gap between the large and small singular values, which makes the choice of the truncation parameter, r, a straightforward matter. Ill-posed problems have gradually decaying singular values with no obvious separation between large and small singular values. Discrete ill-posed problems arise from the discretization of Fredholm integral equations of the first kind of which (1) is a special case. Several numerical regularization methods exist for handling (1) and its discrete form (2) (see [7] for a survey of these methods). One such alternative to truncating the small singular values is to make use of filter factors. Filtering rather than truncating the singular values addresses the noise amplification problem. So the regularized solution is computed as follows:

$$\mathbf{A}_{\mathrm{reg}}^{+} = \mathbf{V}\,\mathbf{F}\,\mathbf{S}^{+}\,\mathbf{U}^{\mathrm{T}} \tag{9}$$

and

$$\mathbf{C}_{\mathrm{reg}} = \mathbf{A}_{\mathrm{reg}}^{+}\,\mathbf{B} \tag{10}$$

where $\mathbf{F}$ is a diagonal $n\times n$ matrix consisting of the filter factors f_n. The truncated SVD is a specific case of applying filter factors where for $n\leq r$, the filter factors, $f_n = 1$ and for $n>r$, $f_n = 0$. Tikhonov regularization is one method that makes use of filter factors. A compelling feature of Tikhonov regularization is that all filter factors are specified with a single parameter, λ:

$$f_n = \frac{\sigma_n^2}{\sigma_n^2 + \lambda^2} \tag{11}$$

Selection of λ is achieved through use of the L-curve criterion.[9] The L-curve plots the norm of the regularized solution versus the norm of the residual on a log-log graph for various values of λ. In most cases, this plot results in a curve with a distinctive L shape. The optimal tradeoff between regularization and fit is achieved at the corner of the L-curve. Unfortunately, this approach does not result in acceptable results, particularly for noisy or highly quantized sensor data. The recovered responses have a noticeable high frequency component, not present in the actual sensor response. The reason for this is that the SVD does not provide an acceptable set of basis functions from which to reconstruct the response. An alternative, as mentioned in previous papers, is to minimize a semi-norm based on the second derivative of the solution rather than the solution norm itself. A formulation that achieves these goals involves the Generalized SVD (GSVD).[7,11] The GSVD of a matrix pair ($\mathbf{A}$, $\mathbf{L}$) is a decomposition of the form:

$$\mathbf{A} = \mathbf{U}\begin{pmatrix}\Sigma & \mathbf{0}\\ \mathbf{0} & \mathbf{I_{n-p}}\end{pmatrix}\mathbf{X}^{-1}, \quad \mathbf{L} = \mathbf{V}(\mathbf{M},\mathbf{0})\mathbf{X}^{-1} \tag{12}$$

where $\mathbf{L}$ is a $p\times n$ matrix. For a discrete approximation of the second derivative, $p=n-2$ and $\mathbf{L}$ is a banded convolution matrix of the Laplacian filter response.[1–21] Consequently, $\mathbf{U}$ is an orthonormal $m\times n$ matrix, $\mathbf{V}$ is an orthonormal $p\times p$ matrix and $\mathbf{X}$ is a nonsingular $n\times n$ matrix. The matrices $\mathbf{\Sigma}$ and $\mathbf{M}$ are $p\times p$ diagonal matrices with elements σ_p and μ_p respectively. The elements of σ_p and μ_p are positive valued and sorted in increasing order. They are also normalized such that $\sigma_p^2 + \mu_p^2 = 1$. The generalized singular values are $\gamma_p = \sigma_p/\mu_p$. The GSVD provides a new basis, $\mathbf{X}$, which allows a smooth solution to be computed using:

$$\mathbf{A}_{\mathrm{i}}^{+} = \mathbf{X}\,\mathbf{F}_{\mathrm{i}}\,\mathbf{\Sigma}^{+}\,\mathbf{U}^{\mathrm{T}} \tag{13}$$

and

$$\mathbf{c}_{\mathrm{i}} = \mathbf{A}_{\mathrm{i}}^{+}\,\mathbf{b}_{\mathrm{i}} \tag{14}$$

where $\mathbf{\Sigma}^{+}$ is a $p\times p$ diagonal matrix with elements $1/\sigma_p$. The matrix $\mathbf{F_i}$ contains the Tikhonov filter factors based on the parameter λ selected using the L-curve criterion for each of the i color filters. The filter factors for the GSVD regularized version are given by:

$$f_n = \frac{\gamma_n^2}{\gamma_n^2 + \lambda^2} \tag{15}$$

Experimental Results

A simulation of the color response recovery problem was implemented in MATLAB. Both RGB and CMY type filters were generated using a pseudo-random filter generator. Programmable levels of noise were added to the generated digital codes. This allowed the evaluation of the regularization methods over a large sample of different color filters under varying noise environments. Reflectance data from a MacBeth ColorChecker Chart under a D65 illuminant measured every 2nm from 380 to 780 nm was reduced via quadrature to 20 spectral lines. The resulting 24×20 array forms the matrix $\mathbf{A}$. Simulated camera responses, $\mathbf{B}$, are formed from the randomly generated filters. The Regularization Toolbox[8] was used to compute the optimal regularized solutions, $\mathbf{C}$, using the L-curve criterion. An example L-curve from the recovery of simulated sensor data quantized to 8 bits is shown in Figure 1. The actual and recovered sensor responses for a typical 8-bit sensor are presented in Figure 2. Clearly, the regularized solutions are very close to the actual solutions. Further experimentation resulted in some observations and uncovered some difficulties with Tikhonov regularization and the L-curve criterion as applied to the color sensor response recovery problem. As would be expected, different noise levels resulted in computation of different regularization parameters. However, for a given set of color filters but with varying noise levels, the regularized solutions were generally very close to each other. The method was also tolerant of offsets in the digital codes. However, it is clearly preferable to adjust for these offsets if at all possible as mentioned in [3]. As mentioned previously, using the solution norm in the L-curve did not generate acceptable results. The computed solutions were generally not as smooth as the actual filter response. Regularization using the second derivative semi-norm proved to be useful in generating smooth solutions. Since the regularization was applied independently to each of the three color channels, the computed Tikhonov parameter, λ, was different for each color channel. This results in different regularization for each channel. This poses a difficulty since the relative amplitudes of the three responses are important, particularly for qualifying sensors on a production line. Also, occasionally, the regularization process arrives at a clearly incorrect solution. This occurs when the L-curve does not have its characteristic L shape. The determination of the corner becomes problematic and an inappropriate value for λ is chosen. In all of these cases an over-regularized or overly smooth solution is returned. It is known that, in general, the L-curve criterion will produce a slightly over-regularized solution with respect to the true optimum.[9] To address this issue and the problem of determining the relative amplitudes of the three channels, a single value for λ is chosen to compute the final solutions. This consists of the minimum λ of those computed for each of the three channels. For most cases the three values of λ are very close, and the resulting solutions are not affected a great deal.

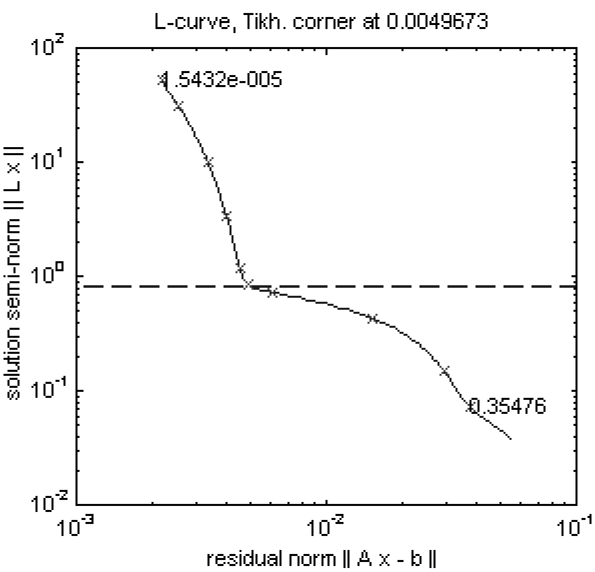

Figure 1. The L-curve for 10 values of the Tikhonov regularization parameter λ. The corner of this curve represents the best tradeoff between the size of the residual norm and the size of the semi-norm of the solution.

The relative amplitudes determined in this way are very close to the actual and an approximation of *g* from (3) can be determined from the peak value over all of the three responses. For those cases where up to two of the computed responses are incorrectly determined, the minimum λ returns an appropriate solution. In the rare case where all three responses are incorrect (determined by clearly over-smoothed solutions or L-curves without an obvious L shape), the only recourse is to choose a new **A**. This can be accomplished by repeating the experiment with a different illuminant or a target other than the MacBeth chart. This infrequent problem suggests that future research should investigate how **A** can be selected such that Tikhonov regularization is more robust across various possible sensor responses.

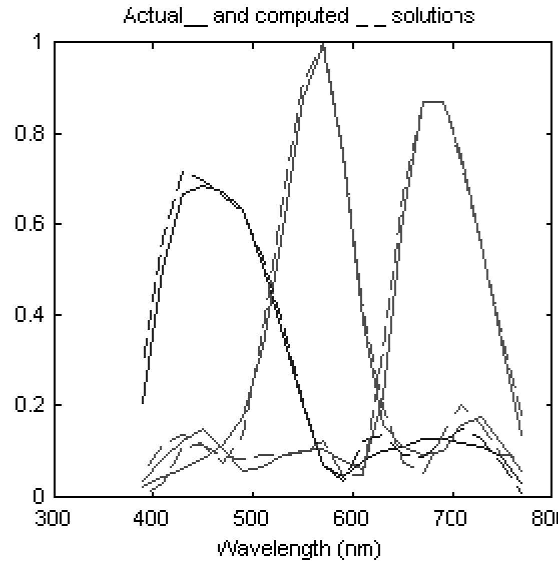

Figure 2. Typical results of recovering the color filter response of a simulated 8-bit sensor using Tikhonov regularization.

Conclusion

Indirectly determining the color sensor response can be a difficult problem, particularly for CMOS sensors that can be noisier than CCD sensors. Previous methods addressing this problem require empirical selection of multiple parameters and are generally not robust to larger noise levels. Tikhonov regularization requires selection of a single parameter that can be chosen via the L-curve criterion. This method is tolerant of higher noise levels. Although Tikhonov regularization is highly effective it occasionally fails. Those instances are obvious when the L-curve lacks an obvious L shape. For most of these occasions, choosing a single regularization parameter from the three computed for the different color channels is sufficient to achieve satisfactory results.

References

1. Gaurav Sharma and H. J. Trussell, Characterization of Scanner Sensitivity, *Proc. IS&T/SID Color Imaging Conference: Transforms and Portability of Color*, pp. 103-107, November 1993
2. P. M. Hubel, D. Sherman, J. E. Farrell, A Comparison of Methods of Sensor Spectral Sensitivity Estimation, *Proc. IS&T/SID 2nd Color Imaging Conference: Color Science, Systems and Applications*, pp. 45-48, November 1994
3. Kobus Barnhard, Computational Color Constancy: Taking Theory into Practice, Masters Thesis, Simon Fraser University, 1995

4. Poorvi L. Vora, Joyce E. Farrell, Jerome D. Dietz and David Brainard, Digital color cameras - 1 - Response models, HP Technical Report, March 1997
5. Poorvi L. Vora, Joyce E. Farrell, Jerome D. Dietz and David Brainard, Digital color cameras - 2 - Spectral response, HP Technical Report, March 1997
6. Graham D. Finlayson, Steven Hordley, Paul M. Hubel, Recovering Device Sensitivities with Quadratic Programming, *Proc. IS&T/SID 6th Color Conference: Color Science, Systems and Applications*, pp. 90-95, November 1998
7. P. C. Hansen, *Rank Deficient and Discrete Ill-Posed Problems*, SIAM, 1998
8. P. C. Hansen, Regularization Tools. A MATLAB Package for Analysis and Solution of Discrete Ill-Posed Problems, Technical Report, Dept. of Mathematical Modeling, Technical University of Denmark, 1998
9. P. C. Hansen and D. P. O'Leary, The use of the L-curve in the regularization of discrete ill-posed problems, *SIAM J. Sci. Comput*. 14, pp. 1487-1503, 1993
10. S. L. Campbell and C. D. Meyer, Jr., *Generalized Inverses of Linear Transforms*, Dover, 1991
11. Gene H. Golub and Charles Van Loan, *Matrix Computations*, Johns Hopkins University Press, 1996