Camera Characterization for Color Research

Kobus Barnard,[1*†] Brian Funt[2]

[1] Computer Division, University of California–Berkeley, Berkeley, CA 94720-1776

[2] Department of Computing Science, Simon Fraser University, 8888 University Drive, Burnaby, British Columbia, Canada V5A 1S6

Received 17 October 2000; accepted 12 July 2001

Abstract: In this article we introduce a new method for estimating camera sensitivity functions from spectral power input and camera response data. We also show how the procedure can be extended to deal with camera nonlinearities. Linearization is an important part of camera characterization, and we argue that it is best to jointly fit the linearization and the sensor response functions. We compare our method with a number of others, both on synthetic data and for the characterization of a real camera. All data used in this study is available online at www.cs.sfu.ca/~colour/data. © 2002 Wiley Periodicals, Inc. Col Res Appl, 27, 152–163, 2002; Published online in Wiley InterScience (www.interscience.wiley.com). DOI 10.1002/col.10050

Key words: camera characterization; calibration; linearization; constructed least squares; ccd noise; color camera spectral response

INTRODUCTION

The image recorded by a camera depends on three factors: the physical content of the scene, the illumination incident on the scene, and the characteristics of the camera. Because the camera is an integral part of the resulting image, research into image understanding normally requires a camera model. The most common use of camera characterization is to predict camera responses given an input energy spectral distribution. This has applications in the development and practical realization of color-related image-processing algorithms, such as computational color constancy algorithms.

In this article we begin by introducing the standard camera model used in color-oriented computer vision.[1,2] Next we discuss previous methods for fitting the parameters of that model (camera sensor response as a function of wavelength) and introduce a new method for obtaining these parameters. We then show how this method can be extended so that a proposed camera linearization function can be fit simultaneously with the camera sensor response functions. This extension takes advantage of the linearization information contained in the data required to estimate camera sensitivity. Furthermore, it is beneficial to allow the errors that result from the inadequacy of the linearization fit to be traded against errors that are a product of the inaccuracies in the sensor function fit. Finally, we provide results for both synthetic and real camera characterization experiments. All data used in this study are available online.[3]

THE CAMERA MODEL

The goal of this work is to develop a model that predicts image pixel values from input spectral power distributions. In this section we discuss the general form of the model. For the moment we assume that all camera controls such as aperture are fixed. Let $v^{(k)}$ be the value for the kth channel of a specific image pixel and let $L(\lambda)$ be the spectral power distribution of the signal imaged at that pixel. Then we model image formation by[1,2]:

$$\rho^{(k)} = F^{(k)}(v^{(k)}) = \int L(\lambda)\, R^{(k)}(\lambda) d\lambda \qquad (1)$$

where $R^{(k)}$ is a sensor sensitivity function for the kth channel, $F^{(k)}$ is a wavelength independent linearization function, and $\rho^{(k)}$ is the linearized camera response. The key assumption is that all nonlinear behavior is independent of wavelength, given the sensor. This model has been verified as being adequate for computer vision over a wide variety of systems.[2,4–8] This model is also assumed for the human

† Work completed while the first author was at Simon Fraser University.

* Correspondence to: Dr. Kobus Barnard, Computer Division, University of California–Berkeley, Berkeley, CA 94720-1776. (e-mail: Kobus@CS.Berkeley.edu).

visual system and forms the basis for the CIE colorimetry standard (here $F^{(k)}$ is the identity function).

As we move around the image plane, the signal is attenuated because of geometric effects, notably vignetting,[1] and a fall-off proportional to the fourth power of the cosine of the off-axis angle.[1] These effects can be absorbed into either $R^{(k)}$ or $F^{(k)}$. We defer these considerations by using only the central portion of the image in our experiments.

Similarly, global effects on the overall magnitude of the responses, such as camera lens aperture, and focal length, can also be absorbed into either $R^{(k)}$ or $F^{(k)}$. In fact, for much of the work in color, absolute light flux is somewhat arbitrary, being under aperture control, and is usually adjusted by the user or the camera system to give a reasonable image. Partly for this reason, work in color often uses a chromaticity space that factors out luminance. The most common such space is (r, g) defined by $(R/(R + G + B), G/(R + G + B))$. In chromaticity space geometric attenuation effects can be ignored. On the other hand, if absolute luminance is important, then these effects have to be taken into account.

Successful use of the above model requires consideration of the function $F^{(k)}$. $F^{(k)}$ reverses added gamma correction, compensates for any camera offset, and corrects for other more subtle nonlinearities. Even if $R^{(k)}$ is not required for an application, $F^{(k)}$ can be important. For example, reliably mapping into a chromaticity space such as (r, g) requires either an estimate of $F^{(k)}$ or confidence that it is the identity function and thus can be ignored. This can be the case with a camera designed for scientific use, but inexpensive consumer cameras usually do not give the operator the option to disable nonlinear behavior.

For the practical application of the above model, continuous functions of the wavelength, λ, are replaced by samplings of those functions. For example, our data were collected with a PhotoResearch PR-650 spectroradiometer, which measures data from 380 nm to 780 nm in 4-nm steps. The function $L(\lambda)$ then becomes the vector $\mathbf{L}$, $R^{(k)}(\lambda)$ becomes the vector $\mathbf{R}^{(k)}$, and Eq. (1) becomes:

$$\rho^{(k)} = F^{(k)}(\upsilon^{(k)}) = \mathbf{L} \cdot \mathbf{R}^{(k)} \tag{2}$$

Using this notation, camera characterization can be defined as finding $F^{(k)}$ and $\mathbf{R}^{(k)}$.

MOTIVATION

We became interested in color camera characterization as part of our research into computational color constancy.[9] Practically all algorithms for color constancy assume that the image pixels are proportional to the input spectral power. This is equivalent to assuming either that $F^{(k)}$ is the identity function or that it is known and has been applied to the data. In other words, color constancy algorithms require $\rho^{(k)}$ as input, as opposed to the more readily available $\upsilon^{(k)}$.

Determining the function $\mathbf{R}^{(k)}$ is also important for computational color constancy. Most algorithms, including the ones we currently think are the most promising, require an estimate of camera responses to the real world with its many different surfaces and illumination conditions. Although it is conceivable to obtain camera responses for a large number of surfaces under a given illuminant, it is impractical to obtain this data for each camera. Furthermore, some algorithms require this data for each possible illumination, including combinations of several sources. It is thus far more effective to first obtain reflectance functions and illuminant spectra and then to use a camera model to predict the wide range of camera responses required by these algorithms.

PREVIOUS WORK

Because $F^{(k)}$ is assumed to be independent of wavelength, it can be determined by stimulating the camera with varying intensities of a single light source obtained with neutral density filters or by simply moving the source. An appropriate function can then be fitted to the data, or as an alternative a smoothed version of the data can be used to generate a lookup table. Vora *et al.*[2] used this method to verify that their Kodak DCS-200 digital camera was linear over most of its operating range, and also to develop a linearization curve for their Kodak DCS-420 digital camera. They then determined $\mathbf{R}^{(k)}$ for those cameras by stimulating them with very narrow-band illumination produced by a monochrometer.[8] This method is conceptually very simple and can be very accurate. However, the equipment required to produce sufficiently intense narrow-band illumination at uniformly spaced wavelengths is expensive and not readily available. Hence, various researchers have investigated methods for characterization that do not use such equipment.

The general approach of these methods is to first measure $F^{(k)}$ and then to measure a number of input spectra and the corresponding camera responses. Let $\mathbf{r}^{(k)}$ be a vector whose elements are the linearized camera responses $\rho^{(k)}$, and let L be a matrix whose rows are the corresponding sampled spectra. Then Eq. (2) becomes:

$$\mathbf{r}^{(k)} = L\mathbf{R}^{(k)} \tag{3}$$

Eq. (3) can be solved by multiplying both sides by the pseudo-inverse of L. However, this does not work very well because L is invariably rank deficient. L is rank deficient because we are trying to determine $\mathbf{R}^{(k)}$ using easily obtainable input spectra, and these tend to be of relatively low dimensionality. If L were of full rank, then we would have a method analogous to the monochrometer method. The number of independent spectra needed for methods based on matrix inversion is, of course, a function of the number of samples for which we wish to solve. Wyszecki[10] reported results using matrix inversion on a similar problem where the rank of L and the number of samples were explicitly matched. In the more general case, where L is rank deficient, results based on matrix inversion (or pseudo-inversion) are very sensitive to noise because it is mainly the noise that is being fitted, and the resulting sensor responses

tend to have numerous large spikes and an abundance of nonnegligible negative values [Fig. 3(a)].

Sharma and Trussell[4,7] improved the prospects for a reasonable solution by introducing various constraints on $\mathbf{R}^{(k)}$. First, instead of solving Eq. (3) exactly, they constrained the maximum allowable error as well as the RMS error. In addition, they constrained a discrete approximation of the second derivative to promote a smooth solution. Finally, they constrained the response functions to be positive. They then observed that the constraint sets were all convex, and so they computed a resulting constraint set using the method of projection onto convex sets.

Hubel *et al.*[11] also recognized that some form of smoothness was necessary for a good solution, and they investigated the Wiener estimation method, as described by Pratt and Mancill,[12] as a method for finding a smooth fit. They found that generally the method produced good results. They noted, however, that the method produced negative lobes in the response functions and briefly mentioned using the projection-onto-convex-sets method to remedy this problem.

Sharma and Trussell's contribution was the starting point for some of our own work on this problem.[6] Rather than constrain the absolute RMS error, we chose instead to minimize the relative RMS error. We then rewrote Sharma and Trussell's other constraints so that the entire problem became a least-squares fit with linear constraints, for which there are standard numerical methods readily available. Once we had a fit for our camera sensors, we noted that they were essentially unimodal and that once the sensors dropped to a small value they remained small. On these grounds we also constrained the sensors to be 0 outside a certain range on subsequent runs. In this particular case this forced the sensors response functions to be unimodal. This last step needs to be applied with care, as it is possible that the sensors are in fact nonzero beyond the points where the main peaks drop to 0.

Recently, Finlayson *et al.* used a similar approach.[13] They constrained smoothness by restricting the sensors to be linear combinations of the first 9–15 Fourier basis functions. They also introduced a modality constraint expressed by peak location and reported results constraining the sensors to be unimodal and bimodal. They determined the best fit for each proposed modality by stepping through all possible peak locations. This method also requires care, as the modality is often unknown. This method makes sense when used in conjunction with Fourier smoothing, as that method can introduce spurious peaks. As part of this work we have implemented the modality constraint and Fourier-based smoothing.

THE CHARACTERIZATION APPROACH

Our proposed characterization method is described in two stages. First, we will explain the basic method, which estimates the response vector for each channel $\mathbf{R}^{(k)}$ on the assumption that the linearization function, $F^{(k)}$, has been found and applied. Second, we incorporate the estimation of

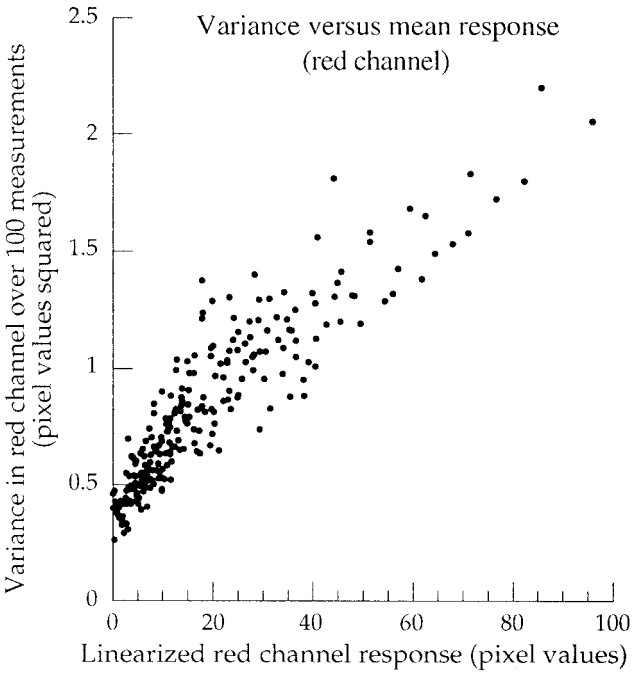

FIG. 1. The variance of red channel measurements versus their values.

$F^{(k)}$ into the fitting procedure. This has the advantage that the error in the two fits can be traded off against each other, and data collected to find $\mathbf{R}^{(k)}$ can be exploited to estimate $F^{(k)}$ more accurately.

In our initial work[6] we minimized the relative RMS error in Eq. (3) subject to positivity constraints, smoothness constraints, and a constraint on the maximum allowable error (and/or relative error). We have since found that it is better to replace the constraint on smoothness with a regularization term added to the objective function. Thus, we were able to minimize the relative error and the nonsmoothness measure together. This allows fitting the error to be traded against nonsmoothness and vice versa. With the hard constraint used previously, there is no recourse in the case in which making the sensor response slightly less smooth at a particular location could substantially reduce the error. Similarly, there is no recourse when a small increase in error beyond the hard limit could substantially increase the smoothness. These observations also apply to the Fourier smoothing approach.

We minimized the relative RMS error for two reasons. First, as discussed in more detail below, the variance of the pixels values increases with their magnitude. And, second, we have found that minimizing relative error better reduces the error in chromaticity, which is difficult to minimize directly but is often of most interest. However, for some applications it may make more sense to minimize absolute error or even to use a weighted combination of both. We have also found that it is not generally necessary to use Sharma and Trussell's constraint on maximum allowable error to get good results, but, again, limiting either the absolute error or the relative error may be called for in some cases and is easily added to the method as detailed below.

To investigate the nature of the error in our pixel values, we made 100 consecutive measurements of the Macbeth ColorChecker® illuminated at 10 intensities. A pixel was chosen for each of the 24 patches, and the mean and the variance was computed for each of the 240 pixel/intensity combinations over the 100 measurements. The means were linearized by the method described later in this article. The results for the red channel are plotted in Figure 1.

We considered the variance to be the sum of the intensity-dependent and the intensity-independent parts. We further assumed that the intensity-dependent variance to be a result of photoelectron shot noise and thus to be proportional to the mean.[14] Thus, we expect that the observed variance would be linear with the intensity, with an offset characterizing the noise from other sources. This is more or less consistent with the graph in Figure 1, but the spread of values indicates that this model is perhaps overly simplified.

For the case where $F^{(k)}$ has already been found and applied to the data, beginning with the formulation that minimizes absolute error, the details are as follows. Let N be the number of spectral samples used. First, the N-2 was formed by the N second derivative matrix S:

$$S = \begin{vmatrix} -1 & 2 & -1 & & & & \\ & -1 & 2 & -1 & & & \\ & & \cdot & \cdot & \cdot & & \\ & & \cdot & \cdot & \cdot & & \\ & & & -1 & 2 & -1 & \\ & & & & -1 & 2 & -1 \end{vmatrix} \quad (4)$$

and then we solved

$$\left|\frac{L}{\lambda S}\right| \mathbf{R}^{(k)} = \left|\frac{\mathbf{r}^{(k)}}{0}\right| \quad (5)$$

in the least-squares sense, subject to linear constraints. This is equivalent to minimizing the objective function

$$\sum_i (\mathbf{L}_i \cdot \mathbf{R}^{(k)} - \rho_i^{(k)})^2 + \lambda \sum_i (\mathbf{S}_i \cdot \mathbf{R}^{(k)})^2 \quad (6)$$

where the first term expresses the error and the second term expresses nonsmoothness and thus provides the regularization. Coefficient λ specifies the relative weight attributed to the two terms. If λ is 0 and there are no constraints, then this becomes the pseudo inverse method. A serviceable value for λ is easily found by trial and error. To ensure positivity, we used the constraint

$$\mathbf{R}^{(k)} \geq \mathbf{0} \quad (7)$$

To specify that the sensor response is 0 outside the range [min, max], we can add the constraint:

$$R_i^{(k)} \leq 0 \quad \text{for} \quad i < \min, i > \max \quad (8)$$

To specify that the absolute error is no more than a specified positive value, δ, we can add the constraint:

$$\rho_i^{(k)} - \delta \leq \mathbf{L}_i \cdot \mathbf{R}^{(k)} \leq \rho_i^{(k)} + \delta \quad (9)$$

This condition was not used for any of the results in this article, but it may be interest and can be used to more closely emulate the method of Sharma and Trussell.[4,7]

To minimize the relative error we need to replace Eq. (6) by:

$$\sum_i \left(\frac{\mathbf{L}_i \cdot \mathbf{R}^{(k)} - \rho_i^{(k)}}{\rho_i^{(k)}}\right)^2 + \lambda \sum_i (\mathbf{S}_i \cdot \mathbf{R}^{(k)})^2 = \sum_i \left(\frac{\mathbf{L}_i \cdot \mathbf{R}^{(k)}}{\rho_i^{(k)}} - 1\right)^2 + \lambda \sum_i (\mathbf{S}_i \cdot \mathbf{R}^{(k)})^2 \quad (10)$$

One way to express this is to use a modified version of L, L_{rel}, which is simply the rows of L divided by the corresponding sensor response. Formally, L_{ref} is given by:

$$L_{\text{rel}} = (\text{diag}(\mathbf{r}^{(k)}))^{-1} \cdot L \quad (11)$$

We then replace Eq. (5) with:

$$\left|\frac{L_{\text{rel}}}{\lambda S}\right| \mathbf{R}^{(k)} = \left|\frac{1}{0}\right| \quad (12)$$

Similar to the constraint of Eq. (9), if we require a constraint limiting the relative error to less than a positive amount, ζ, we can use:

$$1 - \zeta \leq L_{\text{rel}} \mathbf{R}^{(k)} \leq 1 + \zeta \quad (13)$$

where the inequalities are applied to each component of the vector $L_{\text{rel}}\mathbf{R}^{(k)}$. As with Eq. (9), this constraint was not used for any results reported in this article.

Note that minimizing the relative error may need to be modified slightly to deal with a very small $\rho^{(k)}$. Such data is likely to be inaccurate for a variety of reasons. Thus, we need to either ignore small values or give the corresponding data row less weight in the fitting process. Eq. (12) can be interpreted as a weighted version of Eq. (5), with the weighting inversely proportional to $\rho^{(k)}$. Thus, it is natural and easy to put an upper bound on this weighting to safeguard against a small $\rho^{(k)}$ when excluding them outright is not desired. These considerations are not relevant for the experiments with real data reported below, as there was no data with small values of $\rho^{(k)}$ because of the relatively large response of our camera to no light [specifically (11.05, 13.06, 12.36)].

Following Finlayson *et al.*,[13] we can constrain the sensors to be unimodal given a specification of the peak location $i_{\max}^{(k)}$ with the following constraints:

$$R_{i-1}^{(k)} \leq R_i^{(k)} \quad \text{for} \quad i \leq i_{\max}^{(k)} \quad \text{and} \quad R_i^{(k)} \geq R_{i+1}^{(k)} \quad \text{for} \quad i \geq i_{\max}^{(k)} \quad (14)$$

The procedure for multiple peaks is similar. Of course, because the peak locations are not known, this method requires trying all possible peak locations and choosing the peak locations that produce the least error.

Finlayson *et al.*[13] also proposed ensuring smoothness by restricting the sensor response functions to being linear combinations of Fourier basis functions. To implement this approach, we formed a matrix B whose rows are the first D Fourier basis functions, with one period coinciding with the

wavelength range used. Then the sensor functions can be expressed as:

$$\mathbf{R}^{(k)} = B\mathbf{a}^{(k)} \tag{15}$$

where all the $\mathbf{a}^{(k)}$ are the Fourier coefficients. Finlayson *et al.* substituted Eq. (15) into all relevant equations and then used $\mathbf{a}^{(k)}$ as the unknowns in their quadratic programming problem. This has a small advantage in reducing the size of the problem, and we used this method for the results on synthesized data. Unfortunately, using the Fourier basis constraint in this form is problematic when used in conjunction with the linearization fitting described shortly. Thus, to provide results with the linearization fitting, we used a different form of the constraint. Using the orthogonality of the rows of B

$$\mathbf{a}^{(k)} = B^T\mathbf{R}^{(k)} \tag{16}$$

then

$$\mathbf{R}^{(k)} = BB^T\mathbf{R}^{(k)} \tag{17}$$

and

$$(BB^T - 1)\mathbf{R}^{(k)} = \mathbf{0} \tag{18}$$

For the experiments on real data, we used Eq. (18) as equality constraints on the least-squares minimization problem. Alternatively, it is also possible to use Eq. (18) in place of the regularization rows of Eqs. (5) or (12). If Eq. (18) is used in this manner, then the adherence to the Fourier smoothing constraint increases with increasing values of λ.

The methods developed so far assume that the function $F^{(k)}$ has been found and applied as a preliminary step. However, the body of data collected to find $\mathbf{R}^{(k)}$ also contains information about $F^{(k)}$, and because this data set needs to be comprehensive, it makes sense to use it for the final determination of $F^{(k)}$. Therefore, we propose fitting $\mathbf{R}^{(k)}$ and $F^{(k)}$ together. This has the advantage that fitting errors from $F^{(k)}$ and $\mathbf{R}^{(k)}$ can be traded against each other. We first made a rough estimate of $F^{(k)}$, which we used to propose a parameterized expression for it. We then fit the parameters for $F^{(k)}$ and $\mathbf{R}^{(k)}$ simultaneously. Next, we provide a specific example of such a strategy.

The Sony DXC-930 camera that we used for our experiments is quite linear for most of its range, provided it is used with gamma disabled. However, in all three channels it has a substantial response to no light (camera black) as well as a slight nonlinearity for small pixel values. Because of this nonlinearity, a line fitted to the linear part does not intersect the response axis at the camera black, and simply linearizing the camera by subtracting the camera black leads to errors in chromaticity. Therefore, the nonlinearity must be taken into account, even if it is not explicitly estimated. Figure 2 shows the slight nonlinearity for the red channel. The other two channels are similar.

The fit shown in Figure 2 was found using a simple linear fit [see Eq. (20) below] with pixel values greater than 30. This shows that linearization information is available in the data set to be used to find $\mathbf{R}^{(k)}$. To proceed with our strategy

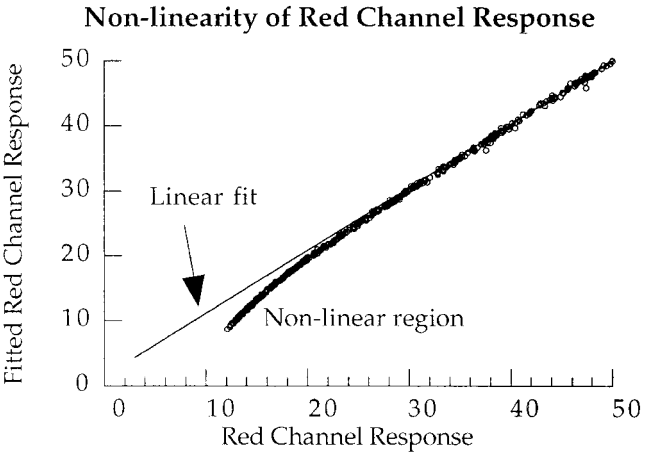

FIG. 2. The nonlinearity of the red channel response for the Sony DXC-930 camera used for the sensor fitting experiments. The fit shown is a simple linear fit on pixel values greater than 30. The other two channels have similar curves.

of explicitly fitting the nonlinearity, we need to parameterize it. The particular form of the parameterization is somewhat arbitrary and will vary substantially from case to case. With a little experimentation we found that the nonlinearity of our camera could be approximated by:

$$F^{(k)}(x) = x - a_0^{(k)} - a_1^{(k)}e^{-C_k(x-b_k)} \tag{19}$$

where b_k is the camera black for channel k and C_k is a constant that must be found by trial and error but was found to be quite stable. If we use the simpler form:

$$F^{(k)}(x) = x - a_0^{(k)} \tag{20}$$

then we would simply be fitting a camera offset simultaneously with $\mathbf{R}^{(k)}$. This would be a reasonable approach for our camera if we did not wish to use smaller pixel values. In general, the parameters of the approximation function must generate a reasonable collection of response functions that roughly fit the nonlinearity so that the overall fitting procedure can find a good estimate of $F^{(k)}(x)$. In addition, the parameters that are fitted must be linear coefficients. For example, we can only directly fit for $a_0^{(k)}$ and $a_1^{(k)}$; C_k must be found by trial and error.

To find the parameters for the approximation of $F^{(k)}(x)$ simultaneously with $\mathbf{R}^{(k)}$ when fitting for absolute error, we replace Eq. (5) with:

$$\left|\begin{array}{c|cc} L & \mathbf{1} & e^{-C_k(\mathbf{r}^{(k)}-b_k)} \\ \hline \lambda S & 0 & 0 \end{array}\right| \cdot \left|\begin{array}{c} \mathbf{R}^{(k)} \\ a_0^{(k)} \\ a_1^{(k)} \end{array}\right| = \left|\begin{array}{c} \mathbf{r}^{(k)} \\ \hline \mathbf{0} \end{array}\right| \tag{21}$$

where the arithmetic in the upper right block of the matrix is done elementwise as needed. Similarly, in the case of fitting for relative error, we replace Eq. (12) with:

$$\left|\begin{array}{c|cc} L_{\text{rel}} & \dfrac{\mathbf{1}}{\mathbf{r}^{(k)}} & \dfrac{e^{-C_k(\mathbf{r}^{(k)}-b_k)}}{\mathbf{r}^{(k)}} \\ \hline \lambda S & 0 & 0 \end{array}\right| \cdot \left|\begin{array}{c} \mathbf{R}^{(k)} \\ a_0^{(k)} \\ a_1^{(k)} \end{array}\right| = \left|\begin{array}{c} \dfrac{\mathbf{1}}{\mathbf{r}^{(k)}} \\ \hline \mathbf{0} \end{array}\right| \tag{22}$$

where, again, the arithmetic in the upper right block of the

TABLE I. Results of fitting experiments on generated data. The first error measure is the relative error (RMS) over the three sensor functions for the separate channels. This error was minimized by the fitting process, and therefore the addition of constraints always leads to an increase in error. the error measure in the second column is an estimate of how well the fitting process estimated the actual sensor response functions used to generate the data. The maximum value of this error measure is 1, which is very nearly reached with the pseudo-inverse method.

Fitting method	Average of relative error over the 3 sensor-response functions (%)	RMS difference between fitted sensor curve and target, normalized by maximum of the fit and target norms, averaged over the 3 sensors
Pseudoinverse	4.158	0.9997
Pseudoinverse with positivity	4.756	0.8157
Pseudoinverse with positivity and modality	4.793	0.2412
Fitting with positivity and smoothing	4.812	0.0495
Fitting with positivity, smoothing, and range	4.820	0.0424
Fitting with positivity, smoothing, and modality	4.815	0.0520
11 Fourier basis functions with positivity	8.445	0.1826
13 Fourier basis functions with positivity	5.788	0.1303
15 Fourier basis functions with positivity	5.160	0.1172
18 Fourier basis functions with positivity	5.021	0.1056
21 Fourier basis functions with positivity	4.876	0.0712
25 Fourier basis functions with positivity	4.821	0.0923
31 Fourier basis functions with positivity	4.807	0.0921
39 Fourier basis functions with positivity	4.796	0.2357
11 Fourier basis functions with positivity and modality	25.558	0.3967
13 Fourier basis functions with positivity and modality	17.063	0.2835
15 Fourier basis functions with positivity and modality	10.137	0.1938
18 Fourier basis functions with positivity and modality	5.973	0.1514
21 Fourier basis functions with positivity and modality	5.193	0.1413
25 Fourier basis functions with positivity and modality	5.000	0.0911
31 Fourier basis functions with positivity and modality	4.868	0.1019
39 Fourier basis functions with positivity and modality	4.814	0.0982
50 Fourier basis functions with positivity and modality	4.804	0.1030
75 Fourier basis functions with positivity and modality	4.798	0.1367
100 Fourier basis functions with positivity and modality	4.793	0.2280

matrix is done elementwise as needed. Note that the response vectors $\mathbf{r}^{(k)}$ now correspond to the observed camera responses, $\upsilon^{(k)}$, in Eq. (2), in contrast to the earlier formulation, where $\mathbf{r}^{(k)}$ corresponded to the linearized camera responses, $\rho^{(k)}$.

In all cases the entire fitting procedure is a least-squares minimization problem with linear constraints, or equivalently, it can be viewed as a quadratic programming problem. Such problems can be solved with standard numerical techniques for which software is readily available. We used the freely available SLATEC Fortran library routine DLSEI. The routine DBOCLS in that library may also be used. A third option is the Matlab routine "qp."

EXPERIMENTS WITH SYNTHETIC DATA

Experiments with synthetic data are useful because the sensor functions sought are known. For these experiments we used a linear camera model with sensors similar to the real camera sensors determined in the next section. Given the sensors and the set of 598 input spectra used in the real calibration experiment, we synthesized responses using Eq. (3). To all responses we added 5% relative Gaussian noise. Under these conditions it should be easy to obtain a relative fitting error of roughly 5%. However, some methods, such as the pseudo-inverse method, are expected to overfit the characteristics of the specific input data set. Thus, a more interesting error measure is the difference between the actual sensors and the computed ones. Because the actual sensors are relatively smooth and nonnegative, we expect methods that promote these characteristics to do better.

The results are shown in Table I, and the sensors obtained using a selection of the methods plotted together with the actual sensors are shown in Figure 3. As predicted, the lowest fitting error was obtained using the pseudo-inverse, as it is the least constrained method, but the resulting sensor response functions are very poor. The results show that adding constraints for positivity and the regularization equations for smoothness do not overly increase the fitting error but significantly reduce the error in the sensor response functions. The best match of the sensor response functions was obtained by in addition using the range constraints. But it note that human input was used in deciding which the limits to use and that this method will be of less use when the nature of the sensors is more in question.

Fourier smoothing proved to be less effective that the simple regularization approach proposed in this work. The results shown in Table I indicate there is no choice in the number of basis functions that produce results comparable to our approach. Fourier smoothing puts constraints on the sensor functions, which are not necessary for simple smoothness, and many good candidates for the sensor response functions cannot be considered. In contrast, the approach proposed in this work allows the degree of

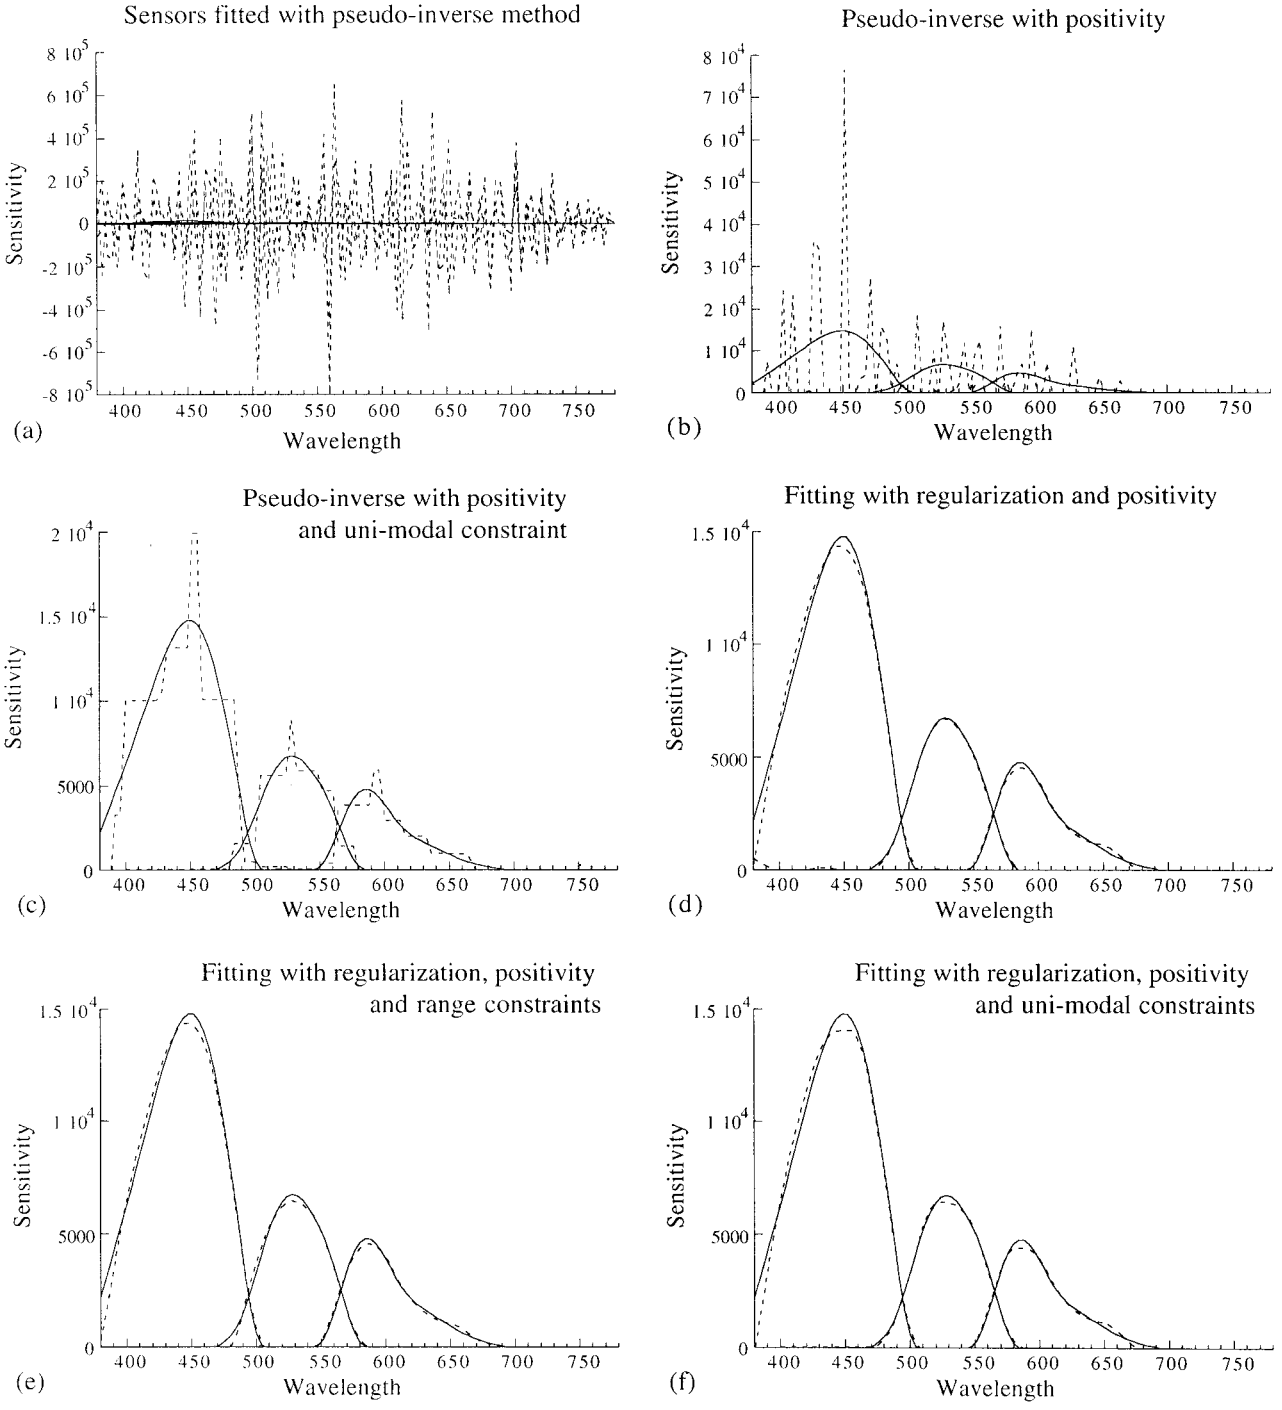

FIG. 3. The results of various fitting methods on synthetic data. The data were generated from idealized sensors based on the actual sensors of our Sony DXC-930 camera; 5% relative Gaussian noise was added. The results of various fitting methods that promote smooth results by constraining the solution as a linear combination of a specified number of Fourier basis functions.

smoothness to be traded against the fitting error, which yields more flexible fitting.

EXPERIMENTS WITH REAL DATA

We investigated camera characterization for a Sony DXC-930 3-chip CCD camera. To obtain a comprehensive set of calibration data, we automated the collection of input energy spectra and the corresponding camera responses. Our target was a Macbeth ColorChecker®, which has 24 different colored patches that we illuminated with a number of illuminant/filter combinations. The black patch of the chart was not used because it did not reflect enough light with the darker illuminants for reliable spectroradiometer measurements. The main criterion for the apparatus was to ensure that the camera and the spectroradiometer measured the

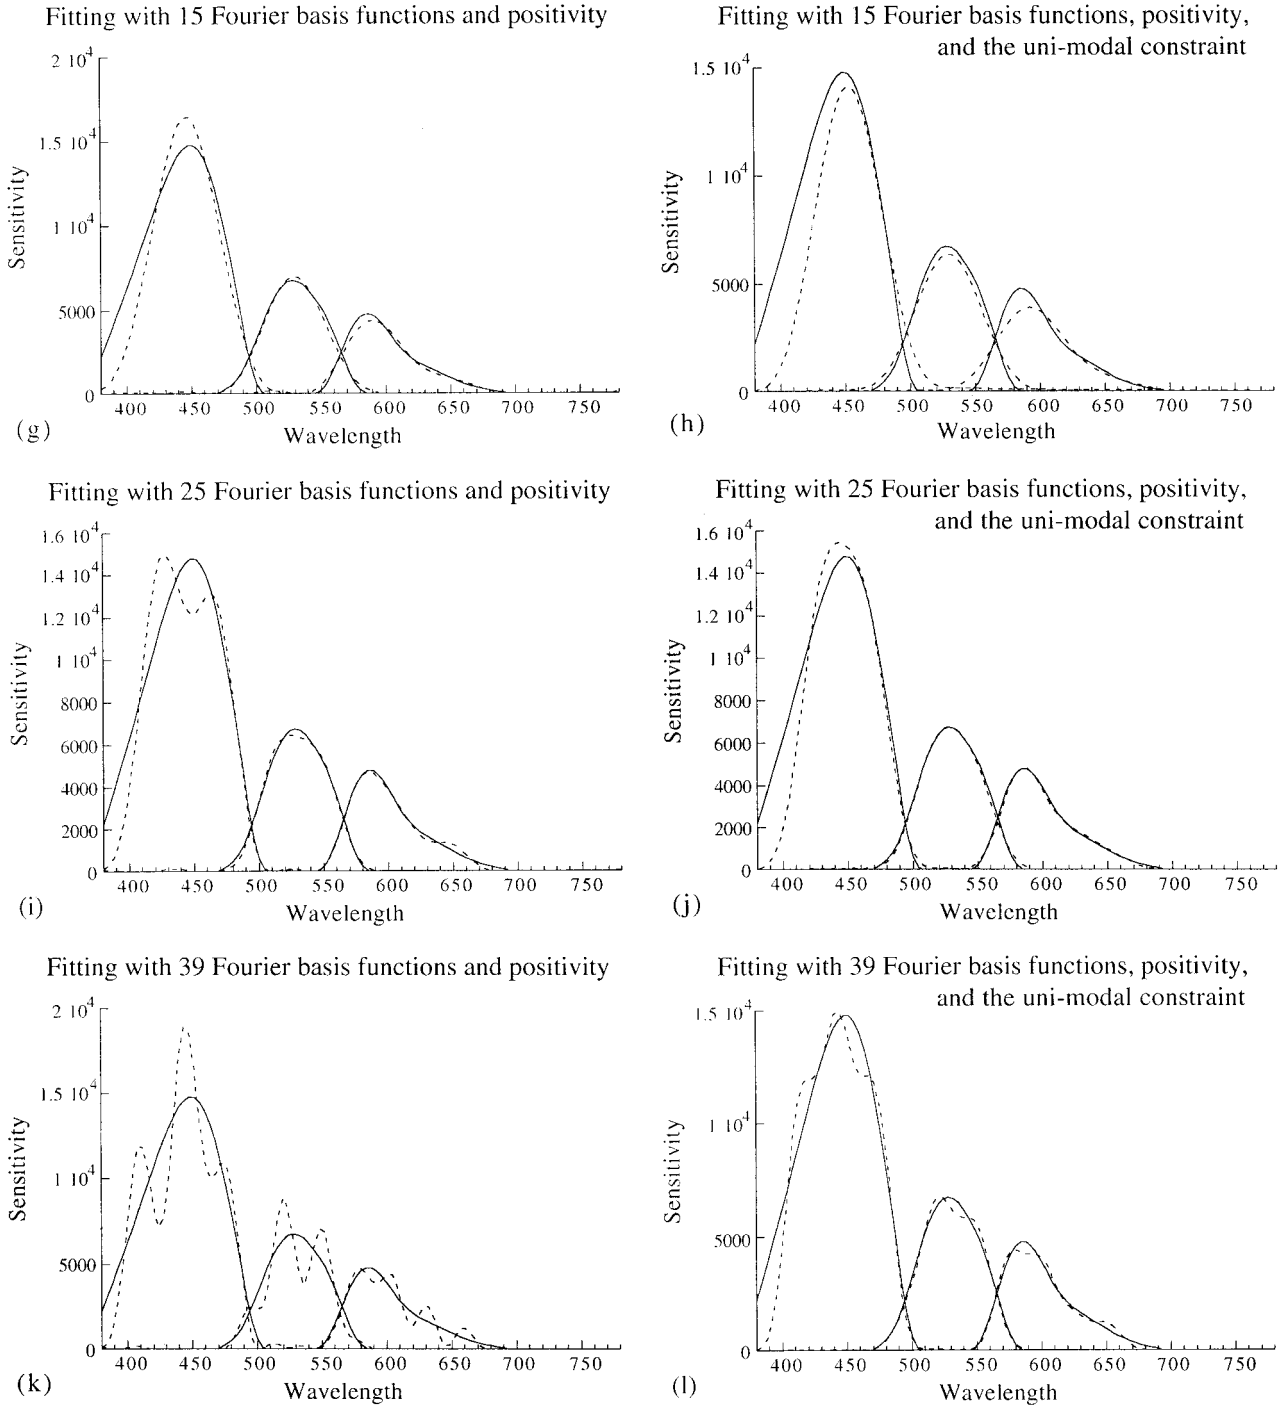

FIG. 3. (*Continued*)

same signal. We also required that the camera data was always measuring the center of the image. Therefore, we mounted the color checker horizontally on an *XY* table that moved it under computer control. The camera and the spectroradiometer were mounted on the same tripod, with their common height controlled with the tripod head-height adjustment mechanism. Rather than aim them simultaneously at the target, we decided instead to set the optical axes to be parallel. This meant that the tripod head had to be raised and lowered between capturing camera data and spectroradiometer data. Thus, we captured an entire chart's worth of camera data before capturing an entire chart's worth of spectra. Twenty-six illuminant/filter combinations were used in conjunction with the 23 patches, providing 598 measurements (available on-line[3]). For all fitting experiments we excluded response values exceeding 240.

We took additional steps to obtain clean data. As indicated above, it is important that the camera and the spectroradiometer are exposed to the same signal. To minimize the effect of misalignment, we made the illumination as uniform as possible. To reduce the effect of flare, the target was imaged through a hole in a black piece of cardboard,

TABLE II. Results of fitting experiments on data captured as explained in the text. Each linear-based fitting method was used in conjunction with three linearity fitting methods.

Linear fitting method	Relative RGB error with linearity fitting limited to subtraction of camera black	Relative RGB error with fitting of camera linearity intercept	Relative RGB error with full linearity fitting
Pseudoinverse	0.0262	0.0240	0.0095
Pseudoinverse with positivity	0.0427	0.0303	0.0107
Pseudoinverse with positivity and modality	0.0434	0.0309	0.0115
Fitting with positivity and smoothing	0.0448	0.0316	0.0117
Fitting with positivity, smoothing, and range	0.0448	0.0322	0.0146
Fitting with positivity, smoothing, and modality	0.0447	0.0317	0.0123
11 Fourier bases with positivity	0.1069	0.0720	0.0606
13 Fourier bases with positivity	0.0699	0.0443	0.0297
15 Fourier bases with positivity	0.0560	0.0360	0.0189
18 Fourier bases with positivity	0.0490	0.0335	0.0166
21 Fourier bases with positivity	0.0473	0.0324	0.0137
25 Fourier bases with positivity	0.0458	0.0317	0.0120
31 Fourier bases with positivity	0.0447	0.0312	0.0116
39 Fourier bases with positivity	0.0439	0.0309	0.0112
11 Fourier bases with positivity and modality	0.2626	0.2402	0.5775
13 Fourier bases with positivity and modality	0.1878	0.1586	0.4253
15 Fourier bases with positivity and modality	0.1215	0.0907	0.0806
18 Fourier bases with positivity and modality	0.0675	0.0458	0.0315
21 Fourier bases with positivity and modality	0.0497	0.0344	0.0196
25 Fourier bases with positivity and modality	0.0474	0.0333	0.0168
31 Fourier bases with positivity and modality	0.0454	0.0320	0.0138
39 Fourier bases with positivity and modality	0.0449	0.0315	0.0123

exposing the region of interest but as little else as was practical. We extracted a 30-by-30 window from the image that corresponded as closely as possible to the area used by the spectroradiometer. The 8-bit RGB values of the pixels in this window were averaged. Finally, the camera measurements were averaged over 50 frames to further reduce the effect of photon shot noise, and the spectroradiometer measurements were averaged over 20 capture cycles.

We considered three approaches to linearization. The most naive method is simply to subtract the camera black from the data but otherwise assume the data is linear. The second method makes the assumption that the camera is linear except at the two extremes. Thus, the intercept of the linear is made to fit a parameter of the fitting process [as in Eq. (20)]. Because of the curvature evident in Figure 2, this method results in the subtraction of an offset that is somewhat less than the camera black used in the first approach. Finally, in the third method the results of parameterizing the nonlinearity are provided, as developed above.

Each linearization method was used in conjunction with a number of methods for fitting the camera response functions. The compiled results are shown in Table II, all of which are based on minimizing the relative error. The regularization smoothing parameter, λ, was initially set by trial and error to a value that provided a reasonably smooth sensor functions. We did not attempt to tune λ beyond a factor of 2, and we used the same value for all variants. For the Fourier smoothing method we provide results for a wide range of choices for the number of basis functions.

The results in Table II show that for our camera, fitting for the linearity in conjunction with the sensor functions substantially reduced the error compared to both subtraction of camera offset and fitting for the intercept. Although we expected some benefit, the extent of the improvement was beyond what we expected, as our camera is actually quite linear.

A comparison for one of the methods of a fitting based on

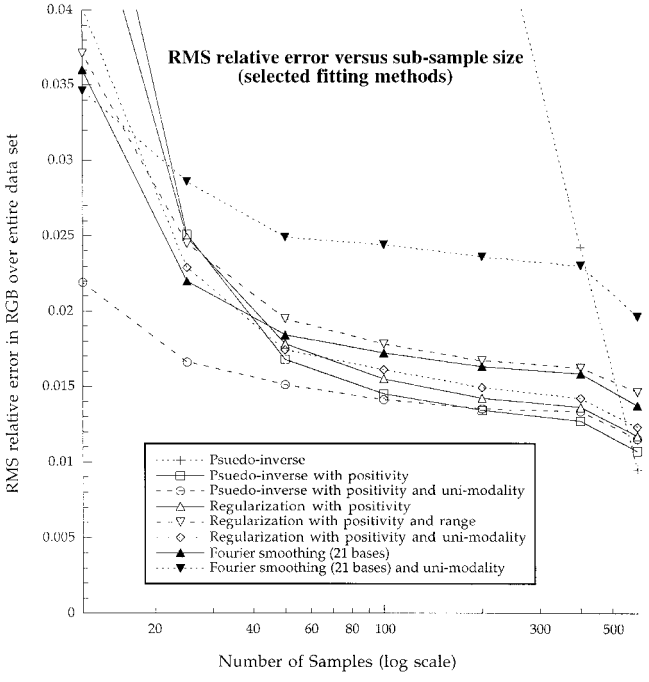

FIG. 4. RMS relative error of various fitting methods versus sample size. The more constrained methods tend to more robust as the number of sample points decreases, but have higher error when the number of points is large. For example, the unconstrained pseudo-inverse method (first curve) has the least error when found using the full data set but degraded rapidly with decreasing sample size. As the number of data points decreased, we became more reliant on prior conception of the shape of the functions.

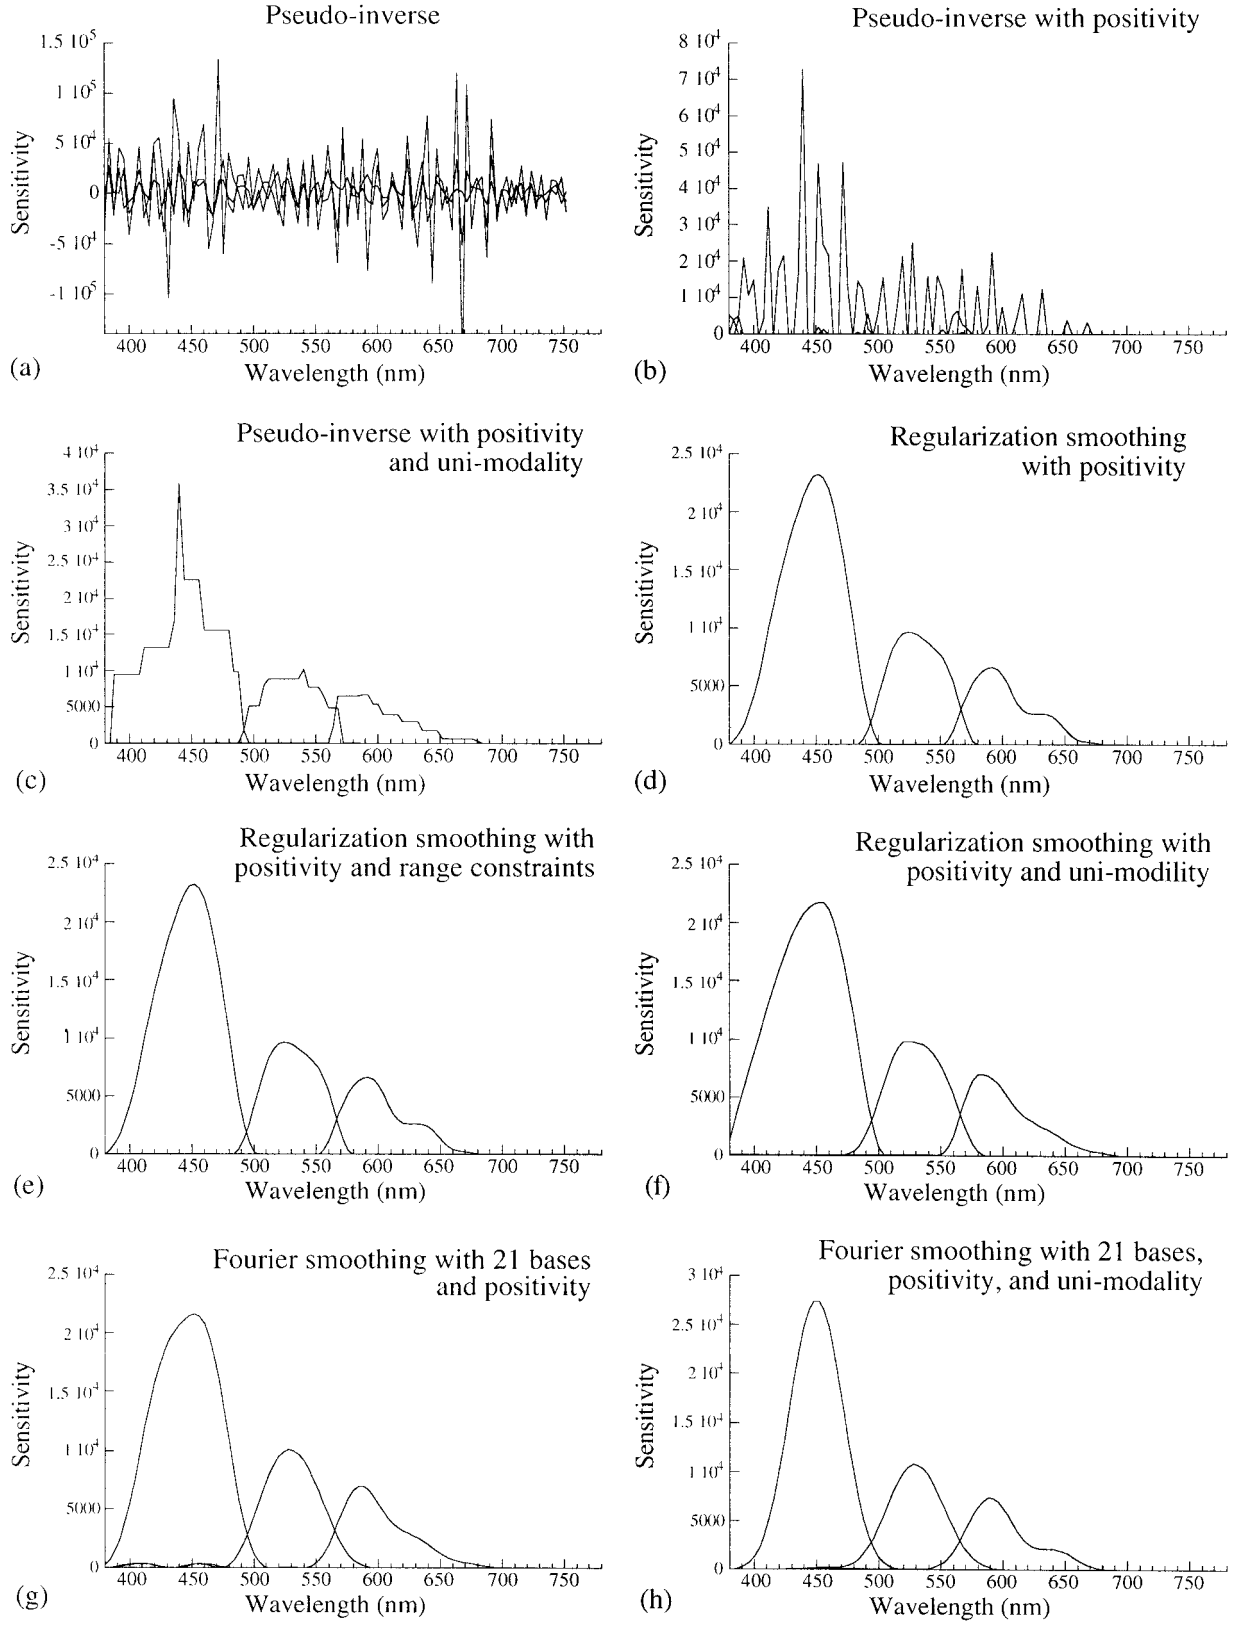

FIG. 5. Sensor response functions found using a variety of fitting of the data collected, as described in the text. All methods were used in conjunction with linearity fitting. These sensors correspond to the results in Figure 4, for the full sample size of 598 points.

absolute error with a fitting based on relative error is shown in Table III. Not surprisingly, when the data were fit using relative error, the relative error was lower and the absolute error higher than when the data were fit using absolute error. More significantly, fitting with relative error substantially reduced the absolute error in (r, g) chromaticity, which is

TABLE III. A comparison of fitting based on relative error with fitting based on absolute error as one of the preferred methods (positivity, smoothing, and unimodality). Linearization is fitted simultaneously with the sensor response functions as in the third column of Table II.

Error minimized	RMS relative RGB error	RMS absolute RGB error	RMS absolute error in $r = R/(R + G + B)$	RMS absolute error in $g = G/(R + G + B)$	RMS $L*a*b$ error
Absolute	0.0141	0.79	0.0052	0.0056	0.303
Relative	0.0123	0.89	0.0027	0.0044	0.285

difficult to minimize directly and is key for many applications. Table III also includes $L*a*b$ error for the two objective functions.

Unlike with the synthetic case, true camera sensor functions are not known. Thus, to investigate the robustness of the fitting methods, we determined the camera model using subsets of the data and computed how well the sensor responses for the entire data set were predicted. We used subsets of sizes 400, 200, 100, 50, 25, and 12, as well as the full data set (598 data points). We averaged the results over 100 random selections of the above subset sizes. Each data subset was augmented with the data for no light. Each fitting method was used in conjunction with the linearization method developed above. In this experiment we restricted our attention to 21 basis functions for the Fourier smoothing method. Figure 4 shows the results for each fitting method plotted against the number of data sample points on a log scale. Figure 5 shows the sensor response functions corresponding to each of these methods when all the data were used.

When the full data set was used for fitting, adding constraints invariably increased the error, as expected. However, as the number of points used for fitting decreased, the more constrained methods proved to be more robust. This was best illustrated by the pseudo-inverse method. It had the least error when fitted using the entire data set (which is exactly the test data set), but its performance deteriorated very rapidly when its parameters were determined using smaller and smaller subsets of the data. Adding positivity led to a big improvement in robustness for all the methods (for simplicity only the pseudo-inverse method is shown without positivity). As the number of points in the subset became very small, the unimodality constraint became increasingly useful. Of course, a small data set could not be used to determine with confidence that sensors are in fact unimodal. Interestingly, the pseudo-inverse method with positivity and unimodality produced surprisingly low error, even though Figure 5 and the synthetic experiments suggest that the corresponding sensors are not likely to be close to the real sensors. This result reflects the relatively low dimensionality of the input spectra relative to the 101 samples provided by the spectroradiometer. The constraint on the range of the sensors also added robustness as evaluated by the deterioration in performance when the sample size is small. However, as in the case of the modality constraint, when the amount of data is small, it is hard to be confident in the range, and therefore it can only be used if it is already available.

CONCLUSIONS

We have developed and tested a new method for fitting a common camera model used in color research. By promoting smoothness and using constraints on the sensor response functions such as positivity, we obtained a result that is both reasonable and robust. We have found that it is best to promote smoothness by adding a regularization term to the minimization expression rather than constraining it, as has been done in earlier work by others and ourselves. We also investigated fitting a small nonlinearity in the camera response simultaneously with the sensor response functions. This was effective because errors from the lack of fit of the two model parts can be traded against each other for a better overall characterization. This approach also takes advantage of the linearization information inherent in the data required to determine the sensitivity functions. Finally, our experiments support the hypothesis that it can be preferable to minimize the relative error, especially if chromaticity accuracy is more important than overall accuracy.

ACKNOWLEDGMENTS

We are grateful for the financial support of Hewlett-Packard Corporation and the Natural Sciences and Engineering Council of Canada. In addition, we acknowledge the efforts of Lindsay Martin, who helped greatly with the data collection.

1. Horn BKP. Robot Vision. Cambridge (MA): MIT Press; 1986.
2. Vora PL, Farrell JE, Tietz JD, Brainard DH. Digital color cameras. 1. Response models: technical report [Internet]. Hewlett-Packard Laboratory; 1997. Contract No: HPL-97-53. Available from: http://www.hpl.hp.com/techreports/97/HPL-97-53.html.
3. Barnard K, Martin L, Funt B, Coath A. Data for colour research [Internet]. Available from: http://www.cs.sfu.ca/~colour/data.
4. Sharma G, Trussell HJ. Characterization of scanner sensitivity. Proceedings of the IS&T and SID's Color Imaging Conference: Transforms & Transportability of Color. Springfield (VA): The Society for Imaging Science and Technology; 1993; p. 103–107.
5. Healey GE, Kondepudy R. Radiometric CCD camera calibration and noise estimation. IEEE Pattern Anal Mach Intell 1994;16:267–276.
6. Barnard K. Computational colour constancy: Taking theory into practice [MSc thesis]. Burnaby [British Columbia, Canada]: Simon Fraser University; 1995. Available from: ftp://fas.sfu.ca/pub/cs/theses/1995/KobusBarnardMSc.ps.gz.
7. Sharma G, Trussell HJ. Set theoretic estimation in color scanner characterization. J Elec Imag 1996;5:479–489.
8. Vora PL, Farrell JE, Tietz JD, Brainard DH. Digital color cameras. 2. Spectral response; Technical report [Internet]. Hewlett-Packard Lab-

oratory; Contract No: HPL-97-54, Available from: http://www.hpl.hp.com/techreports/97/HPL-97-54.html.
9. Barnard K. Practical colour constancy [PhD Thesis]. Burnaby [British Columbia, Canada]; Simon Fraser University School of Computing Science; 1999. Available from: ftp://fas.sfu.ca/pub/cs/theses/1999/KobusBarnardPhD.ps.gz.
10. Wyszecki G. Multifilter method for determining relative spectral sensitivity functions of photoelectric detectors. J Opt Soc Am 1960;50: 992–998.
11. Hubel PM, Sherman D, Farrell JE. A comparison of method of sensor spectral sensitivity estimation. Proceedings of the IS&T/SID 2nd Color Imaging Conference: Color Science, Systems, and Applications. The Society for Imaging Science and Technology; Springfield (VA): 1994; p 45–48.
12. Prat WK, Mancill CE. Spectral estimation techniques for the spectral calibration of a color image scanner. Appl Opt 1976;15:73–75.
13. Finlayson G, Hordley S, Hubel P. Recovering device sensitivities with quadratic programming. The Proceedings of the IS&T/SID Sixth Color Imaging Conference: Color Science, Systems and Applications, Springfield (VA): The Society for Imaging Science and Technology; 1998. p 90–95.
14. Holst GC. CCD Arrays, cameras, and displays. 2nd ed. Bellingham (WA): SPIE Press; 1998.