

## **Kobus Barnard\***

Computer Division  
University of California, Berkeley  
Berkeley, CA 94720-1776

## **Brian Funt**

Department of Computing Science  
Simon Fraser University  
8888 University Drive  
Burnaby, BC, Canada, V5A 1S6

This is a preprint of an article accepted for publication in Color Research and Application © copyright 2001 (John Wiley & Sons, Inc.)

(Preprint URL is  
[http://www.CS.Berkeley.EDU/~kobus/research/publications/camera\\_characterization/index.html](http://www.CS.Berkeley.EDU/~kobus/research/publications/camera_characterization/index.html))

# **Camera characterization for color research**

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 non-linearities. 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 on-line at <http://www.cs.sfu.ca/~colour/data>.

## **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. Since 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.

---

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

We first introduce the standard camera model used in color-oriented computer vision<sup>1,2</sup>. We then 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 the 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 fact that the data required to estimate camera sensitivity also contains linearization information. Furthermore, it is beneficial to allow the errors due to the inadequacy of the linearization fit to be traded against errors due to 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 is available on-line<sup>3</sup>.

## THE CAMERA MODEL

The goal of this work is to develop a model which 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  $k$ 'th 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<sup>1,2</sup>:

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

where  $R^{(k)}$  is a sensor sensitivity function for the  $k$ 'th channel, and  $F^{(k)}$  is a wavelength independent linearization function, and  $\rho^{(k)}$  is the linearized camera response. The key assumption is that all non-linear 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<sup>2,4-8</sup>. This model is also assumed for the human 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 due to geometric effects, notably vignetting<sup>1</sup>, and a fall-off proportional to the fourth power of the cosine of the off axis angle<sup>1</sup>.

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 work in color, absolute light flux is somewhat arbitrary, being under aperture control, and 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 which 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 non-linearities. 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 non-linear behavior.

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

$$\rho^{(k)} = F^{(k)}(v^{(k)}) = \mathbf{L} \bullet \mathbf{R}^{(k)} \quad (2)$$

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

## MOTIVATION

We have become interested in color camera characterization as part of our research into computational color constancy<sup>9</sup>. 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  $v^{(k)}$ .

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

Since  $F^{(k)}$  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 alternatively a smoothed version of the data can be used to generate a look up table. Vora et al<sup>2</sup> used this method to verify 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<sup>8</sup>. 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 which 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 (2) becomes:

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

Equation (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  $L$  was 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 we wish to solve for. Wyszecski<sup>10</sup> 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 since it is mainly the noise that is being fitted, and the resulting sensor responses tend to have numerous large spikes, and have an abundance of non-negligible negative values (see Figure 3(a)).

Sharma and Trussell<sup>4,7</sup> improved the prospects for a reasonable solution by introducing various constraints on  $\mathbf{R}^{(k)}$ . First, instead of solving (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 are all convex, and so they computed a resulting constraint set using the method of projection onto convex sets.

Hubel et al.<sup>11</sup> 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<sup>12</sup>, as a method for finding a smooth fit. They found that the method gave generally good results.

They note, however, that the method produced negative lobes in the response functions, and mention briefly 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<sup>6</sup>. Rather than constrain the absolute RMS error, we chose instead to minimize the relative RMS error. We then re-wrote 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 uni-modal, and that once the sensors dropped to a small value they remained small. On these grounds we also constrained the sensors to be zero outside a certain range on subsequent runs. In this particular case this forced the sensors response functions to be uni-modal. This last step needs to be applied with care, as it is possible that the sensors are in fact non-zero beyond the points where the main peaks drop to zero.

Recently Finlayson et al. used a similar approach<sup>13</sup>. 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 in terms of the peak location and report results constraining the sensors to be uni-modal and bi-modal. 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**

We now describe our proposed characterization method in two stages. First, we will describe 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  $F^{(k)}$  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 initial work<sup>6</sup> we minimized the relative RMS error in equation (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 minimize the relative error and the non-smoothness measure together. This allows fitting error to be traded against non-smoothness and vice versa. With the hard constraint used previously, there is no recourse in the case that 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 minimize 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 minimizing absolute error, or even a weighted combination of both, may make more sense. 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 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 different 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 paper. The results for the red channel are plotted in Figure 1.

We consider the variance to be the sum of intensity dependent and intensity independent parts. We further assume that the intensity dependent variance is due to photoelectron shot noise, and thus is proportional to the mean<sup>14</sup>. Thus we expect that the observed variance to be linear with the intensity with an offset characterizing the noise due to other sources. This is more or less



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

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

with the first term expressing the error, and the second term expressing non-smoothness and thus providing the regularization. The coefficient  $\lambda$  specifies the relative weight attributed to the two terms. If  $\lambda$  is zero and there are no constraints, then this becomes the pseudo inverse method. A serviceable value for  $\lambda$  is easily found by trial and error. To ensure positivity, we use the constraint:

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

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

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

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

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

This condition was not used for any of the results in this paper, but may be interest, and can be used to more closely emulate the method of Sharma and Trussell<sup>4,7</sup>.

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

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

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

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

We then replace (5) with:

$$\begin{vmatrix} \mathbf{L}_{rel} \\ -\lambda \mathbf{S} \end{vmatrix} \mathbf{R}^{(k)} = \begin{vmatrix} \mathbf{1} \\ \mathbf{0} \end{vmatrix} \quad (12)$$

Similar to the constraint (9), if we require a constraint limiting the relative error to less than a positive amount  $\zeta$ , we can use:

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

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

We note that minimizing the relative error may need to be modified slightly to deal with 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. Equation (12) can be interpreted as a weighted version of equation (5), with the weighting being inversely proportional to  $\rho^{(k)}$ . Thus it is natural and easy to put an upper bound on this weighting to safeguard against 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  $\rho^{(k)}$  due to the relatively large response of our camera to no light (specifically (11.05, 13.06, 12.36)).

We now specify the constraints used to implement Finlayson et al.'s method<sup>13</sup>. They constrain the modality of the sensor response functions which requires specifying the peak locations. In the uni-modal case, we specify one peak at  $i_{max}^{(k)}$ . Then we use the simple constraints:

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

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

Finlayson et al.<sup>13</sup> also propose ensuring smoothness by restricting the sensor response functions to be linear combinations of Fourier basis functions. To implement this approach we form 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:

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

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

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

Then

$$\mathbf{R}^{(k)} = \mathbf{B}\mathbf{B}^T \mathbf{R}^{(k)} \quad (17)$$

And

$$(\mathbf{B}\mathbf{B}^T - \mathbf{I})\mathbf{R}^{(k)} = \mathbf{0} \quad (18)$$

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

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 since 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 due to  $F^{(k)}$  and  $\mathbf{R}^{(k)}$  can be traded against each other. We first make a rough estimate of  $F^{(k)}$  which we use to propose a parameterized expression for it. We then fit the parameters for  $F^{(k)}$  and  $\mathbf{R}^{(k)}$  simultaneously. We will now provide a specific example of such a strategy.

The Sony DXC-930 camera which 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 non-linearity for small pixel values. Due to this non-linearity, 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 non-linearity must be taken into account, even if it is not explicitly estimated. Figure 2 shows the slight non-linearity for the red channel. The other two channels are similar.

The fit shown in Figure 2 was found using a simple linear fit (see (20) below) using pixel values greater than 30. This illustrates the point that linearization information is available in the data set which is to be used to find  $\mathbf{R}^{(k)}$ . To proceed with our strategy of explicitly fitting the non-linearity, 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 non-linearity of our camera could be approximated by:

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

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

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

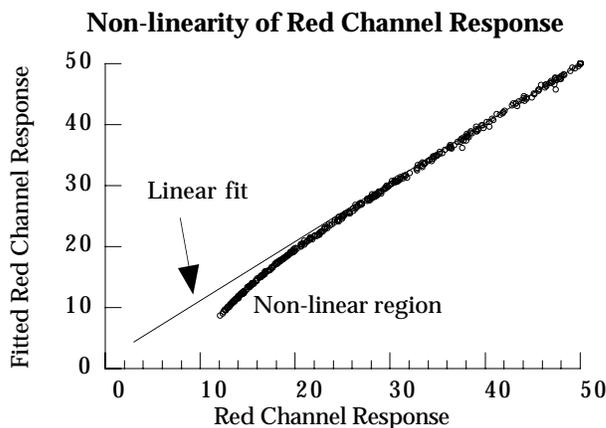


FIG. 2. The non-linearity of the red channel response for the Sony DXC-930 camera used for the sensor fitting experiments. The fit shown is a simple linear fit on pixel values greater than 30. The other two channels have similar curves.

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 which roughly fit the non-linearity so that the overall fitting procedure can find a good estimate of  $F^{(k)}(x)$ . In addition, the parameters which 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 equation (5) with:

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

where the arithmetic in the upper right block of the matrix is done element-wise as needed.

Similarly, in the case of fitting for relative error, we replace (12) with:

$$\left| \begin{array}{c|cc} \mathbf{L}_{rel} & \mathbf{1} & \frac{e^{-C_k(\mathbf{r}^{(k)}-b_k)}}{\mathbf{r}^{(k)}} \\ \hline \lambda S & 0 & 0 \end{array} \right| \bullet \left| \begin{array}{c} \mathbf{R}^{(k)} \\ \hline a_0^{(k)} \\ a_1^{(k)} \end{array} \right| = \left| \begin{array}{c} \mathbf{1} \\ \hline \mathbf{0} \end{array} \right| \quad (22)$$

where again, the arithmetic in the upper right block of the matrix is done element-wise as needed.

Note that the response vectors  $\mathbf{r}^{(k)}$  now correspond to the observed camera responses  $v^{(k)}$  in equation (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 use 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 (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 over-fit 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. Since the actual sensors are relatively smooth and non-negative, we expect methods which promote these characteristics to do better.

The results are shown in Table I, and the sensors obtained using a selection of the methods are plotted together with the actual sensors in Figure 3. As predicted, the lowest fitting error is 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 additionally using the range constraints, but it should be noted 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 than the simple regularization approach proposed in this work. The results in Table I indicate that there is no choice of the number of basis functions which give 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 can not be considered. In contrast, the approach proposed in this work allows the degree of smoothness to be traded against fitting error, which yields more flexible fitting.

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 (percent)	RMS difference between fitted sensor curve and the target, normalized by the 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 and smoothing and range	4.820	0.0424
Fitting with positivity and 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

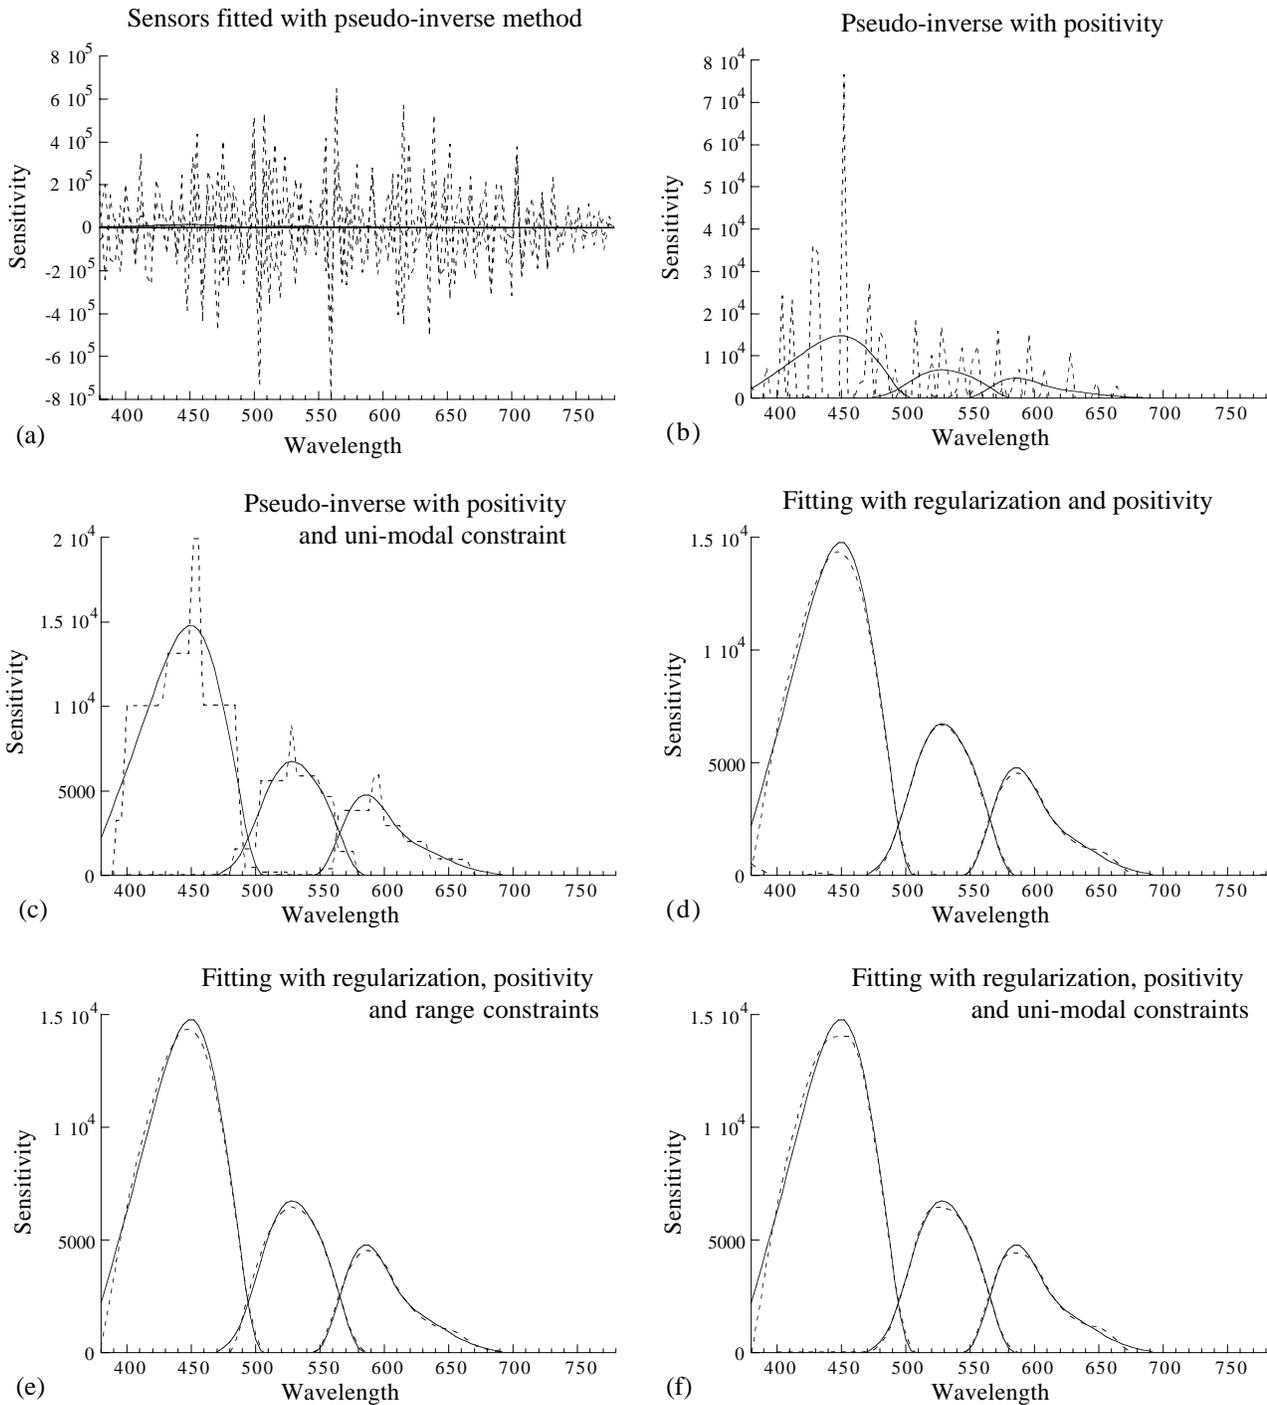


FIG. 3. The results of various fitting methods on synthetic data. The data was generated from idealized sensors based on the actual sensors of our Sony DXC-930 camera. 5% relative Gaussian noise was added.

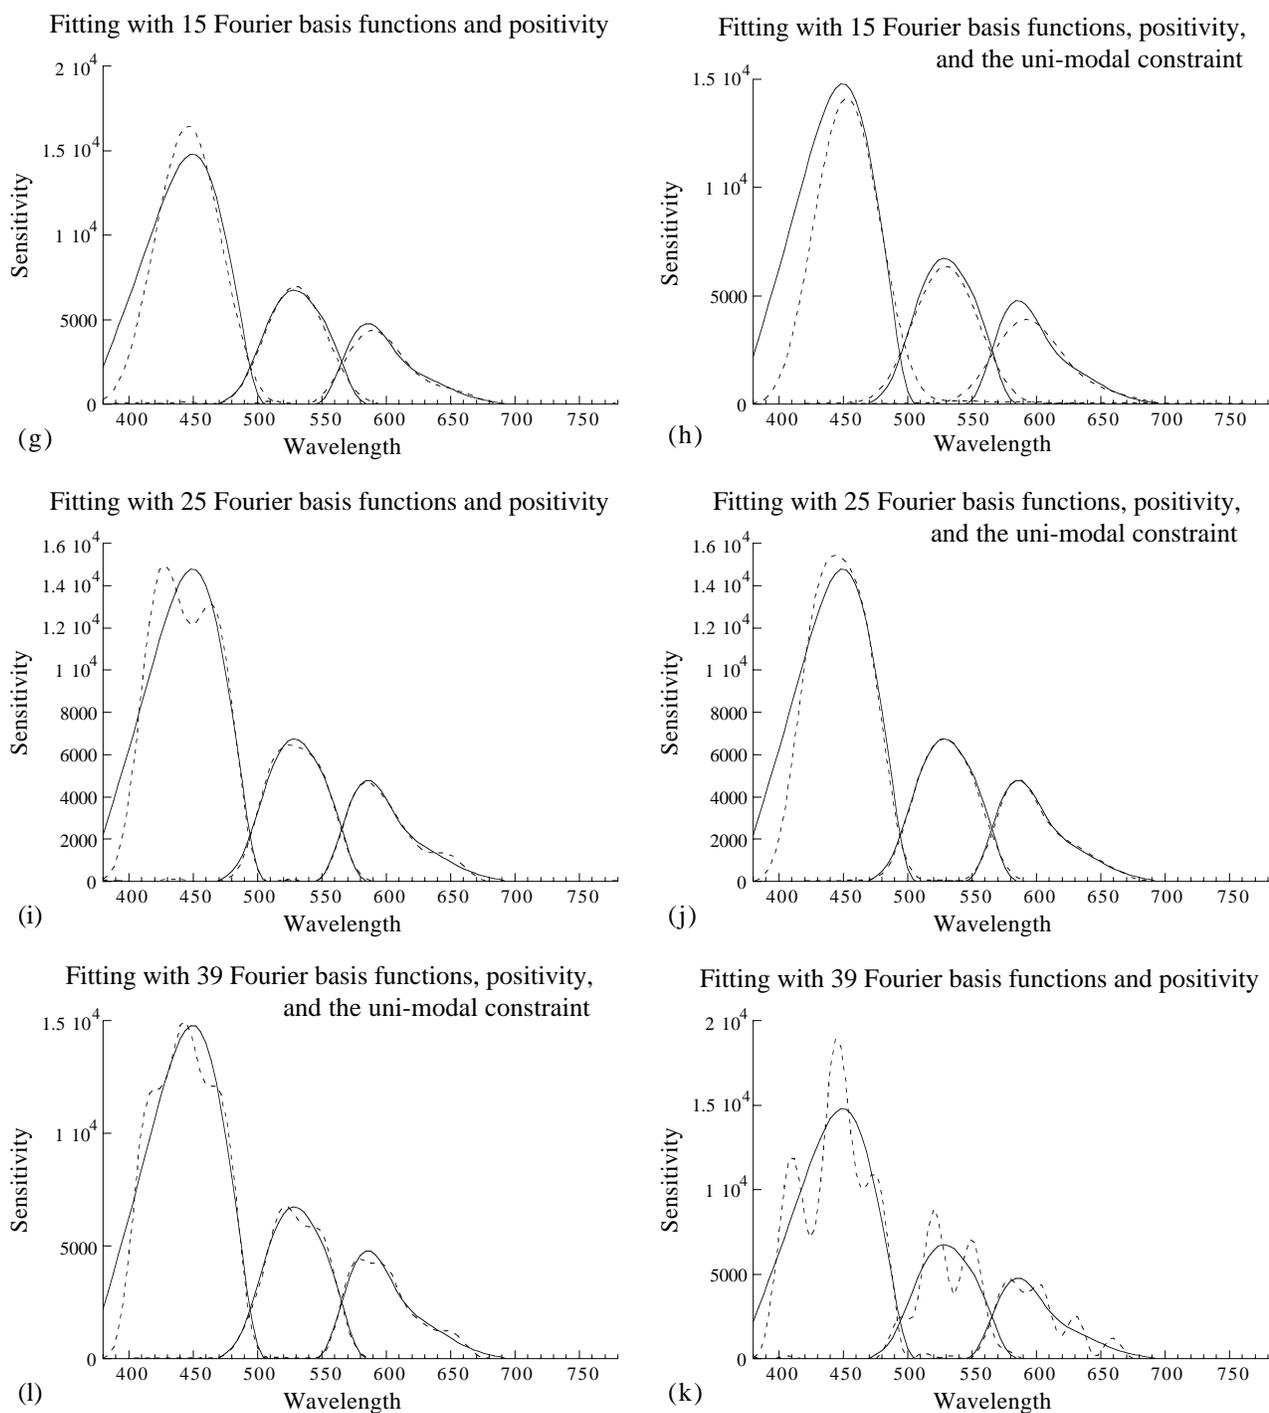


FIG. 3 (Continued). The results of various fitting methods which promote smooth results by constraining the solution to be a linear combination of a specified number of Fourier basis functions.

## EXPERIMENTS WITH REAL DATA

We investigated camera characterization for a Sony DXC-930 3 chip CCD camera. In order 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 which 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 of the apparatus was to ensure that the camera and the spectroradiometer measured the same signal. We also required that the camera data was always for center of the image. Therefore we mounted the color checker horizontally on an XY table which 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 worth of camera data before capturing an entire chart worth of spectra. A total of 26 illuminant/filter combinations were used in conjunction with the 23 patches, providing 598 measurements (available on-line<sup>3</sup>). 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, exposing the region of interest, but as little else as was practical. We extracted a 30 by 30 window from the image which 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 to simply 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 we make the intercept of the linear fit a parameter of the fitting process (as in (20)). Due to the curvature evident in Figure 2, this method results in the subtraction of an offset which is somewhat less than the camera black used in the first approach. Finally we provide the results of parameterizing the non-linearity as developed above.

Each linearization method was used in conjunction with a number of methods for fitting the camera response functions. The results are compiled in Table II. All results in this table are based on minimizing the relative error. The regularization smoothing parameter,  $\lambda$ , was initially set by trial and error to a value which gave reasonably smooth sensor functions. We did not attempt to tune  $\lambda$  beyond a factor of two, 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 reduces 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.

In Table III we compare fitting based on absolute error with fitting based on relative error for one of the methods. Not surprisingly, when the data was fit using relative error, the relative error was lower, and the absolute error higher, than when the data was fit using absolute error. More significantly, fitting with relative error substantially reduced the absolute error in (r,g) chromaticity which is difficult to minimize directly, and is key for many applications. Table III also includes  $L^*a^*b$  error for the two objective functions.

Unlike the synthetic case, the true camera sensor functions are not known. Thus in order 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 was 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 was 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 to 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 uni-modality constraint becomes increasingly useful. Of course, a small data set could not be used to determine with confidence that sensors are in fact uni-modal. Interestingly, the pseudo-inverse method with positivity and uni-modality gives 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 adds 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.

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

TABLE III. A comparison of fitting based on relative error with fitting based on absolute error for one of the preferred methods (positivity, smoothing, and uni-modality). 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

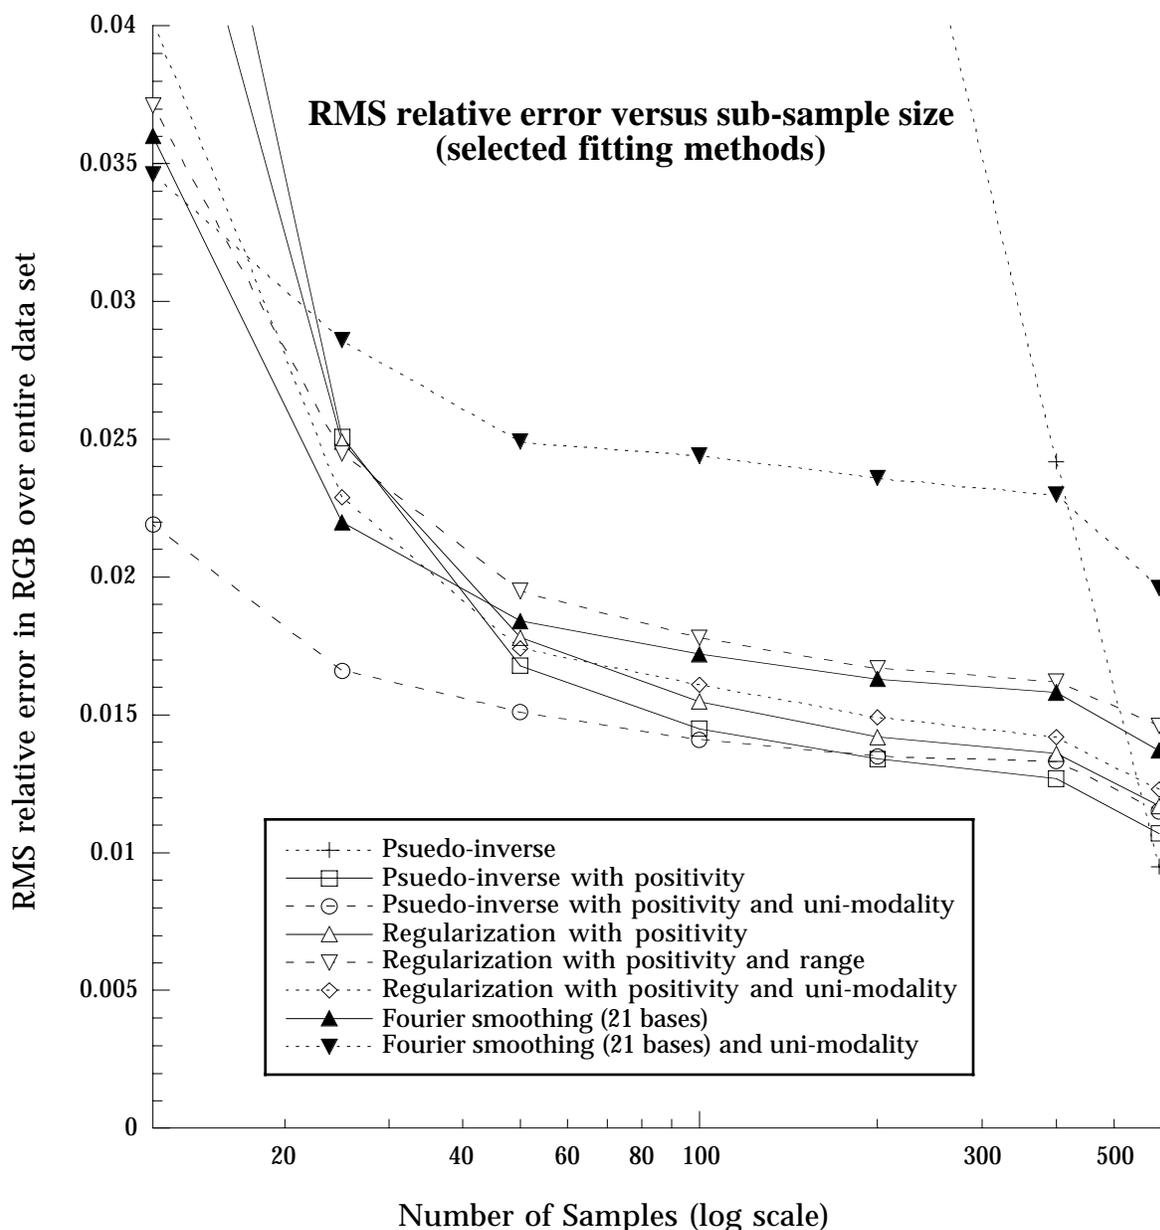


FIG 4. RMS relative error of various fitting methods versus sample size. The more constrained methods tend to more robust as the number of sample points decreases, but have higher error when the number of points is large. For example the unconstrained pseudo inverse method (first curve) has the least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

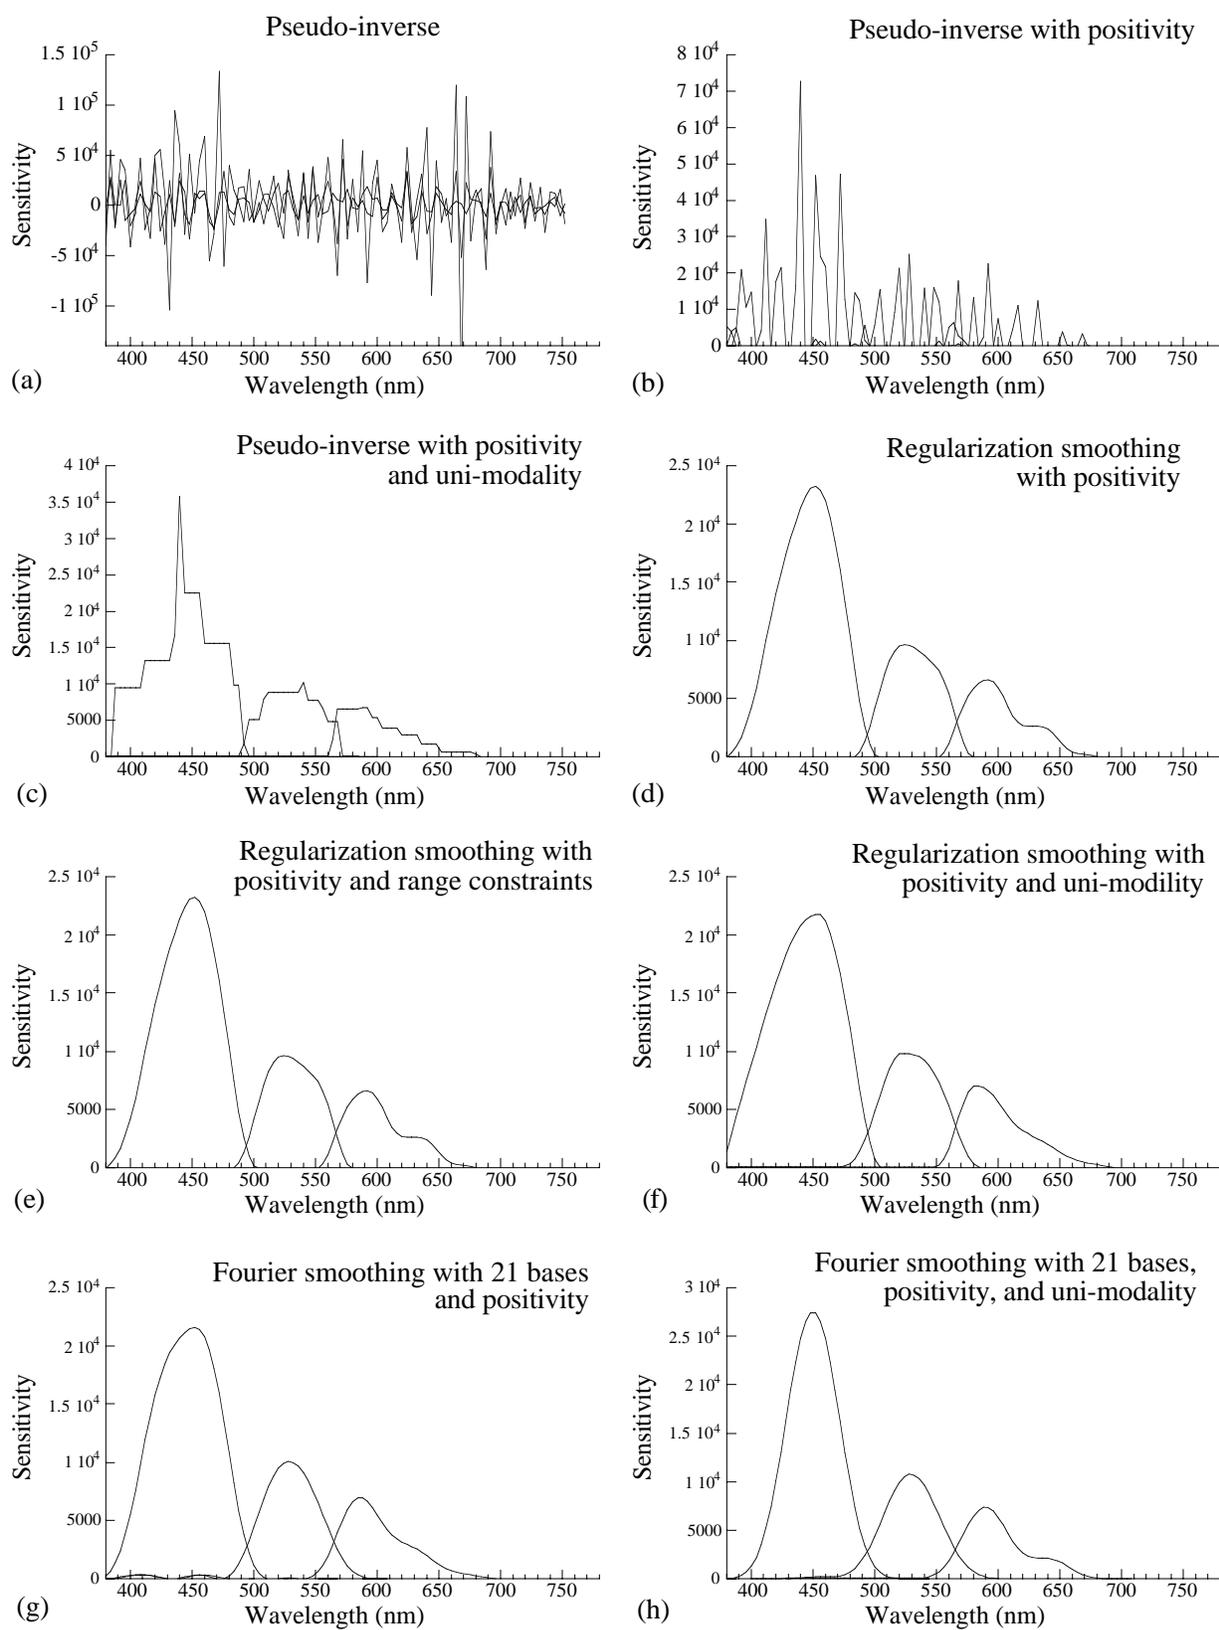


FIG. 5. Sensor response functions found using a variety of fitting for the data collected as described in the text. All methods were used in conjunction with linearity fitting. These sensors correspond to the results in Figure 4 when for the full sample size of 598 points.

## 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 obtain a result which 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 have also investigated fitting a small non-linearity in the camera response simultaneously with the sensor response functions. This is effective because errors due to 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.

## REFERENCES

1. B. K. P. Horn, *Robot Vision*, MIT Press, 1986.
2. P. L. Vora, J. E. Farrell, J. D. Tietz, and D. H. Brainard, Digital color cameras. 1. Response models, Hewlett-Packard Laboratory, Technical Report HPL-97-53, available from <http://www.hpl.hp.com/techreports/97/HPL-97-53.html> (1997).
3. K. Barnard, Lindsay Martin, B. Funt, and A. Coath, Data for colour research, available from <http://www.cs.sfu.ca/~colour/data>
4. G. Sharma and H. J. Trussell, Characterization of Scanner Sensitivity, in *Proceedings of the IS&T and SID's Color Imaging Conference: Transforms & Transportability of Color*, The Society for Imaging Science and Technology, Springfield, Va., 103-107 (1993).
5. G. E. Healey and R. Kondepudy, Radiometric CCD camera calibration and noise estimation, *IEEE Pattern Anal. Mach. Intell.*, **16**, 267-276 (1994).
6. K. Barnard, Computational colour constancy: taking theory into practice, Simon Fraser University School of Computing Science, M.Sc. thesis, available from <ftp://fas.sfu.ca/pub/cs/theses/1995/KobusBarnardMSc.ps.gz> (1995).
7. G. Sharma and H. J. Trussell, Set theoretic estimation in color scanner characterization, *J. Elec. Imag.*, **5**, 479-489 (1996).
8. P. L. Vora, J. E. Farrell, J. D. Tietz, and D. H. Brainard, Digital color cameras. 2. Spectral response, Hewlett-Packard Laboratory, Technical Report HPL-97-54, available from <http://www.hpl.hp.com/techreports/97/HPL-97-54.html> (1997).
9. K. Barnard, Practical colour constancy, Simon Fraser University School of Computing Science, Ph.D. thesis, available from <ftp://fas.sfu.ca/pub/cs/theses/1999/KobusBarnardPhD.ps.gz> (1999).
10. G. Wyszecki, Multifilter method for determining relative spectral sensitivity functions of photoelectric detectors, *J. Opt. Soc. Am.*, **50**, 992-998 (1960).
11. P. M. Hubel, D. Sherman, and J. E. Farrell, A comparison of method of sensor spectral sensitivity estimation, in *Proceedings of the IS&T/SID 2nd. Color Imaging Conference: Color Science, Systems and Applications*, The Society for Imaging Science and Technology, Springfield, Va., 45-48 (1994).
12. W. K. Pratt and C. E. Mancill, Spectral estimation techniques for the spectral calibration of a color image scanner, *Applied Optics*, **15**, 73-75 (1976).
13. G. Finlayson, S. Hordley, and P. Hubel, Recovering device sensitivities with quadratic programming, in *Proceedings of the IS&T/SID Sixth Color Imaging Conference: Color Science, Systems and Applications*, The Society for Imaging Science and Technology, Springfield, Va., 90-95 (1998).
14. G. C. Holst, *CCD Arrays, Cameras, and Displays*, 2 ed, SPIE Press, Bellingham, WA, 1998.