

Image Capture: Simulation of Sensor Responses from Hyperspectral Images

Poorvi L. Vora, Joyce E. Farrell, Jerome D. Tietz, and David H. Brainard

Abstract—This paper describes the design and performance of an image capture simulator. The general model underlying the simulator assumes that the image capture device contains multiple classes of sensors with different spectral sensitivities and that each sensor responds in a known way to irradiance over most of its operating range. The input to the simulator is a set of narrow-band images of the scene taken with a custom-designed hyperspectral camera system. The parameters for the simulator are the number of sensor classes, the sensor spectral sensitivities, the noise statistics and number of quantization levels for each sensor class, the spatial arrangement of the sensors and the exposure duration. The output of the simulator is the raw image data that would have been acquired by the simulated image capture device.

To test the simulator, we acquired images of the same scene both with the hyperspectral camera and with a calibrated Kodak DCS-200 digital color camera. We used the simulator to predict the DCS-200 output from the hyperspectral data. The agreement between simulated and acquired images validated the image capture response model and our simulator implementation. We believe the simulator will provide a useful tool for understanding the effect of varying the design parameters of an image capture device.

I. INTRODUCTION

IN THIS paper, we describe how one can simulate the camera's response to a scene from hyperspectral image data.

The light sensors in many modern image capture devices (e.g. digital scanners and digital cameras) are based on charge-coupled device (CCD) or active pixel sensor (APS) technology. These devices are usually designed so as to have linear irradiance-response functions over most of their operating range [1]. The overall camera system may not exhibit the underlying device linearity, however. For example, there may be a nonlinear mapping between the raw sensor output and the digital responses actually available from the camera. Such a nonlinearity might be designed into a camera system if the quantization precision of the sensor itself is larger than that of the camera. This is the situation with the Kodak DCS-420. It employs a 12-bit internal data representation for measurements that are linear with respect to irradiance, but its standard control software provides

only eight bits of precision and eight-bit output that is nonlinear with respect to irradiance.

Color cameras require multiple classes of sensors with different spectral sensitivities. By placing color filters in series with either CCD or APS sensors, usually on a pixel-by-pixel basis, such multiple classes can be created. When the color filters are placed in a mosaic pattern, one color per pixel, the cameras are referred to as color filter array (CFA) cameras. Evaluation of digital camera design parameters has received considerable attention in the recent literature [2]. These evaluations are based on theoretical models of image statistics and simple image quality metrics. A useful complement to the theoretical approach is to evaluate the performance of different camera designs for actual scenes. A difficulty with this approach is that it is not always feasible. This paper describes a method for constructing, testing and evaluating the performance of an image capture device simulator. A reliable simulator provides a means for evaluating the performance of a complete image capture device design prior to manufacture.

The simulator we describe is based on several simplifying assumptions about the image capture device. These are that the

- 1) optical system is linear and shift invariant;
- 2) response of the sensors to light at varying intensities and wavelengths is known;
- 3) sensor noise is additive.

The input to the simulator is a hyperspectral image of the scene, which provides the full spectral power distribution of the incident light at every image location. The image is acquired with a custom-built hyperspectral camera system described in this paper. Given the hyperspectral image, the simulator computes the response of the image capture device using the response models developed here.

This paper is organized as follows. Section II briefly describes work in camera modeling. In Section III we describe the simulator and Section IV contains experimental results verifying its accuracy. Section V presents conclusions and future directions. The Appendix details the modeling and calibration methods and results.

II. CAMERA MODELS

The range of operation of CCD sensors is highly linear. We assume two models for the entire camera system—the first is a simple linear model, the second a linear model for the sensor with a static nonlinearity. These models are used to calculate the output of a camera given its input and are described briefly in this section before we describe the simulator.

Manuscript received November 13, 1998; revised August 9, 2000. This work was presented in part in [3] and [4]. D. H. Brainard was supported in part by a philanthropic grant from Hewlett-Packard. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Glenn Healey.

P. L. Vora is with the Hewlett-Packard Company, Corvallis, OR 97330 USA (e-mail: poorvi@hpl.hp.com).

J. E. Farrell is with Shutterfly, Redwood City, CA 94065 USA.

J. D. Tietz and D. H. Brainard are with the Department of Psychology, University of California, Santa Barbara, CA 93106 USA.

Publisher Item Identifier S 1057-7149(01)00115-4.

A. Linear Response Model

Assuming linearity, the output of a sensor array at grid position (m, n) maybe be approximated as

$$\mathbf{r}(m, n) \approx e \sum_{k=1}^N \mathbf{s}(m, n, k) \mathbf{c}(m, n, k) \Delta_\lambda \delta^2 + \text{noise} \quad (1)$$

where

e	exposure setting;
argument k	variation with wavelength;
$\mathbf{c}(m, n, k)$	spectral sensitivity of the sensor at position (m, n) ;
$\mathbf{s}(m, n, k)$	irradiance distribution incident on the camera at position (m, n) ;
Δ_λ	wavelength sampling for the irradiance and spectral response functions;
δ	spatial sampling rate (i.e., the distance between contiguous sensors, assumed to be uniform and identical in both horizontal and vertical directions);
noise	sensor measurement noise.

Correct calibration allows us to drop the constant $\Delta_\lambda \delta^2$ in the above sum. In the formulation of (1) we neglect optical blur of the camera. This is justified for the moment because we consider only images of the Macbeth ColorChecker Chart (MCC), a low spatial frequency target, where optical blurring is not a critical factor. We also assume that the spectral response of a single sensor is constant over each pixel and that the wavelength sampling used is fine enough to accurately represent the spectral response.

The Kodak DCS-200 is an example of a camera system which responds linearly to irradiance. Experimental results showing this are reported in [3] and summarized in the Appendix.

B. Static Nonlinearity Model

The behavior of some cameras like the Kodak DCS-420 can be described by a *static nonlinearity model*. For this model, the camera response for a pixel of the i th sensor type pixel is given by

$$\mathbf{r}(m, n) \approx \mathcal{F} \left(e \sum_{k=1}^N \mathbf{s}(m, n, k) \mathbf{c}(m, n, k) \Delta_\lambda \delta^2 + \text{noise} \right) \quad (2)$$

where \mathcal{F} is a monotonically increasing nonlinear function.

The Kodak DCS-420 operating in its standard eight-bit mode is an example of a camera system which responds according to the static nonlinearity model. Experimental results showing this are reported in [3] and summarized in the Appendix.

III. CAMERA SIMULATOR

In [3] and [4], we have shown that a linear model with the application of a static nonlinearity (if necessary) is good at predicting sensor response from digital image sensors. These results are further summarized in the Appendix. In this section, we describe the use of the linear model to simulate the response of a specific digital camera to a specified scene. We also describe a

hyperspectral camera system used to provide input to the camera simulator. The digital camera is defined by the values of a set of camera design parameters, and the scene is defined by irradiance values as a function of space and wavelength.

A. Simulator Description

The input to the simulator consists of a set of images which together represent an approximation of the irradiance incident on the camera lens as a function of space and wavelength. At present, the input is a set of 31 images taken with the hyperspectral camera system described in Section III-B. The 31 images each represent a spatial distribution of the incident irradiance over a narrow range of wavelength values. The simulator can also take a different representation of the input, for example a set of images each of which represents the spatial distribution of coefficients with respect to a basis set of principal components of a database of radiant spectra. The simulator uses (1) or (2) to compute the simulated output image.

The parameters of the simulator are

- 1) sensor spectral sensitivities as a function of wavelength, sampled at the same rate as the representation of the input (or represented in terms of the same basis vectors as the input);
- 2) exposure time;
- 3) noise statistics (mean and variance);
- 4) mosaic pattern;
- 5) number of bits per pixel of the camera sensors;
- 6) function \mathcal{F} .

A final step of the simulation, not described by (1) and (2), is to quantize the simulated output to the same number of bits as the simulated device. The output raw data image of the simulator may be directly compared with the camera image for numerical verification. For visual quality judgements, we need to demosaic and color correct the outputs of the DCS-200 and the simulator and then compare the color images.

B. Hyperspectral Camera

The hyperspectral camera was based on a scientific grade monochrome CCD camera (Photometrics PXL, 2 K by 2 K spatial resolution at 12 bits/pixel direct digital interface, electronic control of temporal integration, liquid cooling reduces dark current, irradiance-response function is linear, 50 mm Nikkor lens) interfaced to a Macintosh host. Hyperspectral images were acquired by sequentially placing interference filters in the optical path of the camera and acquiring monochromatic images through each filter. For the work reported here, the filters were placed directly in front of the camera lens. The filters were narrowband (roughly 10 nm full-width at half-height) with peak wavelengths that evenly spanned the visible spectrum (31 filters, 400 nm to 700 nm, in 10-nm steps). Aperture (f-stop) and exposure duration for each monochromatic image were adjusted so that the image data were within the response range of the camera, and focus was adjusted individually for each monochromatic image. Dark images (camera shutter not opened) were acquired for each exposure duration used, and the appropriate (same exposure duration) dark image was subtracted from each monochromatic image. A spectroradiometer




Fig. 1. RGB Scatter plot of MCC color simulation.

(PhotoResearch PR-650) was used to measure the spectrum from a small, approximately uniform, image region at the same time as the 31 monochromatic images were acquired. This measured reference spectrum was used to calibrate the dark-subtracted monochromatic images. Image data from the region corresponding to the radiometer measurement were extracted and a scale factor for each monochromatic image was computed that brought the extracted data into agreement with the reference spectrum. Similar calibration procedures have been used by others [5], [6]. Since the focus of the present work is on the chromatic properties of the simulator, no attempt was made to correct for small alignment errors between individual monochromatic images and we did not evaluate the spatial point spread function of the hyperspectral camera. A website [7] provides information about the ongoing development of our hyperspectral camera system. Descriptions of similar hyperspectral cameras may be found elsewhere [5], [6], [8]–[10].

IV. SIMULATOR VERIFICATION

This section describes the experiments performed to verify the accuracy of the simulator for color patches. The DCS-200 was used to take pictures of the MCC in a laboratory illuminated by an incandescent source (Kodak 4400 Slide Projector) at different exposure settings: 1/8, 1/15, 1/30, 1/60, and 1/125 s. The hyperspectral camera system was also used to acquire a 31-band hyperspectral image of the same chart under the same conditions. The hyperspectral image, the DCS-200 spectral sensitivities (estimated as described in [4] and summarized in the Appendix), the exposure durations, the measured noise statistics, the bits per pixel value (8) and the calibrated nonlinearity (the identity) of the DCS-200 were used to generate simulated images. The simulated images were compared both numerically and visually.

Fig. 1 is a scatter plot of the real versus simulated R, G, B values over all the shutter speeds studied in the experiment. It is clear that the agreement between real and simulated values is good. Table I lists the numerical error statistics of $R, G,$ and B values averaged over the center of each patch in both real and

TABLE I
STATISTICS OF SIMULATION ERROR FOR
MCC-DCS-200

Exposure Setting in sec.	Mean of Absolute Error	Mean of Absolute % Error	RMS Error	Maximum Absolute Error	Maximum Absolute % Error
$\frac{1}{8}$	1.49	1.22	2.59	9.07	13.23
$\frac{1}{15}$	3.18	4.03	4.11	9.64	30.99
$\frac{1}{30}$	4.96	7.66	5.43	13.21	24.13
$\frac{1}{60}$	3.47	9.39	3.66	6.75	26.19
$\frac{1}{125}$	1.85	7.22	1.96	3.04	15.16

simulated images. Since the DCS-200 is an 8-bit camera, the numerical response values range between 0 and 255. The values in the table are computed from response values on this scale.

For purposes of comparison, the spatial variation in a dark noise image taken with the DCS-200 has a root mean square value of about 0.89 (see the Appendix). Linearity tests for this camera have shown that the root mean square value of the variation from linearity is 1.45 [3] (see the Appendix). Furthermore, the correspondence between predicted and empirical sensor values is consistent across exposure settings, supporting our assumption of linearity with respect to exposure duration—an assumption implicit in (1), and verified through earlier experiments reported in [3] and [4] and summarized in the Appendix.

We visually compared the predicted and empirical sensor data after processing the data with a simple demosaicing routine based on bilinear interpolation. (The complexity of the demosaicing routine is not expected to make a difference to visual quality, as the images consist of large color patches.) The real and simulated images had similar appearances.

To further quantify the accuracy of the simulator with respect to perceptual error, we calculated the perceptual error represented by the real and simulated R, G, B values for one simulated image and the corresponding real image. For this comparison we chose the image that covered best the dynamic range of the camera. We used the estimated camera spectral sensitivities (see the Appendix for details of the estimation procedure) and to calculate a linear minimum mean square error correction of camera RGB values recorded in the tungsten illuminant to CIE tristimulus values for illuminant D65. This correction is equivalent to a regression estimate for simulated camera RGB values under tungsten for the Macbeth chart to CIE tristimulus values under D65. As the data set is known to be the Macbeth Chart, we chose to use this fact in the color correction procedure. Our method is based on [11].

The average absolute error between the real and simulated values is very low for the X and Y values (0.0189 and 0.0193 out of maximum values of 0.6991 and 0.7044 respectively; corresponding to 2.7% and 2.74%, respectively) and slightly higher for the Z values (0.0887 out of a maximum value of 1.3548 corresponding to 6.55%). This is probably because the color correction matrix has a very large term for the conversion from blue to Z , which would multiply the contribution of errors in blue readings to Z values. The large term in the color correction matrix is because of the low blue sensor response of the DCS-200 camera and the fact that the image was taken in tungsten light which further reduces the blue response of the recording system.

We calculated CIE Lab values corresponding to the tristimulus values, assuming a white point of (0.6991, 0.7044, 1.3548)—the XYZ values for the macbeth white square in the real DCS-200 image after correction to XYZ values in D65. The perceptual agreement between real and simulated values—as far as it is represented by similar CIELab values—is fair. The maximum error in CIELab space is 14.79, and the average is 6.6874.

A comparison of Fig. 1 and the bottom panel of Fig. 4 suggests that much of the simulation error arises because of differences between the hyperspectral image data compared to direct radiometric measurements. Such differences could be introduced because our hyperspectral camera calibration does not account for geometric distortions. It is worth noting that for purposes of evaluating camera design it is not critical that the hyperspectral image data exactly measure the input scene, since comparisons will be between performance of different camera designs simulated for the same hyperspectral input.

V. CONCLUSIONS AND FUTURE DIRECTIONS

The use of a linear model with the application of a static nonlinearity if necessary is appropriate for the simulation of the sensor responses of color filter array cameras. The color fidelity of output simulated using the linear model is good for a wide range of exposure settings. In the future we will incorporate a model for the optical system to simulate the effects of lens blur. The simulation of spatial effects (including the effect of a lens blur that may vary as a function of position with respect to center-field and wavelength, inter-sensor charge leakage and the mosaic pattern) on the visual quality of an image will be verified by using calibrated hyperspectral input images of scenes with richer spatial variation. Another direction for future study is the calculation and correlation of the perceptual error of the simulator with observer preferences.

APPENDIX

We summarize relevant experimental results for modeling and calibrating the camera in the Appendix. These results are detailed in [3], [4]. This Appendix covers the camera model linearity assumption, calibration of a static nonlinearity, noise measurement and characterization and spectral calibration for the Kodak DCS-200 and the Kodak DCS-420 cameras.

A. Camera Linearity

To test the linearity of the camera response, we measured the irradiance-response functions at several exposure durations for both cameras.

A typical result is shown in the left image of the first row in Fig. 2. The x -axis shows the irradiance of the incident light (calculated as described in [3]) and the y -axis shows the camera output value (with the expected value of the noise subtracted). The crosses represent actual data points. The straight lines are fit to the data and constrained to pass through the origin. In fitting the data, we excluded saturated points and points with very low intensities. The good agreement between the data and the fit lines indicate that the DCS-200 has a linear irradiance-response function over most of its operating range.

We note that the performance of a second digital camera (the Kodak DCS-420) is not well-described by the linear model, at least when it is operated with the standardly-supplied eight-bit acquisition software [3]. A typical result is plotted in the right image in the first row of Fig. 2. In this figure, the expected value of the dark noise has not been subtracted from the camera output. The relationship between irradiance and response was clearly nonlinear. A probable cause for this nonlinearity is the 12-to-8-bit reduction in the image acquisition software. We describe the calibration of the static nonlinearity of the Kodak DCS-420 in Appendix Section C.

B. Camera Response to Variation in Exposure

To test for linearity with exposure duration in the Kodak DCS-200, we took pictures under fixed illumination at different exposure durations. The left panel of the second row of Fig. 2 shows typical results. As with the figure above it, the crosses represent actual data points with the expected value of the noise subtracted and the lines are fits constrained to pass through the origin. A slight variation from linearity may be due to the fact that the shutter exposure time is not controlled accurately.

As the irradiance-response function of the DCS-420 is not linear, it would be surprising if its output were linear with exposure duration. We roughly calculated the average green sensor (G) value at the center of the image field for images taken at various exposure durations for 525-nm illumination. The right panel of the second row of Fig. 2 shows the results with average noise subtracted (x 's) overlaid on irradiance-response data (re-plotted as o 's). The x -axis represents exposure duration relative to one second and irradiance relative to unity. The two readings corresponding to one second and unit irradiance are replications of the same illumination condition, so that no scaling of the data were required. The close agreement between the two curves suggests that the same nonlinearity mediates both.

C. Calibration of Linearity and Static Nonlinearity

Our data indicate that the linear response model describes the output of the Kodak DCS-200, at least over the output range 20–240 out of a total range of 0–255 camera units. To obtain the parameters describing a single line for all the data, we fit a *calibration line* to the data for the blue sensor readings of the left panel of the first row of Fig. 2, *without subtracting out the average dark noise value*. The calibration line is the solid line in the left panel of the last row of Fig. 2. Measured camera values on a scale of 0–255 are plotted on the x -axis, while linearized fractional input on a scale of 0–1 is plotted on the y -axis.

To verify that the calibration line derived from one irradiance response function describes all the data, we can use this line to normalize all of our data and examine it on a single plot. For each measured irradiance response function, the irradiance measure we used is arbitrary, since we varied both the exposure and wavelength across the different measurements. We can use the calibration line to normalize the data, however. For each data set, we found the highest camera output value in the linear range (below 240) and found its position on the calibration line. We then scaled all the irradiance values of that data set by a single normalization scale factor such that the highest camera output value in the linear range would correspond to the irra-

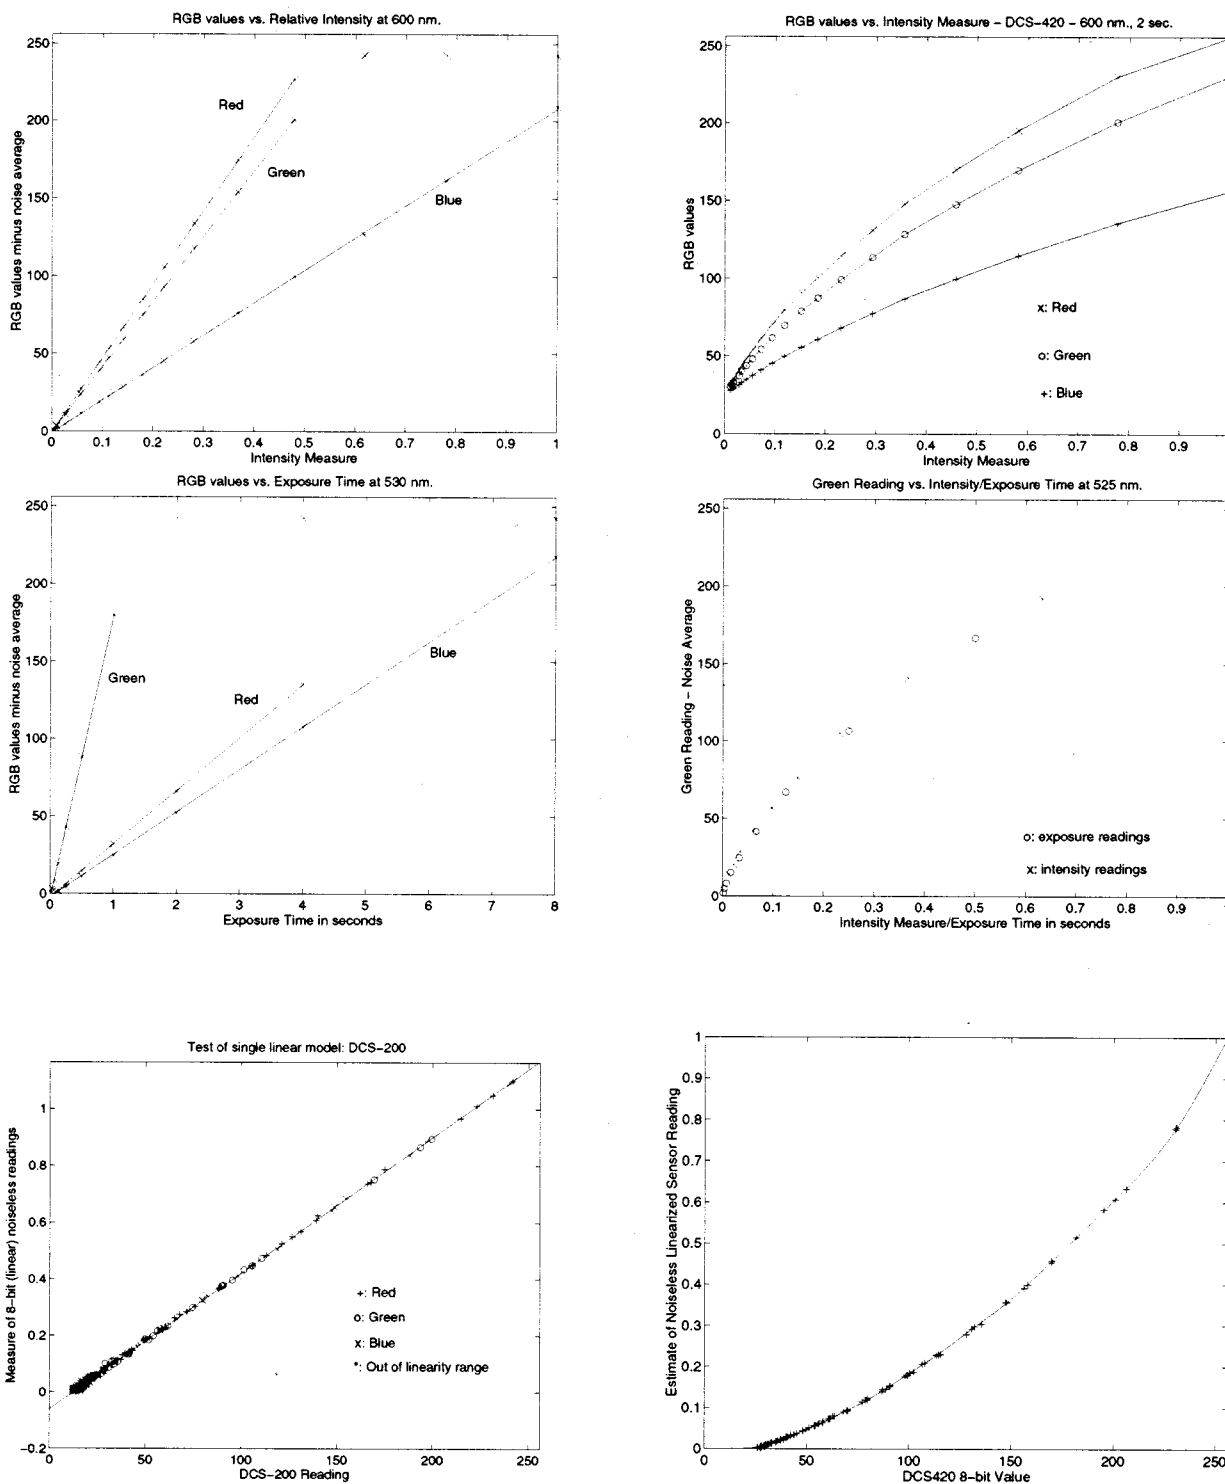


Fig. 2. Intensity and exposure responses.

diance factor obtained by looking at the calibration line. This procedure allows us to compare all of our data to the calibration line, as shown in the left panel of the last row of Fig. 2. The highest camera output value for each data set lies on the line because of the way the normalization is performed. Data points with values below 240 and above 20 all lie close to the line. Data points with values below 20 or above 240 are plotted with asterisks (*) or lie outside the region shown in the plot.

The Kodak DCS-420 is not linear. To examine whether the static nonlinearity response model described its performance, we asked how well a single function \mathcal{F} can describe its output across the conditions we measured. We used the measurements plotted in the right panel, first row, Fig. 2 as a reference. This series covered most of the dynamic range of the camera. By interpolating and extrapolating the reference, we obtain a *calibration curve* for the DCS-420 that maps between sensor values (0




Fig. 3. Variation from calibrated curves and noise.

to 256) to intensities that lie between 0 and 1. This irradiance measure is in arbitrary units but may be calibrated to physical units. The result is tabulated in [3] and graphed as the line in the right panel of the last row of Fig. 2. It represents the value of $\mathcal{F}^{-1}(r) - n$ (or $e \int_{\lambda_t}^{\lambda_r} s(\lambda) i(\lambda) d\lambda$) of (2).

We tested the accuracy of the calibration curve by asking how well it described the rest of our data. Each set of acquired data points has a different irradiance scale. A value of unit irradiance corresponds to the maximum irradiance for the shutter speed used for that test. To check if the other acquired data points lie on the calibration curve, the irradiance values need to be transformed to a single scale. We calculated the scale factor for the conversion for each data set by using the highest measured output value (which corresponds to a unit irradiance for that series), finding its position on the calibration curve, and using the fractional irradiance value thus obtained as the scale factor. The data points from all of our irradiance-response series as well as the exposure data are plotted in the right panel of the last row of Fig. 2 along with the calibration curve. The data all lie along the curve.

D. Variation from Linear Model and Static Nonlinearity

The data points vary slightly from the linear model for the DCS-200 and from the calibration curve for the DCS-420. In this section, we quantify the variation.

To estimate the slight variation from linearity of the DCS-200, we calculated the differences between values predicted by the straight line in the linearity plot of the left panel of the last row of Fig. 2 and actual values, for measured values above 20 and below 240. These differences are plotted in the top left panel of Fig. 3. This calculation assigns zero difference to the maximum value in each data set because of the way placement of all data points on one curve is performed and is thus only approximate. The error statistics reported below were calculated without using the maximum value in each data set, and are transformed from fractional values from 0 to 1 to camera values from 0 to 255. The mean absolute value of the variation is 1.13, and the mean value is 0.58. The average of the noise when estimated from the calibration curve is 12.5. This value is close to the value of 13.6 obtained by directly estimating the dark noise (see Appendix Section E below). The root mean square value of the variation is 1.45. The maximum error is 4.67 and occurs for a green sensor reading.

To quantify the slight variation of the scaled DCS-420 data points from the curve in the right panel of the last row of Fig. 2, we calculated the difference between the data point and the value on the curve corresponding to the scaled irradiance, i.e. the difference between indirectly measured values of $\mathcal{F}^{-1}(r) - n$ and values obtained from the calibration curve. As for the DCS-200, this calculation assigns zero difference to the maximum value in each data set and is thus only approximate. The

error statistics reported below were calculated without using the maximum value in each data set, and without scaling the linearized output to the camera output scale of 0–255.

The average absolute value of the variation was 0.0015. The root-mean-square value of the variation was 0.0021, approximately 0.5 units per 256 (for comparison with the variation for the DCS-200) and the maximum value was 0.0072, approximately 1.8 units per 256. As can be seen from the plots in the top right panel of Fig. 3, the blue has most variation, and the red and green variations are comparable.

E. Noise Measurements

We took a number of dark images at different times during our day-long experiments, and at different exposure durations. The effect of aging is negligible and is detailed in [3]. We discuss the effect of exposure duration on dark current noise.

The data for the DCS-200 are tabulated in Table II. Dark noise shows some variation with exposure duration, up to 4 units, but is quite constant over the different color bands. The mean of the data tabulated is 13.61, 13.63 and 13.61 for red, green, and blue sensors respectively. The overall mean is 13.62. The variances for the three sensor types are 0.78, 0.79, and 0.81 respectively; the corresponding standard deviations are 0.88, 0.89 and 0.90. The overall variance about 13.62 is 0.79 with a standard deviation of 0.89. Variation is greatest for blue sensors and least for red, but the differences are slight.

The mean values may be compared to those obtained from the variation from linearity calculations in Appendix Section D. The value of 12.5 obtained there is close to the measured values. The variation values may be compared to the values obtained in Appendix Section D. The variation from linearity includes the dark noise variation, but is larger because it is not limited to the dark noise variation. It includes other nonlinear aspects of the sensor response, including other noise sources like shot noise.

The bottom panel of Fig. 3 illustrates the fact that the variation of dark noise with exposure duration is not monotonic at low exposure durations. This could be because of inaccuracy in the mechanics of the shutter movement. At exposure durations of 1/4 s and higher, the variation of dark noise with exposure duration is monotonically decreasing. This could be because the effects of dark current are averaged out at higher exposure durations.

The average value of dark current noise is usually subtracted from readings that are known to be linear, i.e., readings predicted by (1). As the DCS-420 sensor outputs are the result of a nonlinear function operating on the CCD measurements, [(2)], the dark noise average cannot simply be subtracted from the sensor readings. In fact, our calibration curve (the right panel in the last row of Fig. 2) provides an indirect estimate of $\mathcal{F}^{-1}(r) - n$ and the variability from this curve provides an estimate of the effective additive noise. Nonetheless, obtaining a direct measure of the dark noise variability seems useful for an estimate of acceptable errors in RGB prediction for camera calibration [4] (see Section F).

We took a few dark images (with the lens cap on) at various stages of the experiment, and at various exposure times. We calculated the average value over the same rectangle in the center field used for other measurements. The average value did

TABLE II
DARK NOISE VS. EXPOSURE DURATION, DCS 200

Exposure time in seconds	Average Dark Noise Value in Camera Output Units		
	Red	Green	Blue
8	12.52	12.58	12.51
4	12.85	12.83	12.84
2	12.92	12.93	12.94
1	13.54	13.56	13.51
0.5	13.90	13.92	13.83
0.25	14.27	14.32	14.32
0.125	14.44	14.47	14.45
1/15	14.31	14.32	14.33
1/30	14.54	14.54	14.54
1/60	14.54	14.56	14.60
1/125	11.90	11.88	11.87

not vary much. Its average over the different images was 25.02, 25.01, and 25.06 over red, green, and blue sensors, respectively. Its overall mean was 25.03. Individual variances about individual means were 0.2375, 0.2207, and 0.2203 for red, green and blue respectively. Its overall variance with respect to the overall mean was 0.2266 and the standard deviation 0.4760.

If we convert the dark noise standard deviation to the linear domain (using the average slope of the calibration curve) we get a value of 0.0019. This is a little lower than the measured deviations of the data from the curve. As with the DCS-200, the difference is explained by the fact that variation from the calibration curve includes the effects of other types of noise besides dark noise.

F. Camera Calibration

To perform calculations, we write a version of (2) that describes the entire calibration data set. Let \mathbf{r} , \mathbf{g} , and \mathbf{b} be vectors representing the R , G , B readings to a series of narrowband lights. The vectors \mathbf{r} , \mathbf{g} , and \mathbf{b} have K_r , K_g and K_b entries respectively, one for each of the narrowband stimuli used to calibrate the corresponding sensor. Let the full spectrum of the i th narrowband light be $s_i(\lambda)$, and let the unknown camera spectral sensitivities be $c_r(\lambda)$, $c_g(\lambda)$ and $c_b(\lambda)$. From (2) we have

$$\mathbf{r} = \mathcal{F} \left(\begin{bmatrix} e(1) \sum_j c_r(\lambda_l + j\Delta\lambda) s_1(\lambda_l + j\Delta\lambda) \Delta\lambda \\ \vdots \\ e(i) \sum_j c_r(\lambda_l + j\Delta\lambda) s_i(\lambda_l + j\Delta\lambda) \Delta\lambda \\ \vdots \\ e(K_r) \sum_j c_r(\lambda_l + j\Delta\lambda) s_{K_r}(\lambda_l + j\Delta\lambda) \Delta\lambda \end{bmatrix} + \mathbf{n} \right) \quad (3)$$

where

- \mathbf{n} vector representing measurement noise with variation about the average dark noise value;
- $\Delta\lambda$ wavelength sampling for the radiometric measurements;
- $e(i)$ exposure setting for the i th measurement.

The function \mathcal{F} is applied pointwise to each component of the vector it acts on. It is the identity for the DCS-200 and the cali-

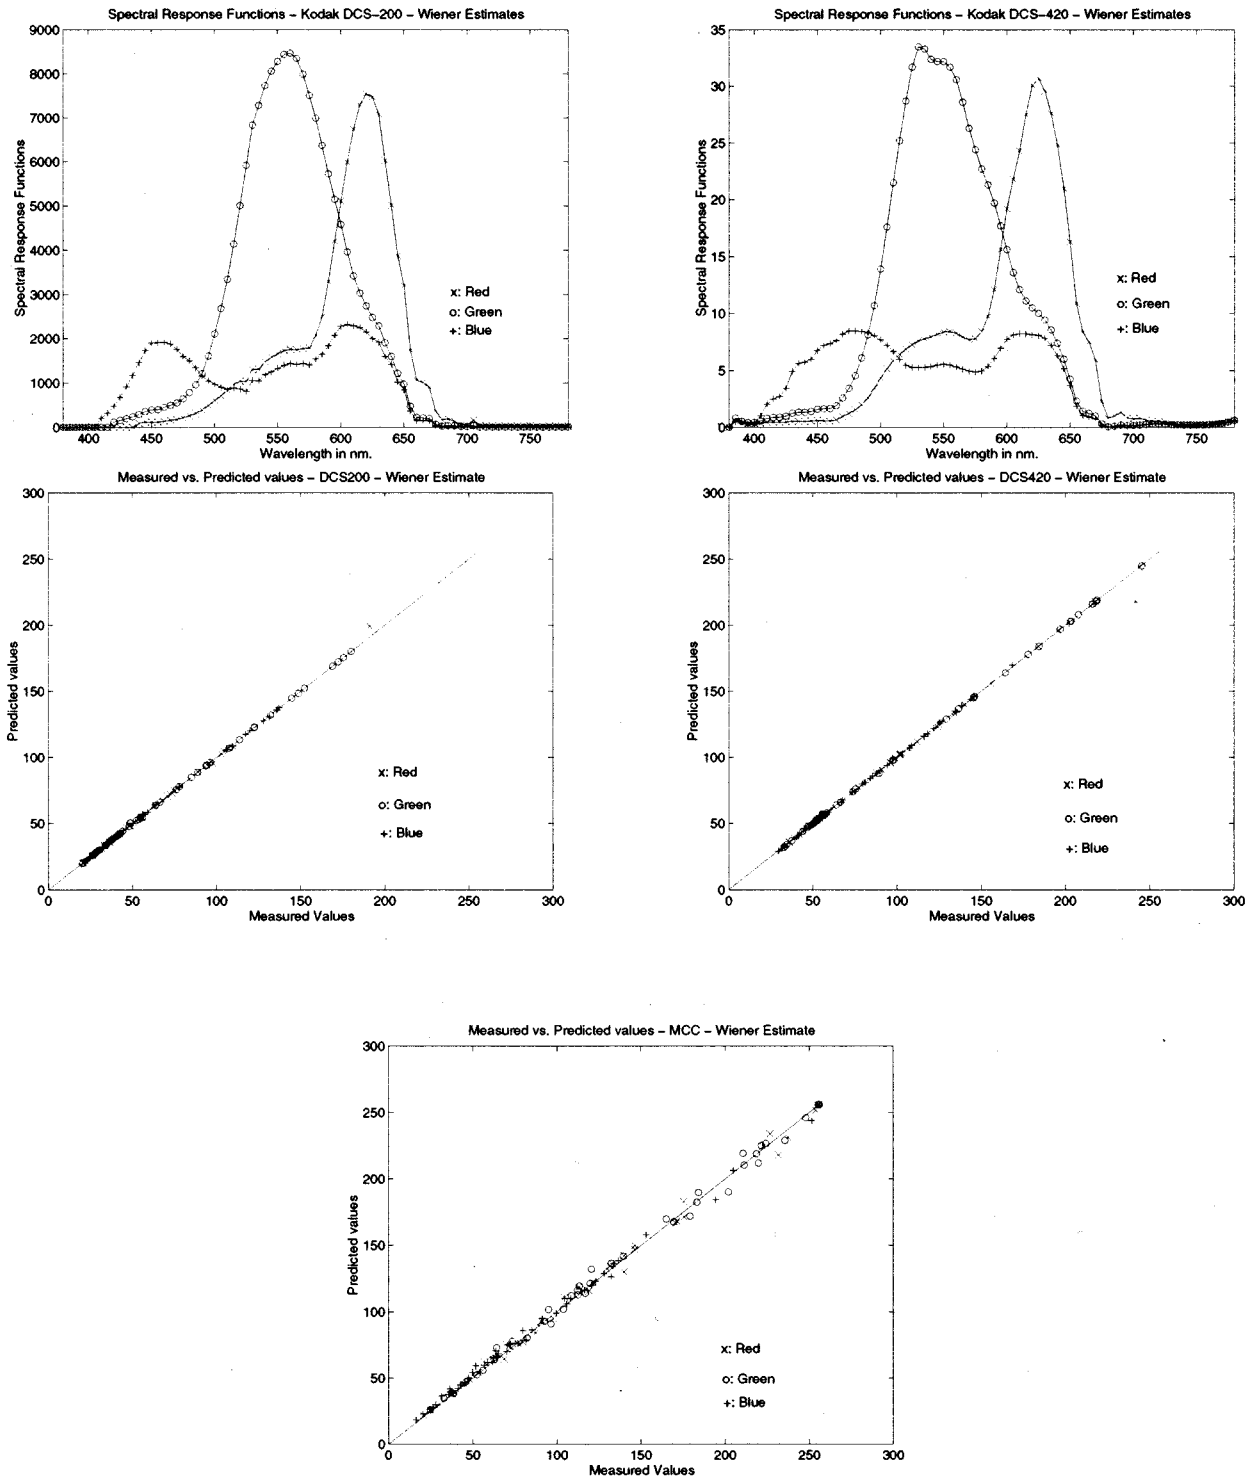


Fig. 4. Wiener estimates and simulator scatter plot.

brated static nonlinearity for the DCS-420. Equations similar to the one above can be written for the readings \mathbf{g} and \mathbf{b} . The equations for $c_r(\lambda)$, $c_g(\lambda)$ and $c_b(\lambda)$ may be solved in a number of different ways. In the rest of this section we discuss some possibilities.

We used a variant of the Wiener estimate that is guaranteed to produce all positive estimates. The Wiener procedure requires that we regard the quantity to be estimated (say \mathbf{x}) as a Gaussian random variable with known mean and covariance. We took the

mean $\bar{\mathbf{x}}$ to be the estimate obtained by assuming the spectral response is constant over the bandwidth of the narrow-band illuminant used for measurement and constructed the covariance matrix $\mathbf{K}_{\mathbf{x}}$ by assuming that \mathbf{x} was the result of a first-order discrete Gauss-Markov process whose variance was equal to the variance of the entries of $\bar{\mathbf{x}}$ and whose entry-to-entry correlation was equal to the correlation between neighboring entries of $\bar{\mathbf{x}}$. We assumed that the entries of \mathbf{n} were independently and identically distributed with mean zero and variance equal to 2% of the

TABLE III
DARK NOISE VS. EXPOSURE DURATION, DCS 200, LATER READINGS

Exposure time in seconds	Average Dark Noise Value in Camera Output Units		
	Red	Green	Blue
0.125	14.05	14.04	14.02
1/125	13.08	13.15	13.08

TABLE IV
STATISTICS OF ESTIMATION ERROR—DCS-200—WIENER ESTIMATES

Sensor Type	Mean of Absolute Error	Mean of Absolute % Error	RMS Error	Camera Noise Std. Deviation	Maximum Absolute Error	Maximum Absolute % Error
Red	0.53	0.65	1.42	0.88	8.07	4.22
Green	0.57	0.55	0.46	0.89	1.34	2.65
Blue	0.28	0.62	0.41	0.90	1.58	4.20

TABLE V
STATISTICS OF ESTIMATION ERROR—DCS-420—WIENER ESTIMATES

Sensor Type	Mean of Absolute Error	Mean of Absolute % Error	RMS Error	Camera Noise Std. Deviation	Maximum Absolute Error	Maximum Absolute % Error
Red	0.40	0.51	0.62	0.49	2.47	3.69
Green	0.35	0.55	0.46	0.47	1.34	2.65
Blue	0.47	0.66	0.59	0.47	1.50	2.63

TABLE VI
STATISTICS OF ESTIMATION ERROR FOR MCC—DCS-200—WIENER ESTIMATES

Sensor Type	Mean of Absolute Error	Mean of Absolute % Error	RMS Error	RMS of Variation in Patch	Camera Noise Std. Deviation	Maximum Absolute Error	Maximum Absolute % Error
Red	2.80	2.63	3.97	2.08	0.88	13.63	9.97
Green	2.85	2.41	4.27	1.96	0.89	11.70	13.63
Blue	2.81	4.68	3.60	1.83	0.90	9.95	15.80

maximum linearized sensor response (after correction for non-linearity, mean noise level, and exposure duration). We did this for each sensor and obtained results slightly better than those obtained by assuming the spectral response is fixed over the bandwidth of the narrow-band illuminants used for measurement. Graphs of the Wiener estimates for the DCS-200 and the DCS-420 are presented in the left and right panels of the first row of Fig. 4. The spectral response estimates are interpolated, and the numerical values may be found in [4]. Scatter plots of measured values vs. values calculated from the Wiener spectral sensitivity estimates for the DCS-200 and the DCS-420 are plotted in the left and right panels respectively of the second row of Fig. 4. Tables III and IV list the statistics of the estimation errors for the DCS-200 and the DCS-420, respectively.

In contrast to our previous attempts to estimate the spectral sensitivities of the Kodak digital camera [12], the error for the estimates are low and close to the rms value predicted by the noise statistics for both cameras. Presumably one factor driving

the small error is that we used many narrowband lights to calibrate the sensors.

Non-linear estimation methods like projections onto convex sets (POCS) [13] are used when the Wiener estimation method gives results that clearly do not satisfy prior knowledge of the solution. For example, POCS would be used if the Wiener estimates gave unreasonable errors in the RGB values. Our estimates satisfy the three known constraints: the set of measured and predicted RGB values agree leaving room for reasonable noise; the filters are reasonably smooth; the filter transmissivities are nonnegative. Hence, we did not attempt more complicated nonlinear (particularly constrained) estimation methods.

G. Verification of Estimates

For the DCS-200, we tested the spectral sensitivity estimates by collecting two images of the MCC under a tungsten illuminant. We compared the actual R , G , and B responses for the 24 color checker patches with values predicted from the spectral sensitivities of the camera and direct radiometric measurements of the light reaching the camera from each patch. The procedure is detailed in [4]. The last panel of Fig. 4 shows the predicted vs. measured values for the Wiener estimates. Table V lists the error statistics. It is clear that the estimates are excellent, and perform well on data sets that were not used for the calibration.

REFERENCES

- [1] T. S. Lomheim and L. S. Kalman, "Analytical modeling and digital simulation of scanning charge-coupled device imaging systems," in *Electro-Optical Displays*, M. A. Karim, Ed. New York: Marcel Dekker, 1992.
- [2] G. Sharma and H. J. Trussell, "Figures of merit for color scanners and cameras," *IEEE Trans. Image Processing*, vol. 6, pp. 990–1001, July 1997.
- [3] P. L. Vora, J. E. Farrell, J. D. Tietz, and D. H. Brainard, "Digital Color Cameras—1—Response Models," Hewlett-Packard Co., Tech. Rep. HPL-97-53, Mar. 1997.
- [4] —, "Digital Color Cameras—2—Spectral Response," Hewlett-Packard Co., Tech. Rep. HPL-97-54, Mar. 1997.
- [5] G. Brelstaff, A. Parraga, T. Troscianko, and D. Carr, "Hyper-spectral camera system: Acquisition and analysis," in *Proc. SPIE Conf. Geographic Information Systems, Photogrammetry, Geological/Geophysical Remote Sensing*, 1995, pp. 150–159.
- [6] M. A. Webster and J. D. Mollon, "Adaptation and the color statistics of natural images," *Vision Res.*, vol. 37, pp. 3283–3298, 1997.
- [7] D. H. Brainard. (1997) "Hyperspectral image data" [Online]. Available: <http://color.psych.ucsb.edu/hyperspectral/>
- [8] C. A. Parraga, G. Brelstaff, T. Troscianko, and I. R. Moorhead, "Color and luminance information in natural scenes," *J. Opt. Soc. Amer. A*, vol. 15, pp. 563–569, 1998.
- [9] D. L. Ruderman, T. W. Cronin, and C. Chiao, "Statistics of cone responses to natural images: Implications for visual coding," *J. Opt. Soc. Amer. A*, vol. 15, pp. 2036–2045, 1998.
- [10] "Links to hyperspectral & multi-spectral (sensor, algorithms, & data processing applications) resources on the web" [Online]. Available: <http://www.techexpo.com/WWW/opto-knowledge/ISresources.html>
- [11] M. J. Vrhel and H. J. Trussell, "Filter considerations in color correction," *IEEE Trans. Image Processing*, vol. 3, pp. 147–161, Mar. 1994.
- [12] P. M. Hubel, D. Sherman, and J. E. Farrell, "A comparison of methods of sensor spectral sensitivity estimation," in *Proc. IS&T/SID 2nd Color Imaging Conf.: Color Science, Systems, Applications*, vol. 48, 1994, p. 45.
- [13] H. J. Trussell and M. R. Civanlar *et al.*, "The feasible solution in signal restoration," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-32, pp. 201–212, 1984.



Poorvi L. Vora received the B. Tech. degree in electrical engineering from the Indian Institute of Technology, Bombay, in 1982, the M.S. degree in mathematics from Cornell University, Ithaca, NY, in 1990, the M.S. degree in electrical engineering and the Ph.D. degree in computer engineering from North Carolina State University, Raleigh, in 1988 and 1993, respectively.

She is with Hewlett-Packard Laboratories, Corvallis, OR. Her current interests are in the applications of cryptography to problems in the fields of privacy protection and digital asset commerce.



Joyce E. Farrell received the Ph.D. degree in experimental psychology and visual psychophysics from Stanford University, Stanford, CA, in 1981.

From 1981 to 1985, she conducted postdoctoral research on visual perception at several institutions, including Copenhagen University, Denmark, New York University, the NASA Ames Research Center, Stanford University, and the Xerox Palo Alto Research Center. From 1985 to 1999, she was with Hewlett-Packard Laboratories, Corvallis, OR, where she was responsible for developing metrics and methods for image quality evaluation. She is presently a Member of the Technical Staff at Shutterfly.com, Redwood City, CA.



Jerome D. Tietz received the B.A. degree in psychology from the University of Washington, Seattle, the M.A. degree in psychology from San Jose State University, San Jose, CA, and the Ph.D. degree in psychology from the University of California, Santa Barbara.

He is a Computer/Networking Technologist with the Psychology Department, University of California, Santa Barbara.



David H. Brainard received the M.S. degree in electrical engineering and the Ph.D. degree in psychology from Stanford University, Stanford, CA, in 1989.

He is Professor of psychology with the University of California, Santa Barbara. His research interests include both the study of human color vision and the study of algorithms for processing digital color images.