Colour-Imager Characterization by Parametric Fitting of Sensor Responses

M. Thomson,* S. Westland

Colour and Imaging Institute, University of Derby, Kingsway House, Derby DE22 3HL, United Kingdom

Abstract: This study describes a novel method for characterizing the colorimetric and photometric properties of three-channel color imaging devices. The method is designed to overcome some undocumented aspects of the imager-characterization problem: The effective spectral sensitivity profiles of the imager's color channels depend on the level of radiant input energy, and these profiles must be known in order to determine the true intensity-response characteristics of the three channels. By fitting the response distributions of the three color channels explicitly with low-dimensional models, the method takes these dependencies into account, and may, therefore, offer several advantages over other imager-characterization methodologies, particularly where illuminant-independent characterization is required. An application of the technique is detailed, in which a CCD camera is characterized using only the Macbeth ColorChecker and a number of artificial illuminants. © 2001 John Wiley & Sons, Inc. Col Res Appl, 26, 442–449, 2001

Key words: color-camera characterization; parametric fitting; sensor spectral sensitivity

INTRODUCTION

The ubiquity of consumer- and professional-grade three-channel CCD color cameras and scanners has motivated the development of several alternative methodologies for the colorimetric characterization of such devices.[1–3] These characterization procedures usually attempt to determine the relationship between digitized color images and their real-world counterparts in some multidimensional space, typically one of the standardized three-dimensional color spaces whose bases have been related to the cardinal axes of human color perception (e.g., CIE $L*a*b*$, CIE $L*C*h*$, CIE $L*u*v*$). For many applications, however, it is desirable to leave perceptual issues outside of the camera-characterization problem: The relationships between the properties of acquired images and those of the corresponding real-world scenes are then best described in some physical, rather than psychophysical, space. A purely physical characterization procedure must relate the digital values output by the imager to spectrophotometric measurements of the corresponding color signals and could thus operate in one of two ways: It may attempt to predict color signals given known imager outputs, or it may attempt to predict the imager outputs given known color signals.

The first of these two problems is in effect the problem of color-signal recovery. This is often considered the Holy Grail of device-independent imaging, but in its general form (i.e., unless illuminant and surfaces are subject to strong constraints) belongs to the category of inverse problems for which no unique solution is guaranteed. The second problem is that of (physical) imager characterization: In seeking to characterize a given imager, one can do no better than predict the imager's outputs for a given color signal of known spectral composition. If such a characterization can be achieved, perfect color-signal recovery will be possible for a special class of color signals: specifically, an intrinsically three-dimensional set of color signals whose bases (which could be determined through linear systems analysis[4]) are actually a linear combination of an imager's own bases. We use the term, imager bases, to denote the effective overall spectral sensitivity profiles of the imager; these can be usefully thought of as the imager's equivalent of the color-matching functions defined for the CIE standard observers.

The present study considers only the problem of imager characterization and begins by showing that it is not trivial: Predicting the imager's responses to known color signals involves estimating the imager's sensitivity as a function of

---

* Correspondence to: Dr. Mitch Thomson, Colour and Imaging Institute, University of Derby, Kingsway House, Derby DE22 3HL, United Kingdom (e-mail: m.g.a.thomson@derby.ac.uk)

wavelength and intensity, and the effects on the imager's outputs of varying these two parameters are not only nonlinear but interdependent. We then present a novel computational method designed to address these problems; the method is based on a nonlinear modelling technique that uses simple parametric equations to fit the imager bases. The method does not require the use of a monochromator (such devices are expensive, require frequent calibration, and are usually available only in a laboratory environment). It uses the Macbeth ColorChecker[5] as a characterization stimulus and assumes either prior knowledge of, or facilities to measure, the surface reflectances of the ColorChecker patches. Finally, we detail an application of the method to the characterization of a three-channel color camera.

BACKGROUND

At the heart of the imager-characterization problem lies the dual dependence of the imager's sensitivity to wavelength and intensity. The simplest approach would be to perform two separate characterizations, that is, to measure how the outputs of a given sensor class depend on intensity (a procedure often termed gamma calibration), then to measure how the outputs of a given sensor class depend on wavelength. In practice, however, these two parameters (intensity and wavelength) interact with one other, and this complicates the characterization problem as follows. Consider the process of determining the input–output characteristics of the three color channels: This involves measuring, for each channel, the relationship between output and some estimate of input, and the resulting curves are often assumed to show a power-law nonlinearity quantified by the power-law exponent or gamma. The output is, of course, the digital values produced by the imager; the issue here is how to obtain the input estimate. A true input–output characterization would use an estimate of the radiant energy input to a given color channel; yet this is unknown until the spectral sensitivity of that channel has been determined. Often, estimates of input energy are formed by simply integrating the color signal over the range of visible wavelengths, but these luminance-γ calibrations tend to be highly illuminant-dependent, as the following example illustrates. Consider a blueish-greenish swatch strongly illuminated by a tungsten source. The overall level of illumination might be very high, yet the blue sensor class would probably capture very little of the energy in the color signal. A plot of luminance against sensor output determined in this way would be a poor approximation of the true input–output characteristics of the blue sensor class; it would fail to predict the outputs of the same sensor class for the same swatch under, say, a daylight-tube source.

Another insight into the behavior of CCD-based imagers comes from careful consideration of the nature of the gamma-like input–output nonlinearity: The very slowly accelerating part of the curve near the origin means that the CCD operates, in effect, as a threshold device, since below a certain level of radiant energy input, only the dark current will flow. Thus, the overall spectral sensitivity of each channel is not determined purely by the relative transmittances of the three types of colored filter in the imager: These transmittance profiles are cut off by the finite absolute sensitivity of the CCD. Even without changes in the relative physical spectral sensitivities of either CCD silicon or filters, then one may observe changes in the shape of the overall effective spectral sensitivity of the device as the level of ambient illumination changes.

Given the interdependence of wavelength and intensity calibrations, the problem of camera characterization may be usefully thought of as the problem of having to map the interior of a 3-D space (sensor output as a function of wavelength and intensity) to sufficient accuracy such that the characterization is not rendered invalid by a change in imaging conditions, for example, by an illuminant change. The approach adopted in the present study is to fit the effective spectral sensitivity profiles explicitly with a low-dimensional parametric model, repeating the procedure at a variety of illumination levels. This is not the only way of fitting channel responses explicitly, but the adoption of a low-dimensional model makes the minimization algorithm so efficient that it is perfectly possible to improve the quality and reliability of the fitted parameters by increasing the amount of data used to estimate the bases. For example, it follows from the discussion above that some illuminants—those whose spectral irradiance distributions approach zero anywhere over the visible range—would make it impossible to map the wavelength-intensity space accurately (since the data available to any characterization algorithm would be inaccurate or nonexistent over certain ranges of wavelengths). Data from these illuminants can be combined with data collected under other illuminants *before* the modelling is conducted: The minimization need only be performed once to yield improved estimates of the imager bases.

The next section describes the algorithm in detail.

ALGORITHM

Theory

Suppose a single uniform surface of known spectral reflectance, S_λ, is captured under an illuminant with known spectral irradiance, I_λ, by three color channels whose spectral responses r_λ, g_λ, and b_λ are unknown; the energy entering the red channel, for example, will then be given by $\Sigma_\lambda I_\lambda S_\lambda r_\lambda$. Assume that the digital output of the red channel P_r is recorded and then normalized so as to correct for the camera exposure time to produce an exposure-independent output level L_r. Assume also that each channel is subject to a (possibly gamma-like, but unknown) input–output nonlinearity, which is represented by a function, f. Thus, one may predict the digital output of the red imager channel by the quantity, $f_r(\Sigma_\lambda I_\lambda S_\lambda r_\lambda)$. This expression makes the imager-characterization problem explicit: how to proceed when we have two unknown nonlinear functions (f_r, which is a function of input energy, and r_λ, which depends on wavelength). The approach followed here is to collapse the

two nonlinearities into one, assuming, again for the case of the red color channel, that there exists a whole family of nonlinearities R—indexed by both wavelength λ and output level L_r—such that

$$L_r \approx f_r(\sum_\lambda I_\lambda S_\lambda r_\lambda) \approx \sum_\lambda I_\lambda S_\lambda R_{\lambda,L_r}.$$

Similar families of nonlinearities, G_{λ,L_g} and B_{λ,L_b}, are assumed to exist for the green and blue channels, respectively, and it is exactly the union of these entities—each of which represents a set of effective spectral sensitivity profiles for the device as a function of output level—that we use to model the bases of the imager.

The task of camera characterization is now reduced to estimating these bases. Assume that the output level L_r of the red imager channel is recorded, and that a suitable low-dimensional parametric model can be found for each unknown red basis R_{λ,L_r}. Then the task of determining R_λ at a given output level L_r can be identified as a nonlinear modelling task in which the parameters of the model are varied so as to minimize the error term E_r over the visible range of wavelengths λ:

$$E_r = L_r - \sum_\lambda I_\lambda S_\lambda R_{\lambda,L_r}. \qquad (1)$$

Of course with data available from only one surface, the fitted model parameters would be meaningless, since it would be impossible to disconfound the effects of intensity from those of spectral shape. A multiple-surface calibration stimulus, however, can be used to increase the degrees of freedom (d.o.f.) in the data used to estimate the color-channel response distributions. The Macbeth ColorChecker[5] comprises 18 uniform chromatic surfaces drawn from the Munsell[6] set as well as six graduated achromatic surfaces. If the spectral reflectance of each chromatic surface is known, and if these surfaces are digitized to produce a mean digital output value L_r for each surface, 18 different measures of reconstruction error E_r can be obtained (the achromatic surfaces are not used since they would not increase the d.o.f. in the data used to estimate the response distributions); simultaneous minimization of these error terms will produce estimates of the fitted model parameters. The procedure can then be repeated, using the same calibration stimulus, for blue and green color channels.

Parametric Model

A key issue is the selection of an appropriate parametric model for the response distributions of the three sensor classes. Published data on the response characteristics of color CCDs suggest that although they are generally smooth, bandlimited, uni- or, possibly, bi-modal functions, their response distributions may be asymmetric with respect to the wavelength axis.[7] In addition, some imaging devices make use of matrix-mixing techniques to transform the CCD sensor spectral absorption curves into a more efficient basis for color representation; this can produce negative lobes in the response distributions of the color channels. Set against these considerations, however, is a desire to keep the parametric model as low-dimensional as possible, so that noise in the digital output data does not alter significantly the fitted parameters. With this in mind, the response characteristics of the color channels were fitted with basis functions of the Gram-Charlier expansion[8]; this probability law uses Hermite polynomials to describe departures from pure gaussian shape. Only the first two Hermite polynomials were used: Each response distribution was treated as possibly skewed, possibly kurtosed Gaussian functions. In total, the model had five parameters: the peak wavelength, amplitude, and width of the Gaussian, together with the skewness and kurtosis terms. If these parameters are denoted as p_r, a_r, w_r, s_r, k_r respectively, then, taking the case of the red sensor class as an example, the equation used to fit R_λ at a given value of L_r is:

$$R_\lambda \approx (s_r\lambda + 2k_r\lambda^2)a_r\exp\left(-\left(\frac{\lambda - p_r}{w_r}\right)\right). \qquad (2)$$

Minimization

Nonlinear fitting was performed using the Levenberg–Marquardt (L–M) technique,[9] which uses a mixture of first- and second-derivative methods to minimize the error between actual and reconstructed pixel values. The algorithm implemented here differed from standard L–M in two ways: First, the derivatives of the fitting function w.r.t. the fitting parameters had to be evaluated by means of finite-differencing, since these depend on the spectral distributions of the color signals and are, therefore, not known analytically. In addition, the fit was a multidimensional one, attempting to minimize reconstruction error E_r simultaneously for all 18 chromatic ColorChecker surfaces (or, optionally, for combined data obtained using more than one illuminant).

The L–M technique is iterative and requires initial estimates of all fitted parameters. Manufacturer's data on the camera's CCD absorption spectra were not available to us; accurate data on the spectral properties of CCD chips are rarely available and, as discussed above, may be a relatively poor predictor of the overall response properties of a given color channel. Moreover, a useful index of the robustness of the model is to test for any dependency on the initial parameter estimates of the algorithm's convergence and the final parameter estimates. With this in mind, three alternative sources were used for initial estimates of the position, width and amplitude terms (skewness and kurtosis terms were always initialized at zero): data on the spectral sensitivities of human cones,[10] published data on typical CCD sensor response properties,[7] and naive estimates, which assumed that the sensor distributions were gaussians of unit amplitude, 100 nm space constant, and positioned at the center (green, 550 nm) and extrema (red, 700 nm; blue, 400 nm) of the visible range. After convergence, the final estimates of the fitted parameters and the residual sum of squares were found to be the same for all choices of initial parameter estimates; thus, the model appeared insensitive to

major changes in the initial estimates of the fitted parameters.

One further source of information was also incorporated into the fitting algorithm: In addition to measuring the mean digital output associated with each ColorChecker patch, it was also possible to estimate the standard deviation associated with the spatial distribution of digital outputs across that patch. This measure was used to weight the minimization of reconstruction error such that surfaces captured with relatively less noise were relatively more important in determining the fitted parameters.

Procedure

The characterization procedure may be summarized as follows:

1. Each image of the entire ColorChecker was automatically segmented into its 24 constituent patches.
2. The RGB digital outputs of the imager in response to each ColorChecker patch were corrected for exposure time and neutral-density (ND) conditions, then normalized to maximum intensity (65535).
3. The mean and standard deviation of the spatial distribution of these digital values across each patch were computed and stored in a database alongside the corresponding color signal data. Recall from "Background" that a true intensity-response characterization is impossible without estimates of the imager bases, which in turn may vary with the level of illumination. If the characterized imager is to be useful for color-signal recovery, the procedure must therefore produce estimates of the imager bases as a function of imager output. This can be achieved by means of a lookup table in which the fitted parameters p, a, w, s, k produced by the minimization algorithm are indexed by output level, which is constructed as follows.
4. Considering the case of the red imager channel, denote a possible digital output value by M_r. For each of the 18 chromatic patches in the ColorChecker, the database was searched to find the conditions under which that patch gave rise to red-channel digital output values L_r most similar to M_r; the imager-output and color-signal data for that patch under those conditions were then saved. Repeating the process for the other 17 patches produced a single dataset capable of characterizing the red imager basis at an output value of M_r.
5. This single dataset was fed into the L–M algorithm, which then converged to produce estimates of the fitted parameters of the red image basis at the given output value M_r, i.e., estimates of the parameters of R_{λ, M_r}.
6. Steps 4 and 5 were repeated for 40 different values of M_r equally spaced over the 12-bit output range of the camera.
7. Steps 4–7 were repeated for the green and blue imager bases. The result of these stages was a lookup table which stored the fitted parameters p, a, w, s, k of the imager bases as a function of M_r; thus, the imager bases

can now be reconstructed for any given output level. What still remained was to estimate the input–output characteristics of the three color channels, and these were determined as follows.

8. A new database was constructed to store the imager-output and color-signal data for the six achromatic ColorChecker swatches under every imaging condition.
9. For each digital output value in the achromatic database, the lookup table of fitted parameters was consulted to find the best parameter set for that digital output value (i.e., the parameter set obtained at the output level closest to the output value read from the database).
10. These fitted parameters were used to reconstruct the three imager bases.
11. These bases were multiplied by the known color signals of the achromatic patches; the resulting data were summed across the visible range of wavelengths so as to produce estimates of the relative energy inputs to the three color channels;
12. These inputs were plotted against the digital outputs recorded for the six achromatic ColorChecker swatches, producing estimates of the true input–output characteristics of the three color channels.

Characterization was then complete, providing estimates of (a) the parameters fitted to the imager bases as a function of imager output and of (b) the true input–output characteristics of the bases. For validation purposes, a reconstruction algorithm was also designed, with the aim of predicting the imager's outputs in response to an arbitrary, known color signal. Given digital output values, the reconstruction algorithm uses these to select the best estimates of the fitted parameters, reconstructs the three imager bases, multiplies these by the measured color signal, then corrects the result for the input–output characteristics of each channel.

APPLICATION

This section describes an example application of the technique to the characterization of a high-spatial-resolution three-channel color camera.

Equipment

Measurements were performed in a purpose-built lighting booth fitted with three illuminants with spectral irradiance distributions similar to those of CIE-defined sources: i1, a daylight tube (CIE illuminant F11, correlated color temperature 4000 K); i2, tungsten bulbs (CIE illuminant A, correlated color temperature 2856 K); and i3, a blue fluorescent tube (CIE illuminant F6, correlated color temperature 4150 K). These were switched on 1 h before the calibration procedure to minimize flicker noise, and tests were conducted to ensure that the illumination of the ColorChecker (Macbeth [Kollmorgen]) was spatially uniform. The relative positions of the ColorChecker and camera were adjusted so that the imaging geometry was 45/0. The camera was a Leaf

Lumina (Scitex), with spatial and chromatic resolution of 2700 × 3380 pixels and 12 bits-per-channel, respectively. A Spectrascan SV650 spectroradiometer (Spectravision) was used to measure the color signal of each ColorChecker swatch; a barium sulphate 99% reflectance tile was used to measure the spectral irradiance distributions of the illuminants. The ColorChecker was digitized at 40 different exposure levels and used three neutral-density (ND) filters, ×1, ×2, and ×4, producing a total of 120 conditions for each illuminant. The digitization and spectrophotometry were then repeated for the other two illuminants.

Images of the entire ColorChecker were captured and stored on an Apple Power Macintosh, but were transferred to a Sun SPARCstation (Sun Microsystems) for subsequent processing.

Results

Fig. 1 shows examples of the imager bases recovered by the algorithm under three different conditions: low, intermediate, and high output levels. Notice how the shapes of the imager bases are not independent of output level: Changes in channel bandwidth and amplitude can be quite significant (particularly for the green channel).

Fig. 2 shows the intensity-response characteristics of the three bases. The abscissa shows recorded digital output values for the six achromatic ColorChecker swatches under a number of different exposure/ND conditions; the ordinate shows the channel energy inputs estimated by combining the corresponding color signals with the appropriate channel

FIG. 2. Input–output characteristics of the three color channels under each of the three illuminants i1, i2, i3. Since the data for all three channels would otherwise overlie each other, the data for both blue (top) and green (middle) channels have been shifted up the ordinate.
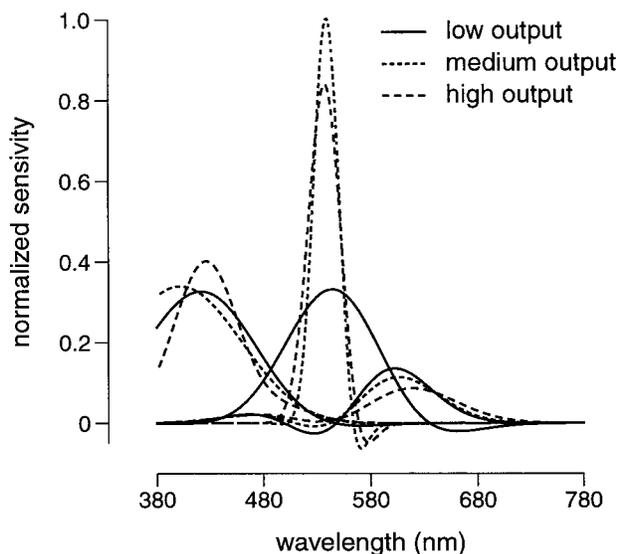

FIG. 1. Estimates of the relative spectral sensitivities of the three sensor classes at low, medium, and high pixel intensities. These data were derived by fitting the parametric model described in the text to measurements of the Color-Checker swatches under illuminants i1, i2, and i3. Data like these were obtained for 40 different output levels; the three output levels shown here represent the mid-range and extremal output levels, and variations in the shape of these functions across output levels were smooth.
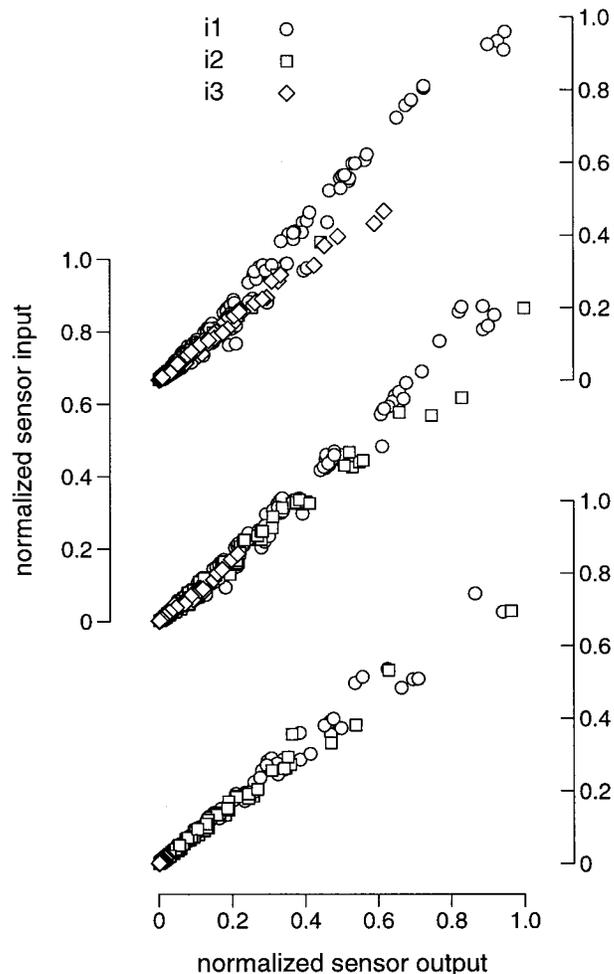
response distributions. Data are shown for all three illuminants and for each color channel, and two points are worthy of note. First, although there is some scatter toward the higher output values, the input–output characteristics are close to linear. The reason for this is that the nonlinear component of the dependence of the camera channel outputs on their inputs appears instead as changes in the amplitudes and shapes of the bases (Fig. 1). Second, although the ranges of output values are different for each channel, the trends obtained under different illuminants overlie each other; these input–output characteristics are illuminant-independent. This behavior can be usefully compared with that of input–output estimates obtained using the standard luminance-γ technique, i.e., by following steps 8–12 in "Algorithm" but without taking the channels' spectral sensitivities into account (the input energy measures are simply the integrals of the color signals over the visible band). The abscissa of Fig. 3 show digital outputs recorded for the six achromatic swatches at an intermediate exposure level (ND × 1), and the ordinate shows the channel energy inputs estimated by integrating the corresponding color signals.
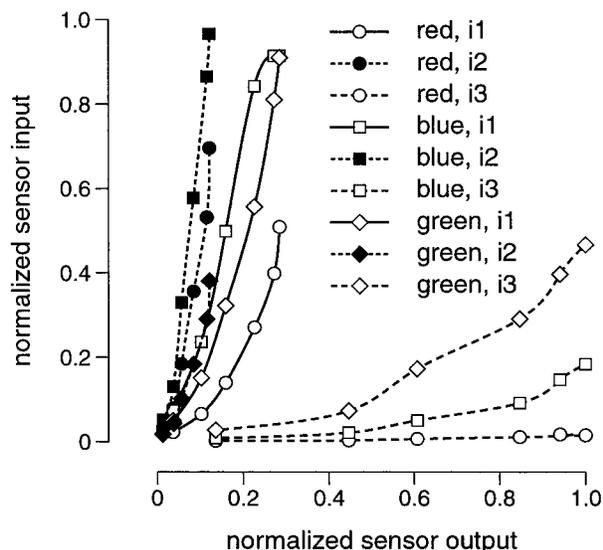
FIG. 3. Input–output characteristics of the three color channels using the luminance-γ technique. Data were obtained using each of the three illuminants i1, i2, i3, under a single exposure and ND condition.

Again, data are shown for all three illuminants and for each color channel. Notice that the input–output curves now exhibit marked nonlinearities. More important, however, is the effect of changing the illuminant: Different illuminants produce completely different curves, so a system characterized in this manner would perform poorly if not recharacterized following an illuminant change.

Validation

There is no guarantee, of course, that the estimated bases shown in Fig. 1 are in fact correct: They could be poor estimates of the true functional sensitivities of the imager. Validation of the procedure ultimately requires that predicted camera outputs be compared with actual camera outputs. More specifically, since all the information supplied to the characterization algorithm is collapsed into a low-dimensional space—only five parameters are determined for each imager basis—there is no guarantee that the model is accurate enough to permit good predictions of the camera's outputs under all imaging conditions. Thus, even though only ColorChecker data were used to estimate the imager bases, circumstances may arise under which the validation algorithm will not accurately predict the camera outputs for each ColorChecker patch.

Fig. 4 shows the results of attempts to predict the digital outputs of the camera to each of the digitized ColorChecker patches. These data were obtained at an intermediate exposure level, ND × 1, and the error bars on the ordinate refer to the standard deviation of digital output values across each patch. For the standard predictions, the L–M algorithm was allowed to use data captured under all three illuminants (i1, i2, i3) to estimate the imager bases. For comparison purposes, two alternative predictions were attempted: In the single-illuminant prediction, the algorithm used data cap-

tured under illuminant i1 only, and in the luminance-γ condition, the standard gamma-calibration technique was used (i.e., the input–output characteristics were the same as those used to generate the data in Fig. 3).

The plots in Fig. 4 show the relationships between the actual Colorchecker data (ordinate) and each of the three types of prediction (abscissa); if a prediction algorithm was 100% successful, all data should lie along the line $y = x$, shown as a dotted line on the plots. For the blue (top left) and green (top right) channels, all three types of prediction fit the actual camera outputs rather well. For the red sensor class (bottom left), however, the single-illuminant reconstruction fails badly. More detailed examination of the look-up-table data revealed the reason for this: There was simply not enough energy in illuminant i1 at the red end of the spectrum to recover accurately the relative spectral sensitivity of the red sensor. In effect, this particular area of the wavelength-intensity space described earlier was rather poorly mapped. The traces in the remaining plot (bottom right) show a different kind of failure: They show how calibrations performed under a single illuminant may provide rather poor characterisations under a second, different illuminant. The data were obtained for the green sensor class under illuminant i2; notice that both the single-illuminant or luminance-γ predictions provide rather poor fits, although the standard prediction still works well under this second illuminant. In fact the range of correlation coefficients (actual v. predicted) measured for the standard predictions over all three illuminants and all three color channels was 0.94–0.97.

DISCUSSION

The simulations presented in the previous section show that illuminant-independent calibration of three-channel imagers can be achieved by means of the parametric fitting of the imager's bases. Three factors appear particularly important contributors to the accurate estimation of the bases: using more than one illuminant, modelling changes in the bases with the level of illumination, and taking into account the effect of the spectral characteristics of the bases on the intensity-response calibration.

Methods that determine explicitly the response characteristics of the three channels appear to offer advantages over methods that determine a system transfer function in some arbitrary multidimensional space. For example, Lee[1] derives an equation similar to Eq. 2 for each ColorChecker swatch under a single illuminant and shows how these can be combined to form a linear system; this is then solved by standard matrix techniques to yield estimates of the system transfer at a number of discrete wavelengths. With only 18 chromatic patches, however, one can only solve nominally for 18 unknowns, which produces a rather coarse sampling in wavelength space. To make matters worse, row degeneracy will occur in the solution matrix if one or more of the equations is a linear combination of the others, an outcome that could well be obtained given the low intrinsic dimensionality of the Munsell samples that make up the Color-
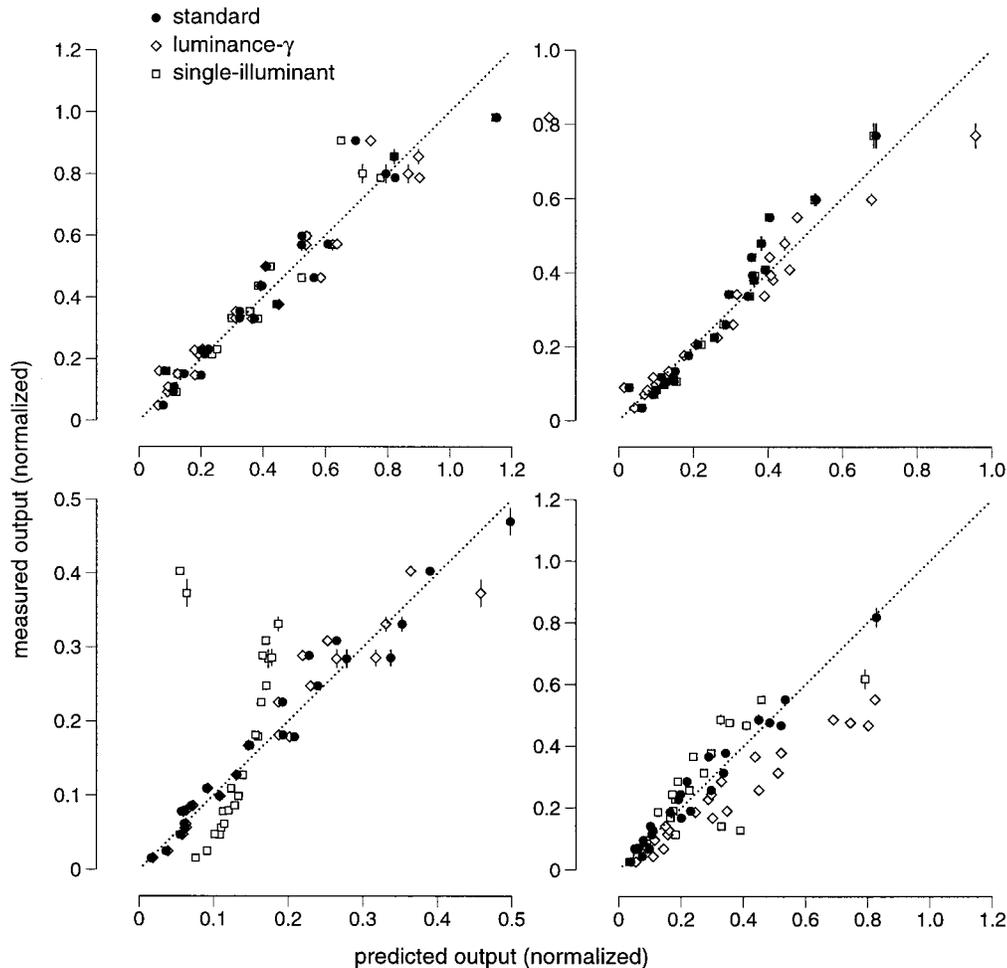
FIG. 4.   Measured and predicted pixel intensities of ColorChecker swatches for red (bottom left), green (top right) and blue (top left) color channels; measurements and predictions were both under illuminant i1. The remaining plot (bottom right) shows an attempt to predict the green-channel pixel intensities under illuminant i2. In all four plots, the ordinate shows the values recorded by the camera, with error bars showing the standard deviation of the spatial noise (these are mostly so small as to be invisible). The three symbols denote three different predictions: a standard triple-illuminant prediction, a single-illuminant prediction, and a triple-illuminant luminance-γ prediction. The data have been sorted along the abscissa in order of increasing pixel intensity in each case.

Checker swatch set (it is largely accepted that the set of Munsell surfaces can be accurately represented using only three slowly varying principal components; see, e.g., Cohen[6]). In practice, our own experience with the matrix approach is that under some illuminants, attempts to solve for more than four or five unknowns produce a near-singular matrix and result in gross errors in the estimated system-transfer-function coefficients.

Recent advances in the matrix technique have been made by Finlayson and Hordley,[11] who impose constraints of positivity, modality, and band-limitedness on the solutions of the linear system. These constraints appear to render the problem tractable in situations where the d.o.f. of the fitting matrix would otherwise be too low to solve the linear system. The method set out in the present study achieves a similar robustness, but by imposing constraints in a slightly different way: through the adoption of a low-dimensional parametric model. Data are not available to permit a direct comparison of the techniques, but the low dimensionality of the parametric model confers several additional advantages.

First, it is reasonable to expect that a model having so few degrees of freedom would be particularly robust in the presence of noise. Second, since it is possible to track explicitly the dependence of the fitted model parameters on the imaging conditions (choice of illuminant, level of illumination, etc.), one may determine rather easily whether the imager is performing within the operating range appropriate to a particular application. Third, as discussed above, having so few parameters in the model itself allows color-signal data obtained under more than one illuminant to be used to obtain illuminant-independent estimates of the imager bases. Since different illuminants project the surface reflectances of the ColorChecker swatches onto different signal subspaces, using additional illuminants can increase the d.o.f. in the data used to estimate the fitted parameters of the color-channel distributions, and this may be particularly useful when, as here, the surface reflectances of the calibration data are very smooth.

Finally, although the primary motivation for determining imager bases is in characterizing imager performance, hav-

ing reasonably accurate estimates of these bases is useful in other ways. The bases can be thought of as the axes of the space in which the imager operates, and this allows us to define the set of colors perfectly reproducible by the imager (those which could be synthesized from a linear mixture of the three imager bases). Algorithms for performing illuminant recovery also require that the imager bases be accurately specified.

CONCLUSION

This study has described a novel method for characterizing three-channel imagers by means of a parametric fitting technique. The method does not require the use of a monochromator; it assumes only prior knowledge, or facilities to measure, the surface reflectances of the ColorChecker swatches. The calibration algorithm can be performed as often as required and at a variety of exposure levels by the end user on any three-channel imager with reasonably well-behaved color-channel response distributions. The technique works by first fitting the response distributions of the three color channels with bases of the Gram-Charlier expansion, taking into account changes in channel shape as a function of intensity. The resulting estimates of channel spectral sensitivity are then used to compute the true input–output characteristics for each channel. These characteristics, together with the estimated response distributions, can be used to predict the camera's responses to data whose color signals are known. A useful feature of the technique is that it allows data obtained under multiple illuminants to be combined for the purposes of estimating the response distributions; an example application of the technique shows that this may be critical in producing illuminant-independent characterizations. The technique may, therefore, have particular applicability to those multispectral imaging methods that are based on multiple-illuminant image capture.

1. Lee R Jr. Colorimetric calibration of a video digitizing system: Algorithm and applications. Color Res Appl 1998;13:180–186.
2. Dalton CJ. The measurement of the colorimetric fidelity of television cameras. J Inst of Electronic and Radio Engineers 1998;58:181–186.
3. Strachan NJC, Nesvadba P, Allen AR. Calibration of a video camera digitising system in the CIE $L*u*v*$ color space. Patt Recog Lett 1990;11:771–777.
4. Maloney LT. Evaluation of linear models of surface spectral reflectance with small number of parameters. J Opt Sc Am A 1986;3:1673–1683.
5. McCamy CS, Marcus H, Davidson JG. A color rendition chart. J Appl Photog Eng 1976;2:95–99.
6. Cohen J. Dependency of the spectral reflectance curves of the Munsell color chips. Psychon Sci 1964;1:369–370.
7. Engelhardt K, Seitz P. Optimum color filters for CCD digital cameras. Appl Opt 1993;32:3015–3023.
8. Frieden BR. Probability, Statistical Optics and Data Testing. Berlin: Springer-Verlag; 1983.
9. Marquardt DW. J Soc Ind Appl Math 1963;11:431–441.
10. Wysecki G, Stiles WS. Color Science: Concepts and Methods, Quantitative Data and Formulae. NY: Wiley; 1982.
11. Finlayson G, Hordley S, Hubel PM. Recovering device sensitivities with quadratic programming. Proceedings of the 6th color imaging conference Scottsdale, Arizona: Color science, systems, and applications, 90–95; 1998.