A Novel Method for Colorimetric Calibration of Color Digitizing Scanners

J A Stephen Viggiano and C Jeffrey Wang
Imaging Division
RIT Research Corporation

R · I · T Research Corporation

A paper presented at TAGA's 1993 Annual Technical Conference
Minneapolis, Minnesota, 25 - 28 April 1993

Cite this paper: J A Stephen Viggiano and C Jeffrey Wang, A Novel Method for Colorimetric Calibration of Color Digitizing Scanners, *1993 TAGA Proceedings,* p 143 - 160.

A Novel Method for Colorimetric Calibration of Color Digitizing Scanners

J A Stephen Viggiano,* C Jeffrey Wang*

A novel method for calibrating desktop color scanners so their output may be converted into device-independent color coordinates, such as CIELAB, is derived. It is a combination of models for the color-related performance of color photographic materials and simple scanners.

Our method is based on a scanner-specific characterization of the scanner's response to a series of neutral objects, a material-specific characterization of the dyes and color patches in a test target, and a characterization of the interaction between the scanner and the material. By separating these three components, duplication of effort is minimized, so more efficient calibration should be possible.

Once the scanner's response to a series of spectrally neutral objects has been determined, the calibration can be performed with as few as four patches on the material of choice. A detailed description of the new technique is provided.

Introduction

Color digitizing scanners, which produce Red-Green-Blue (RGB) output, are important components of desktop color publishing systems. Compared to high-end scanners, they offer advantages in compactness, lower price, and, allegedly, in ease of use. They would seem an attractive way to obtain digital color. Their popularity should continue to increase dramatically. However, they can present some special problems in obtaining good color. This paper discusses a technique for colorimetrically calibrating such devices, enabling better quality results. The aim is to convert the RGB values of a scanned image to the XYZ tristimulus values of the original transparency or print being scanned.

[1] RIT Research Corporation, Imaging Division

One could ask, "What is to be gained by colorimetrically calibrating my scanner? I don't care what tristimulus values correspond to each point on my original, so why be bothered?" The advantages of having such a calibration are manifold, and include:

- Ability to exchange picture data with a variety of systems
- More consistent, quality results
- "Open" versus "Closed" formats are possible
- Consistent format for ease of interfacing with both existing and future systems
- Accuracy of results
- Elimination of scanner metamerism

Overview of Technique

The calibration technique described in this paper is designed to work best with scanners having narrow spectral responses. If something is known about the material being scanned, the RGB values may be used to reconstruct the spectral curve of the pixel being scanned. This spectral curve may be integrated to obtain the tristimulus values X, Y, and Z. Figure 1 is a data flow diagram of the calibration model.

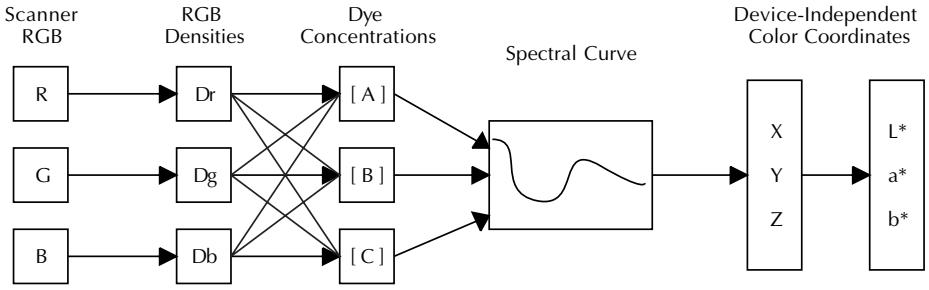

Figure 1.
Flow of data through the calibration model.

The first step involves calibrating the scanner so that the RGB output values may be converted to Red, Green, and Blue densities; a set of empirical equations may be used. These densities are converted to dye concentrations via a linear transformation similar to the masking equations. Once the dye concentrations are known the spectrum of the object being scanned is computed. Finally, the spectrum is integrated to yield the tristimulus values X, Y,

and Z, from which other color space coordinates, such as CIELAB or calibrated RGB, may be computed.

Empirical Approaches in Prior Art

Other approaches to colorimetric calibration of scanners have been taken. Examples of some empirical approaches involve multi-dimensional polynomials [1] and multi-dimensional interpolation. [2] These approaches, though flexible (and hence well suited for quick results from new technologies), suffer from some serious drawbacks. Polynomial models are susceptible to undesirable local minima and maxima; the tendency increases with the order of the model. Gerald offers some dramatic examples on data containing cusps, [3] while Scott and Sylvestre of Kodak provide an example of undesirable local maxima while fitting a very smooth (infinitely differentiable) monotonic curve, even with second- and third-order polynomials. [4] (See Figure 2.)

The sample-and-interpolate method is susceptible to local error caused by an anomalous data point. A large sample size is also required. Spooner recommends a minimum of 1070 samples for calibrating hard copy output using this approach, [5] and this figure is predicated on an optimal, non-uniform and non-orthogonal spacing. For a uniformly spaced, orthogonal sampling of input space, Spooner recommended 7000 samples. Neither the polynomial nor the sample-and-interpolate technique is suitable for extrapolation.

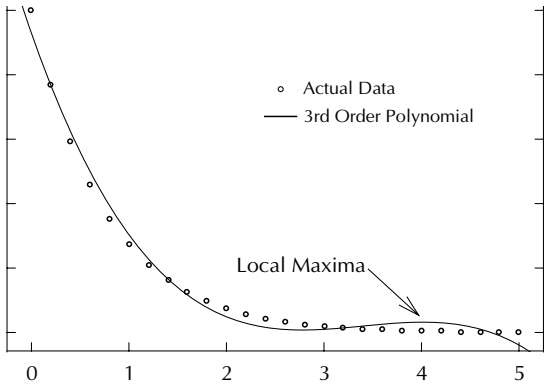

Figure 2.

A local maxima in a least-squares polynomial fit is illustrated. The original data are monotonically decreasing, but the polynomial exhibits a local maxima near X = 4. In an imaging application, this could result in a contrast reversal and false contours.

Analytical Models

At RIT Research, we have found that analytical models, based on careful consideration of the physical attributes of an imaging system, offer considerable advantages over empirical ones. These include:

- Less calibration data required
- Ability to extrapolate beyond calibration data set
- Simplified re-calibration if conditions (*e.g.,* exposure or contrast settings) change
- Less sensitivity to local errors
- Smoothness and monotonicity of results
- Ability to efficiently reduce effect of measurement error through averaging several replicates

Disadvantages of analytical models are that more work is required up front (some cogitation is required), and there is no guarantee that the analytical model will be more accurate than an empirical one. The first consideration means that an investment must be made, but our experience is that the ease of re-calibrating an analytical model makes for an extremely quick payback period. If the second situation prevails, it usually indicates that the modeling work is not completed.

A Novel Method for Calibration of Color Digitizing Scanners

In order to construct an analytic model of a device's color response, we have found it extremely useful to handle two components separately: the device's *amplitude response,* and its *color mixing model.* For scanners, the color mixing model is comprised of the scanner's spectral sensitivities and the spectral characteristics of the material being scanned. These components, as they apply to color digitizing scanners, are discussed in detail in the sections which follow.

Our calibration technique consists of four components:

- A *scanner-specific, material independent* component, which involves the amplitude response of the scanner;
- A *material-specific, scanner independent* component, which characterizes the dyes used in the photographic originals being scanned;

- A *scanner- and material-specific* component, which addresses the interaction between scanner and dyeset; and

- An *implementation-specific* component, in which multi-dimensional lookup tables are constructed.

Scanner-Specific, Material Independent Characterization
There are two aspects of the scanner which may be determined independently of the material being scanned. These are the scanner's spectral sensitivities, and the scanner's response to a series of spectrally neutral objects (which we refer to as the scanner's *amplitude response).*

These steps may be performed without knowledge of the material being scanned. This offers a number of distinct advantages:

1. The process need not be repeated for each material for which the scanner is to be calibrated.

2. It is easy to repeat these steps for various scanner settings (exposure, contrast, etc.) without having to perform a complete recalibration.

3. These steps need not be repeated if new materials become available after the original calibration effort.

Scanner Spectral Sensitivity: In our method, the spectral sensitivities are assumed to be narrow in bandwidth. Scanners are often designed this way in order to obtain improved saturation in lieu of negative sensitivity lobes. (See the discussion of the Ives-Abney-Yule compromise in Hunt.) [6] Further, many color digitizing scanners use fluorescent light sources. Fluorescent lamps have prominent mercury lines at 436 and 546 nanometers. This also promotes a narrow-band spectral sensitivity in the Green and Blue bands.

Other than the assumption of narrow bandwidth, our technique does not require specific determination of spectral sensitivity.

Scanner Amplitude Response: Suppose we scan a gray scale for which we know the reflectance of each step. Will the numbers used to represent the level of gray be proportional to the reflectance once the gray scale is scanned in? Or will they be proportional to the density of the step which they represent? We refer to this as the scanner's *amplitude response,* and have found it efficacious

to study it separately from the spectral sensitivity. In our scanner calibration technique, the amplitude response forms the scanner-specific, material independent component.

It may be measured by scanning a series of spectrally non-selective tiles, and determining the average response to each. We have a set of reference tiles from the National Bureau of Standards (now NIST), but virtually any set of objects having reasonably flat spectra may be used, provided they have similar gloss and translucency to the material being scanned.

Because images are compared, manipulated, edited, and corrected while viewed on a CRT monitor, many scanners incorporate an amplitude response function which attempts to properly render shades of gray when displayed on the CRT.* This is sometimes referred to as "gamma correction." In its simplest form, it would simply mimic the CRT's amplitude response function, so that the two amplitude response functions would cancel. However, it is also necessary to compensate for over- and under-exposed originals, color casts, and other undesirable features of the original.

Offered as an example of an amplitude response function for digitizing scanners, adapted from models for the amplitude response functions of CRT monitors, [7] is:

$$Dr = -\log\left|\, a + b \cdot \left(\frac{R}{255} + c\right)^{\gamma} \,\right| \qquad (1)$$

where R is the red channel output from the scanner;
a, b, c, and γ are constants fit by nonlinear least squares; and
Dr is the Red channel density of the object being scanned.

Similar models may be used for the Green and Blue channels, substituting G or B for R, and Dg or Db for Dr, respectively.

Again, by way of example, one of the scanners used in our investigation had the following parameterization:

---

[1] Some scanners fail to do this. Others do it, but after quantizing the data to 8 bits per channel. In either case, quantization artifacts, such as false contours, appear — after performing the "gamma correction" operation in the latter case. It is imperative that this operation be performed *before* quantization to 8 bits per channel.

Channel	a	b	c	γ
Red	0.0029	0.8293	-0.0914	1.0956
Green	-0.0049	0.8071	-0.0649	1.2112
Blue	0.0044	0.8226	-0.0896	1.1698

There are two additional points we would like to make in conjunction with the amplitude response functions. First, the values which may be digitally encoded are bounded by a minimum and a maximum. These must be considered when a value is to be digitally encoded. When estimating the parameters of the amplitude response function to a set of data, values which have been "clipped" should be excluded.

Secondly, there is a minimum level of transmittance or reflectance (zero) which makes physical sense. A transmittance of zero may be encoded as a positive digital value. For example, in the Green channel parameterization given above, a transmittance of zero is encoded by a value of 20 (rounded from 20.32). Green values below 20 would have no significance in the physical world; reflectances of zero (or a very small positive value) may replace them.

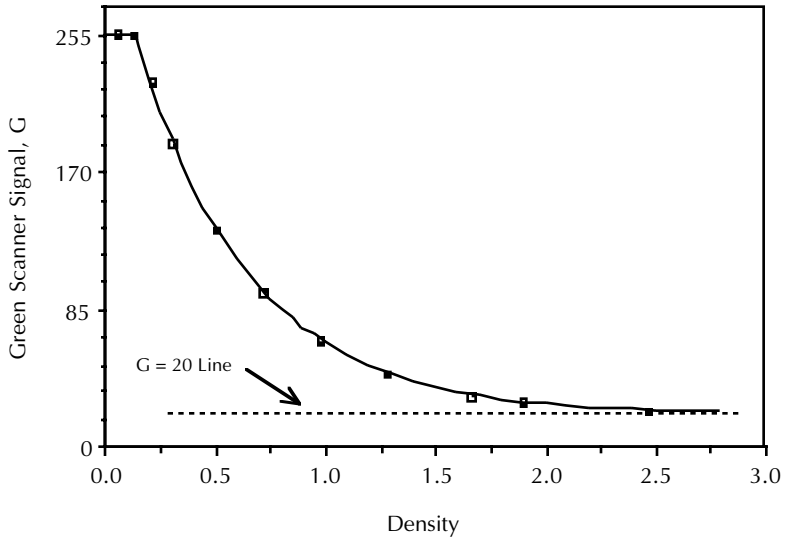

Figure 3.
The amplitude response of a scanner's Green channel is plotted. According to the parameterization given above for this channel, a scanner signal value of 20 corresponds to an infinite density. Contrast the behavior of this function near the asymptote with that of the polynomial depicted in Figure 1.

A recent paper described some significant advantages to using what was referred to as "Gray Balanced RGB" in scanner calibration using polynomials. [8] This is quite similar to the scanner's amplitude response function, though the earlier paper did not use a material independent characterization suitable for all materials.

Material-Specific, Scanner-Independent Component

This component may be performed without knowledge of the scanner to be calibrated. This permits more efficient calibration of a number of scanners, because these steps need not be repeated for each scanner. Further, they will be valid when new scanners become available.

Spectral Properties of Color Photographic Materials: Most of the materials presented to scanners are photographic. Photographic dyes exhibit little optical scatter, and closely follow the Bouguer-Lambert-Beer model, which states that the density (above that of the base) spectrum of a point is linearly related to the concentrations of the dyes. This may be written:

$$D(\lambda) = D_{base}(\lambda) + [c]\, D_c(\lambda) + [m]\, D_m(\lambda) + [y]\, D_y(\lambda) \qquad (2)$$

where $D_{base}(\lambda)$ is the spectral density of the base (*e.g.,* a D-min patch);

$D_c(\lambda)$, $D_m(\lambda)$, and $D_y(\lambda)$ are the dye density spectra of the Cyan, Magenta, and Yellow dyes at some arbitrary unit concentration;

[c], [m], and [y] are the concentrations of Cyan, Magenta, and Yellow dyes, respectively; and

$D(\lambda)$ is the spectral density of the combination.

As a consequence of this, if one knows the density above the base at three wavelengths, and the density spectra of base and the three dyes at unit concentration, one may compute the entire density spectrum. The density spectrum may be converted into a transmission (or reflectance) spectrum. Such a spectrum may be integrated to yield the tristimulus values, a set of device independent color coordinates.

Once the density spectrum of the base is subtracted, the density spectra of photographic materials tend to be three-dimensional. Three principal components (eigenvectors) will often span the range of spectra a photographic material can produce. Further development of this idea is presented in the appendix.

Because of the three-dimensional, linear nature of Equation (2), it is not necessary to use actual isolated dyes. Linear combinations of them may also be used. This will affect the concentrations. Nevertheless, it is not necessary to use perfectly isolated dyes when applying Equation (2).

This leads to the question, "What density spectra should be used for the cardinal dye density curves in Equation (2)?" In other words, should a fully-saturated Cyan be used for one of them? Using a fully-saturated Cyan patch has some disadvantages. First of all, if there is any departure from this simple linear behavior, it will tend to occur at high dye concentrations. [9] It is indicative of the dye behavior at only one concentration, and in relative isolation. This would bias the results. Additionally, it would be measured from a single (or small number) of patches on a calibration target, so it would be sensitive to several sources of error. If there is a defect in the target, even one too small to be visually detected, it will result in further bias.

A better strategy might be to average all the Cyan patches, at various dye concentrations, in the target. This is better, because it includes samples of the Cyan dye in different concentrations, is less influenced by defects, and should have lower measurement error. The same thing may be done with the Magenta and Yellow dyes.

We suggest performing a Principal Component Analysis on the density spectra of *all* patches in the calibration target. This would provide three imaginary dye density spectra, but they would be computed from the net average of all the patches in the calibration target. This offers the maximum reduction in noise and minimizes any systematic error.

If the colorants scatter, as they do with printing inks, three imaginary dyes (vectors) will not span the entire range of density spectra the colorants can produce. Further, if a fourth (or subsequent) colorant is added, at least an additional vector will be needed. During the analysis, the number of principal components to retain may be determined dynamically, [10] provided it is at least three and fewer than the number of patches.

As long as these limits are not violated, there is no real penalty in extracting an extra vector or two. Only the most significant vectors need be used later, so the decision of how many vectors to extract may be deferred if it is "hedged" in this manner.

Another small advantage of using the principal component approach is that the "dyes" will be orthogonal to each other. This simplifies the computation of dye concentrations, and adds numerical stability.

Normally, the principal components are extracted relative to the average curve. For our purposes, the analysis should be performed relative to the density spectrum of the material's D-min. Another option which is sometimes offered is to perform the analysis on the correlation matrix. This option should not be used.

A popular color photographic paper was analyzed in this manner. The first three principal components accounted for 99.92 percent of the total variance in the system. From a statistical point of view, this is considered quite good. More relevant to colorimetric calibration purposes is another lack-of-fit metric, the average ΔE^* value between measured spectra and spectra reconstructed from the three vectors. This was 0.95, close to the noise present in the measurement system.

Test Target: A test target, such as the Kodak Q60, or the IT8-7.3/1, may be used. These targets have a systematic sampling of the colors within the material gamut, D-min patches, neutral scales, and other desirable features. The concentrations of each dye (whether real or a principal component) should be computed for each patch in the target. This may be done using ordinary least-squares. For actual dyes, the concentrations may be computed as follows:

$$\mathbf{c} = (\mathbf{D}^t \cdot \mathbf{D})^{-1} \cdot \mathbf{D}^t \cdot (\mathbf{d} - \mathbf{d}_{base}) \qquad (3)$$

where **D** is a matrix whose columns contain the cardinal dyes of the system;

d is a column vector containing the density spectrum of a patch;

$\mathbf{d}_{base}$ is a column vector containing the density spectrum of the material's D-min; and

c is a column vector containing the dye concentrations.

If the cardinal dyes are principal components, the initial factor drops out, and the concentrations may be computed via the simpler formula:

$$\mathbf{c} = \mathbf{D}^t \cdot (\mathbf{d} - \mathbf{d}_{base}) \qquad (3a)$$

Equation (3a) is valid even if more than three principal components are used. If *n* principal components are used, **D** will have *n* columns, and **c** will have *n* rows.

Scanner- and Material Specific Component

Once a scanner / material combination have been selected, the previous two components may be combined. The first step in this component is to scan, under the same conditions as were used to determine the amplitude response, the target for which the dye concentrations have been determined. Values of R, G, and B are selected to represent each patch. A average of the R values within a patch may be used for the representative R, etc. We have used Adobe Photoshop® for this purpose.

The RGB values for each patch must then be converted to densities, using the Amplitude Response functions. The example amplitude response function given in Equation (1) converts digital count into reflectance or transmittance; the logarithm (base 0.1) is taken to obtain densities.

The product of this component is a simple model for computing or predicting dye concentrations from these densities. A simple linear model works well for materials which exhibit little optical scatter. In order to compensate for any DC shift in the scanner's amplitude response, which would be present to compensate for brightness, an intercept term should be included. If three dyes are used, the result would be a 3 x 4 matrix.

In order to compute this matrix, arrange the densities in columns. Put the Red densities in the first column, the Green densities in the second, and the Blue densities in the third. Prepend a column of ones to this matrix; this represents the intercept term. We shall refer to this matrix as **S**, for "Scanner." It will have as many rows as there are patches in the calibration data set. If higher-order terms are desired, they should be appended at this time.

Likewise, arrange the dye concentrations in a matrix. Again, each row corresponds to a patch in the calibration data set. We shall refer to this matrix as **C**, for "Concentration." The order of the patches must be the same as in the matrix containing the densities.

Compute the matrix relating densities and dye concentrations, again using ordinary least squares:

$$\mathbf{F} = (\mathbf{S}^t \cdot \mathbf{S})^{-1} \cdot \mathbf{S}^t \cdot \mathbf{C} \qquad (4)$$

The dye concentrations for a given patch may then be computed as follows:

$$\mathbf{c}^t = \mathbf{s} \cdot \mathbf{F} \qquad (5)$$

where **s** is a row vector containing a one followed by the densities (and any additional higher-order terms); and

c is, as before, a column vector containing the dye concentrations.

These concentrations may then be used to estimate the density spectrum of each patch using Equation (2). These density spectra may then be converted into reflectance or transmittance spectra, thence to tristimulus values, and CIELAB coordinates. The CIELAB coordinates predicted for each patch may then be compared to the CIELAB coordinates measured for each patch. This will indicate how well (or poorly) the calibration model performed.

Implementation-Specific Component

This component involves computing a set of device-independent color coordinates for a series of scanner RGB combinations. Each RGB combination is converted into a set of three densities. A set of dye concentrations are computed from these densities, using Equation (5). Once again, the dye concentrations are used to compute a density spectrum using Equation (2). This is then converted into a reflectance or transmittance spectrum, which is integrated to obtain tristimulus values. The steps are essentially the same as for evaluating the calibration model.

However, the details of this component depend upon the specific color management system being employed. As an example, a sampled multidimensional lookup table may be built. A regular 17 x 17 x 17 grid of scanner RGB, for example, may be used to address the model resulting in a sampled lookup table. More specific details are highly dependent on implementation, and are beyond the scope of this paper.

Experimental

A reflection scanner was selected for calibration. It was interfaced to a Macintosh Quadra computer using a SCSI interface, and driven from within Adobe Photoshop software using a plug-in software module provided with the scanner. The amplitude response of the scanner was determined using the set of 10 NBS tiles.

Though the reflectance spectra of these tiles are flatter than neutrals made from photographic dyes, there was still some selectivity. The densities of the tiles were measured with a Macbeth RD-514 reflection densitometer through the Red, Green, and Blue filters. Although the spectral products of the densitometer do not match those of the scanner, certainly this is better than a blind assumption of absolutely horizontal reflectance spectra. The wide band nature of this instrument makes it less sensitive to differences between its peak sensitivity wavelengths and those of the scanner.

The tiles were scanned at a spot on the scanner bed determined to offer very uniform illumination and freedom from obvious array defects. This was about 4 cm from the left edge, at the top of the bed.The RGB values for each tile were integrated over a large area to obtain average values using the Histogram tool in Photoshop. Each tile was scanned twice, and the results of the two scans were entered into software written in SAS for non-linear least squares determination of the parameters in Equation (1). The instances in which clipping of the values had occurred were not included in the analysis. This was the lightest tile in all instances, and the second lightest tile for one scan.

The results appear in the table which followed that equation. The densities were predicted for each tile, and converted into CIELAB coordinates using an approximation* based on our technique. [11] The densities actually measured for each tile were likewise converted into CIELAB coordinates, and the color difference was computed for each replicate of each tile. The average ΔE^* value was 1.4.

A Kodak Q60C target was chosen as the calibration object. It is based on a popular photographic print material. All 236 patches were measured twice with a Gretag SPM-100 spectrophotometer and the spectral reflectances for the two measurements were averaged then converted into densities.

A principal component analysis was performed on the density spectra after the density spectrum of the base was subtracted. The SAS procedure PRINCOMP was used with the "No Intercept" and "Covariance" options active. Although five eigenvectors were extracted, the first three accounted for 99.92 percent of the total variance in the spectra and were able to reproduce the density spectra of each of the 236 patches with an average ΔE^* of only 0.9. No reduction was expected if additional vectors were used. Concentrations for the first three eigenvectors were computed for each patch.

The target was scanned in 16 times, and the images were averaged using the Photoshop "Calculate" tool. The target was placed in the same area on the scanner bed as were the tiles. So that all scans would have equal weight in the average, pairs of scans were first averaged. This left eight images, each an average of two. Pairs of these images were then averaged, leaving four images. The process was repeated until a single image, the average of all 16, remained.

Average RGB values for each patch were recorded, again using the Photoshop Histogram tool, after removing a small number of defective pixels. Care was taken to avoid the edges of the patches where adjacency effects,

---

[1] The densitometer used is not colorimetric, but the spectral reflectance curves of the tiles were flat enough for this to not be a practical concern.

misregistration, and scanner line spread function could affect patch uniformity. The RGB values were converted to densities using the logarithm of Equation (1).

The least-squares fit between RGB densities and dye concentrations was determined, and the matrix was used to produce predicted values for the dye concentrations in each patch. These dye concentrations were then used in Equation (2) to estimate the density spectrum of each patch. These spectra were converted into reflectance spectra from which tristimulus values were computed, thence CIELAB coordinates. These were compared to the CIELAB coordinates measured for each patch. An average difference of 4.1 units was observed. The 90th percentile of these values was 6.2.

It is believed that these figures of merit would be improved with a better amplitude response calibration. Both scans of only eight tiles were included; more tiles might improve the results. Some of the error, however, is attributable to the bandwidth of the scanner's sensitivities, particularly in the red channel. This is a shortcoming of the technique that might be addressed by using a higher order model to relate the densities to the dye concentrations. Such a refinement would also be necessary for media which do not obey the Bouguer-Lambert-Beer law, such as printing inks.

A further refinement is to postulate a bichromatic sensitivity for one or more of the channels. This is a simple way of accommodating wider spectral sensitivity curves. The logical place to start is with the Red channel. The simple linear relationship between dye concentration and density may be replaced with a more complicated expression, which assumes that the sensitivity of the Red channel is a mixture of monochromatic sensitivities at two distinct wavelengths. Equation (4) will be replaced with a more general non linear counterpart. Although this would require more complicated non linear estimation techniques, the calibration model will still be in closed form, as opposed to a recently patented technique from Donnelley, [12] which uses an iterative solution.[*]

This may be simplified somewhat if the assumption is made that the density spectrum of the base does not vary significantly throughout the Red region, and further if it is assumed that the Magenta and Yellow dye density curves are flat in this region, as well.

---

[1] Although the technique presented here has some distinctions (the use of the amplitude response functions, and the avoidance of the iterative solution, for example) from the one claimed in the Donnelley patent, we must caution users to examine the Donnelley claims before implementing our technique.

Conclusions

A method for colorimetric calibration of color digitizing scanners was described. It involves three distinct characterizations: A device-specific, material independent characterization, in which the attributes of the scanner are studied; a material-specific, device independent characterization, in which the dyes or colorants in the material being scanned are studied; and a device- and material-specific characterization, in which the interaction between the preceding two factors is addressed. Additional steps, which depend upon the specific implementation, complete the process.

Keeping these characterizations distinct reduces the amount of labor needed for calibration. Because more than one type of material is usually encountered, the separate calibrations required for each will require less effort because less duplication will be necessary. The scanner-specific component may be performed without prior knowledge of the materials being used. Similarly, a variety of scanners may be calibrated for a certain material more easily using the new technique. Again, this component may be performed without knowledge of the scanners being calibrated. This makes it easier to accommodate new scanners.

The new technique assumes the scanner's spectral sensitivities are monochromatic. For desktop scanners with fluorescent sources, this is often not an unreasonable assumption for the Green and Blue channels. A refinement which is based on the assumption of a bi-chromatic response for each channel is introduced.

Analyzing the density spectra of the patches in the calibration target using Principal Component Analysis is a unique aspect of the technique. By performing this, and introducing a higher-order model, colorants which do not obey the Bouguer-Lambert-Beer law may be accommodated.

Acknowledgements

The authors would like to extend their sincere thanks to the Imaging Division of the RIT Research Corporation for the opportunity to conduct this investigation, and their support and encouragement during the preparation of the paper. Mr. Irving Pobboravsky assisted in many ways from the beginning.

Ms. Lisa Reniff of the Munsell Color Science Laboratory assisted us in the collection of test data. We also thank the MCSL for the use of the scanner, as well as for many stimulating discussions.

Appendix

Equation (2) may be written in matrix form:

$$\mathbf{d} = \mathbf{d}_{base} + \mathbf{D} \cdot \mathbf{c} \qquad (2a)$$

where $\mathbf{d}_{base}$ is a column vector containing the density spectrum of the material's D-min;

D is a matrix whose columns contain the cardinal dyes of the system;

c is a column vector containing the dye concentrations; and

d is a column vector containing the density spectrum of a patch.

Suppose we had a new set of dyes which were some mixture or linear combination of the old ones. One could write $\mathbf{D} = \mathbf{D}_{new} \cdot \mathbf{U}$, where **U** is the matrix which defines the linear combination. If there was a matrix $\mathbf{U}^{-1}$ such that the product $\mathbf{U} \cdot \mathbf{U}^{-1} = \mathbf{I}$ (the identity matrix), we could compute a new set of dye concentrations $\mathbf{c} = \mathbf{U}^{-1} \cdot \mathbf{c}_{new}$. Substituting these two expressions into Equation (2a), we would have:

$$\mathbf{d} = \mathbf{d}_{base} + \mathbf{D}_{new} \cdot \mathbf{c}_{new} \qquad (2b)$$

because the **U** and $\mathbf{U}^{-1}$ will cancel. This means it is not necessary to have the density spectra for the pure isolated dyes, provided what we do have is a linear combination of them, and the matrix which defines the linear combination has at least an inverse to the right.

A special type of transformation is of particular interest. It has been established that the single best linear combination is the first *principal component* of the density spectra of the patches. Because they are eigenvectors of the inner product matrix, principal components are also referred to as *eigenvectors.* Similarly, the best 2-dimensional linear combination of the dye density spectra is given by the first two principal components. Other optimal properties of principal components are discussed in Jackson [10].

Our interest lies with linear combinations of dimension 3 or greater. If there is any departure at all from the Bouguer-Lambert-Beer law, the first three principal components are a better alternative to the actual dye density curves, provided the test target adequately samples the material gamut. If the departure from Bouguer-Lambert-Beer behavior is significant, additional principal components may be extracted.

The density spectrum of the base appears as a fixed component in all mixtures; it must be subtracted out before the analysis is performed. No further "intercept compensation," such as subtracting the mean, should be performed.

The principal components play the role of imaginary dyes in our method, and we shall refer to them as such. The first principal component almost always represents a combination of all three dyes in nearly equal amounts. The second and subsequent principal components will be differences between dyes. In the case of the Kodak Q60C, the second will represent Cyan versus Magenta and Yellow, while the third will be primarily the difference between the Magenta and Yellow dyes.

Because they are orthogonal, the second and subsequent principal components will tend to have both positive and negative values. However, when properly combined in the proportions appropriate for colors within gamut, no significant negative amount should remain.

A recent scanner calibration study used principal component analysis, but the analysis was performed on spectral reflectances, rather than spectral densities. [13] Because the tristimulus values are linear in reflectance, this greatly simplifies the math – it is necessary to perform tristimulus integration only on the eigenvectors. However, Equation (2) indicates that the relationship between spectra is more linear in density space than it is in reflectance space. We have exploited that relationship in our method. Fewer eigenvectors are required to obtain the same level of accuracy.

Literature Cited

1. North, Amy, Characterizing the colorimetric properties of a flat-bed scanner using multiple-linear regression. *Munsell Color Science Laboratory Technical Report.* Rochester: Rochester Institute of Technology, 1990.

2. Hung, Po-Chieh, Colorimetric calibration for scanners and media. *SPIE Vol. 1448 Camera and Input Scanner Systems,* 1991. p 164 - 174.

3. Gerald, Curtis F., *Applied Numerical Analysis,* second edition. Reading, MA: Addison-Wesley, 1978. p 474 - 475.

4. Scott, Richard R., and E A Sylvestre, Model building and the grid technique. *Journal of Quality Technology,* 1979. 11 : 2 : 55 - 65.

5. Spooner, David L., Measurement of the transfer function of hardcopy color reproduction systems: a metric for comparison. *1992 TAGA/ ISCC Proceedings,* volume 2, p 917 - 927.

6. Hunt, R. W. G., *The Reproduction of Colour in Photography, Printing, and Television,* Fourth edition. Tolworth, England: Fountain Press, 1987. p 89 - 92.

7. Motta, Ricardo J., *An analytical model for the colorimetric characterization of color CRTs.* Master of Science thesis, Rochester Institute of Technology, 1991.

8. Kang, Henry R., Color scanner calibration. *Journal of Imaging Science and Technology,* 1992. 36 : 2 : 162 – 170.

9. Feiock, Ron, Photographic Engineer, Project Coordinator, Eastman Kodak Company. Personal communication, 1988.

10. Jackson, J. Edward, *A User's Guide to Principal Components*. New York: Wiley, 1991.

11. Viggiano, J. A. Stephen, and C. Jeffrey Wang, A proposed method for computing colorimetric densities. *1991 TAGA Proceedings,* p 196 - 215.

12. United States Patent 5 149 960, issued 22 September 1992.

13. Wandell, Brian A., and Joyce E. Farrell, Water into wine – converting scanner RGB to tristimulus XYZ. *1993 IS&T/SPIE's Symposium on Electronic Imaging Science and Technology.*