

# Scanning Calibration Targets

R. Victor Klassen, Xerox J.C. Wilson Center for Research and Technology, Webster, New York, USA

## Abstract

*Characterizing a scanner well enough to calibrate a printer continues to be challenging. Barriers include the inherent noisiness of scan bars, and the lack of colorimetric response, leading to substantial sensitivity to black addition strategy as well as changes in materials (media or colorant set). Most scanner characterization work has been designed to improve the utility of the scanner for capturing arbitrary print input, an inherently difficult problem given that neither the materials nor the black addition strategy are in general known. If these are both known, the problem becomes tractable, and noise is the largest remaining issue. This paper deals with reducing the errors in patch averages when the patches are scanned from prints made on printers that use halftoning.*

## Introduction

The usual goal of scanner characterization is to obtain the best possible result when scanning pages of unknown origin. Such pages may or may not be halftoned, and they typically contain image content, rather than simple collections of constant patches. Lack of information regarding paper colour (and fluorescence) and colourant materials limits the quality of such characterizations. Our interest here is in determining how well we can characterize a scanner for a known input. We are asking the question: if we know the printer, and its colorant set, and we completely control what is printed, can we use a scan bar to calibrate that printer, rather than using a spectrophotometer? To be successful, we need a very good scanner characterization, but it may be thoroughly tailored to that printer.

Sharma *et al*<sup>0</sup> describe the impact an unknown amount of black substitution can have on scanner characterization. In the worst case, they found errors as high as 4.5 (mean  $\Delta E$ ), while when it was fully known, mean errors dropped below 1.2. When calibrating printers, there is generally no black substitution. It suffices to characterize the scanner for this case.

Many others have characterized scanners. Ostromoukhov *et al*<sup>2</sup> obtained results of 2.37 (mean  $\Delta E$ ). One reason for their poorer results is that the printer was a desktop inkjet printer, with more noise and lower stability than the Xerographic printer used by Sharma *et al*. They noticed neighbourhood effects, and attempted to reduce their impact by using large patches. As we show later this gives incorrect results. Hardeberg<sup>3</sup> optimized a third order (3×20) matrix, obtaining  $\Delta E$  1.4 on two scanners, with less good results on others. Previously, Haneishi *et al*<sup>4</sup> had obtained  $\Delta E=2$  using a second order (3×10) matrix regression. Rao<sup>5</sup> obtained similar values. Hardeberg's thesis<sup>6</sup>, describes an experiment (p. 37ff.) in which a single scanner is characterized with a mean  $\Delta E$  of 0.92, a max of 4.67 and a 95<sup>th</sup> percentile of 2.25 on a set of 288 patches

(the same set used to calibrate). He also characterized and tested on (disjoint) subsets (p. 51), and found that when he used 144 patches to train, and the other 144 to test, the mean  $\Delta E$  rose to 0.96, but the max (of the test set) fell to 3.36 (the max  $\Delta E$  for the training set was higher, at 3.9).

Because scanners are not colorimetric, they may exhibit metamerism: colours that appear identical to a scanner might appear different to a human observer. For fixed media and black substitution strategy metamerism is not a problem. However, the conversion from RGB to XYZ varies throughout colour space. As compensation, scanners may measure far more patches per minute than spectrophotometers: we can afford to sample colour space substantially more densely. Using small patches has its own problems, however, including halftone noise and integrating cavity effect. The main contribution of this paper is techniques for compensating for these two effects.

The majority of this paper is devoted to image processing techniques improving the data used for scanner characterization. Application of a characterized scanner to grey-balance printer calibration is described briefly at the end, as a measure of the value of scanner based calibration.

## Scanner Characterization

To characterize a scanner for printer calibration, we print a number of patches of various colours, measure them with a known instrument (such as a spectrophotometer), scan them and compute the patch averages. Finding the mapping from patch average to measured value is discussed later. This section covers reducing the noise in the patch average estimated reflectances. The three primary noise sources relevant this problem, in order of spatial scale, are point process noise due to printer noise and sensor noise in the scanner, halftone noise in halftoning printers, and integrating cavity effect, causing the apparent reflectance as observed by the scanner to depend on the reflectances of neighbouring pixels.

The only difference between characterizing a scanner for printer calibration and characterizing it for arbitrary prints from that printer is the sampling of colour space. While there is no need to characterize the scanner in regions of colour space which are not used for printer calibration, all regions must be sampled in the more general case. Restricting the portion of colour space to sample offers us the opportunity to increase the sampling rate without increasing the number of sheets printed or scanned.

The remainder of this section discusses methods for reducing the effect of the three primary noise sources, in order of increasing scale.

### Point Process Noise

Printer-induced point process noise will always be a limiting factor, but it is effectively modeled as Gaussian, with an amplitude that depends on separation-wise coverage. Some specific halftone levels may be particularly noisy (e.g. where halftone dots begin to appear, and where they begin to touch); we may avoid these levels if they are few and precisely known.

Scanner-induced point noise includes dark noise and shot noise, and may be modeled or characterized as well. For example, Wach and Dowski<sup>7</sup> discuss sensor noise modeling and characterization. One may also have a relationship with the designer and/or manufacturer. For the purposes of this work, scanner noise may be well approximated as Gaussian, with an amplitude that depends on reflectance. Lacking specific knowledge of the sensor, it suffices to use an empirically derived square root form:  $\sigma \sim a + bR^{1/2}$ , for some parameters  $a$  and  $b$ . Kodak Application Note MDT/PS-0233<sup>8</sup> describes a variety of noise sources, which are either independent of signal (the lumped term  $a$  above) or proportional to its square root ( $bR^{1/2}$ ). Scanner noise at a single pixel can be as large as 0.01 (reflectance) at  $1\sigma$ , at the bright end; less at the dark end.

Printer noise can be as much as 2-3 times as large as scanner noise, for a total noise as high as 0.04. Considering the effect these values would have on a characterization, they seem large, however, they are averaged over regions of 50x50 pixels or more. The standard error of the mean is then  $0.04/50 = 0.0008$ . The noise is typically largest at a reflectance somewhere between 0.2 and 0.5, for which an error of 0.0008 corresponds to 0.4%. To reduce the worst case error from point process noise, we estimate the noise levels as a function of printer coverage and use patches with linear dimensions proportional to the estimated noise. If the expected levels of noise at (say) 0.5 and 0.6 are  $x$  and  $x/2$ , we would print a patch having four times the area (twice the width and height) at 0.5 as compared to the patch printed at level 0.6. This brings the worst case close to the average case, at about 0.05%.

A reflectance measurement error of 0.0005 at 0.03 reflectance gives an  $L^*$  error of just under 0.2. This is the specification for repeatability of the Gretag SPM 100 spectrophotometer. At a reflectance of 0.2, the same error corresponds to  $\Delta E < 0.06$ .

### Halftone Noise

Most printers halftone to create the illusion of continuous tone. Halftones may come in the form of blue noise, for printers capable of printing isolated pixels; or in the form of clustered dots, typically periodic with a spatial frequency of from 75-200 per inch. Standard practice in finding patch averages is to compute the mean RGB value of the central portion of the patch (edge pixels are avoided, lest part of a neighbouring patch contaminate the value). Taking an average with no regard to halftone period can lead to sampling errors. Usual practice is to increase the patch size making these errors small enough to be unimportant. We would rather use page real estate to increase the number of sample colours, and hence the quality of the characterization. Therefore we wish to minimize patch size, especially the minimum patch size, where the size varies.

CMYK halftoned images contain noise at the frequencies of the halftone dots used. Each channel contributes some noise to the result, however yellow contributes substantially less than the other two. Each of the RGB channels has two or three contributing halftone screens: red is affected by magenta and black, and to some extent yellow; green is affected by cyan and black, and to some extent yellow; blue is affected by magenta, cyan and black. The pattern of halftones for any one separation has a regular repeat frequency corresponding to a few (typically less than 10) pixels in each dimension; using high addressability and supercell screens increases the intensity resolution from what would be less than 100 levels to an acceptable number, however the dominant screen frequency is generally greater than 60 lines per inch, which means that the repeat pattern is less than 10 in a 600 spi scan. A simple box filter, of width equal to the dominant screen frequency, eliminates most of the halftone-induced noise. In a general de-screening application such a filter would blur edges unacceptably; for finding the average reflectance of a nearly constant patch—we filter only within the region defined by each patch—there are no edges to blur.

Were individual separations the entire story, it would suffice to average multiples of a single halftone cell. However multiple separations are involved. To capture an entire repeat pattern would require averaging multiples of the least common multiple of the cell repeat patterns for all separations involved (possibly omitting yellow, there are still typically three in blue). However, it is not necessary to go that far. Consider the case where cyan and black have repeat patterns of 7 and 9 pixels respectively (LCM=63). Any sequence of 7 pixels will have the same cyan contribution (to the extent halftone noise is the only noise source). We can average groups of 7 to obtain an “image” with the cyan halftone effectively removed. Another way to describe this is to apply a 7 pixel box filter to the image. It will still have noise, but not noise induced by the cyan halftone screen. Similarly the average of any sequence of 9 pixels will have the same black contribution. The key observation is that the average of any sequence of 9 pixels from the filtered image will have both the cyan and black screen frequencies removed. This ignores the effect of unwanted absorptions, but it works well enough [unwanted absorptions, caused by (e.g.) cyan and magenta having overlapping absorption bands, lead to some remaining halftone noise at the least common multiple].

All of this applies in two dimensions. The frequencies in the two dimensions typically differ, so for each separation the method amounts to applying an appropriately sized box filter to the image, discarding pixels too near the edge to be covered by less than a complete filter. Using summed area tables<sup>9</sup> accelerates box filtering, for large enough filters.

Once the filtered pixels have been obtained, outliers are removed, and the remaining pixels are averaged. Our approach has been to average everything between the 5<sup>th</sup> and 95<sup>th</sup> percentile, but other choices could be used. At this stage a ‘pixel’ is the average of a halftone-cell sized neighbourhood; outliers removed in this process are typically print defects.

Other than the non-linear operation of removing outliers, averaging the filtered pixels is equivalent to taking a weighted average of all the pixels in the original patch image. All pixels obtain the same weight except those around a patch boundary.

For large patches, a small error is incurred by simple averaging. Only for small patches would one expect a significant difference. One measure of the error incurred is the variation from window to window when averaging various sub-patches. In an essentially uniform patch, all sub-patches should have the same average. As a test, we scanned a series of patches and averaged both using the new (halftone-aligned) method and using the standard method. Figure 1 shows results for one sample patch. Each point represents the average of a  $52 \times 52$  sub-window of the same patch; sub-windows are offset from each other by 0.15 steps in each of X and Y. Open squares represent simple averaging; filled triangles represent halftone-aligned. The standard deviation of the simple method's averages was  $3.6 \times 10^{-3}$ ; the standard deviation of the halftone-aligned averages was  $6.05 \times 10^{-4}$ : a factor of 6 better. Halftone-aligned averaging is considerably more reliable when estimating the average reflectance of a small patch. The variation across the abscissa of the graph in the triangles reflects variation across the patch itself (it is  $1.48 \times 10^{-3}$  – less than 1% of the mean), while variation within horizontal position reflects variation in the slow scan direction of the patch ( $5.31 \times 10^{-4}$ ). Printer variation is the dominant cause of these variations.

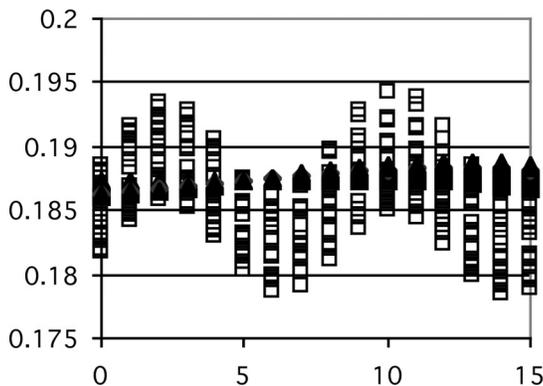


Figure 1. Open squares are averages computed using the simple approach; filled triangles are the averages computed using the new approach. The horizontal axis corresponds to the horizontal location of the window; the spread in the vertical corresponds to multiple vertical locations of the window.

For comparison, we also measured a patch of similar reflectance using a Gretag spectrophotometer. We took 225 measurements, in a  $15 \times 15$  array of locations, spaced at the minimum (0.1 mm) spacing available from the mechanical stage. The values measured were quantized (by the device) to four values, but the histogram appeared close to a Normal distribution. At  $2\sigma$ , the reflectance variation was  $9 \times 10^{-3}$ . The corresponding figure for halftone-aligned averaging is  $1.2 \times 10^{-3}$ . The spectrophotometer has a 4.5 mm diameter aperture (15.9 mm<sup>2</sup>). The scanner values were computed with a 2.33 mm square window (5.4 mm<sup>2</sup>). With halftone-aligned averaging, measurements 7.5 times more reliable

are obtained with averaging windows of just over 1/3 the area. Scaling to equal sized windows, one would expect the  $2\sigma$  value for halftone aligned averaging to be  $7 \times 10^{-4}$  using a 15.9 mm<sup>2</sup> window.

### Integrating Cavity Effect

Keith Knox<sup>10</sup> first pointed out the integrating cavity effect (ICE) in scanners. ICE causes pixel values to be closer to the values of their neighbours than they would be were their reflectances measured independently. Part of the light reflected off the page is reflected back to the page, so the reflectance of the local region affects the apparent illumination. Ignoring ICE when measuring the scanner's response to printed targets introduces unnecessary errors, likely to appear as noise, into the data. Depending on the arrangement of patches, this might introduce a systematic error into the data.

Previous solutions to the problem of integrating cavity effect have been either to change the printed target, or apply a raster based correction. Changes to the printed target include grouping similar colours together, and repeating patches to randomize their surroundings: neither approach gives the correct result.

Knox gave a raster based approach, which requires a large filter and depends on the local average changing very slowly. At the pixel level, it does, and hence the large filter. We use a patch-based approach, and hence much lower resolution. We do not require a slowly varying local average, and due to our lower resolution, apply a significantly smaller filter. The resolution is typically at least 60 times smaller. The approach requires possibly as many as 5-10 iterations (depending on the scanner), but it is still substantially faster than full resolution filtering.

### ICE Calculation

ICE results from an engineering compromise. To reduce power, a reflector focuses the light from a lamp on the page. This is true of most desktop scanners, and some higher end scanners. The reflector lamp serves as an "integrating cavity": as Knox noted, the paper repeatedly reflects a fraction of the light it receives back into the reflecting cavity. On each bounce the light is more dispersed (in position) and less of it remains, as some is absorbed by the paper or misses the mirror. Equation (1) shows an expression for the light striking the page.

$$I = \sum_k (\alpha \langle R \rangle)^k = 1 / (1 - \alpha \langle R \rangle) \quad (1)$$

where  $\langle R \rangle = G * R(x)$  is the local average reflectance, i.e. the reflectance convolved with some unknown averaging kernel  $G$ ;  $\alpha$  is a scale factor related to the initial illumination and the mirror geometry, and  $I$  is the illumination at any given point. The apparent reflectance of a given location is then

$$M(x) = R(x) / (1 - \alpha G * R(x)) \quad (2)$$

Here  $M$  must be expressed in the same units as  $R$ . There is an implied normalization, in that the integral of  $G$  must be unity.

CCD based sensors have a near linear response to light, at least over the range used in scanners. However as Eq. (2) shows, scanners appear non-linear when used to measure the average reflectances of large constant patches. At all light levels, the

denominator is less than one, corresponding to the fact that the illumination is higher than it would be in the absence of ICE, and the apparent reflectance is exaggerated (non-linearly). Instruments designed for colour measurement do not have ICE, and hence measure the actual reflectance; as a result, no simple linear scale when applied to scanner measurements can accurately match the reflectance as measured by an instrument.

ICE correction aims to find the reflectance we would have measured with the scanner, had we been measuring a very small patch, surrounded by the darkest possible black. In the limit of small patch sizes, the average surround approaches zero, and the denominator in Eq. (2) approaches 1. This is equivalent to measuring with an ICE-free scanner. In practice we cannot arrange to give every patch a black surround (and still have a reasonable number of non-black patches). Even if it were possible, the contributions of the patches themselves to ICE would introduce some error. If all patches had sufficiently similar surroundings to their own values, Eq. (2) could be inverted by assuming  $G^*R(x) = R(x)$ , but this would still be difficult near the edge of the page. Randomizing the surround of multiple instances of a patch does not produce a black “average surround”: it is closer to mid-grey. Done thoroughly enough, the average surround of all representatives of a given patch colour will approach a constant, and it might be possible to substitute a constant into the denominator of Eq. (2). However this requires many repetitions of each patch: page real estate that might be better used otherwise.

Knox provided one method of correcting: assume that the local average reflectance is slowly varying, and calculate the reflectance from the local average of the measured values. The assumption of a slowly varying local average hinges on the width of the averaging kernel being large. Herloski<sup>11</sup> derives an approximate analytical expression for the integrating kernel, and in the example he shows, the kernel exceeds 10% of its maximum value 40-60 mm from the centre. For a 600 spot per inch scan, this represents a filter width of 1900-2800 pixels. The width of the filter depends on the scanning geometry (primarily the distance between the paper and the reflector), so it may not always be this large, but at this resolution the local average reflectance is well approximated as constant. On the other hand, at patch-resolution, the filter is closer to 30 pixels wide, which is still quite large, but the assumption of a constant average is weaker. In our experience, the filter is less than 10 pixels wide, at the low resolution, making the locally constant approximation untenable.

Herloski’s expression for the shape of the kernel is separable (as the geometry would suggest), and based on a power series expansion of the integrand of a moderately complex integral. It is based on a simplified geometry and the resulting shape is well approximated by a Gaussian out to the point at which it becomes less than 0.25 of its original height. At this point 88% of the area under the kernel is accounted for. Herloski’s expression matches a ray traced simulation of the same geometry to within 0.033, while the Gaussian fits Herloski’s expression to within 0.06 over the entire range. Given that Herloski’s expression is specific to the geometry in his simulation, it seems preferable to use the more empirical Gaussian shape. The shape of the filter is symmetric

along the bar direction, and asymmetric along the direction of motion of the scan bar (perpendicular to the axis of the mirror). Using the Gaussian formulation, three parameters specify the shape of the filter: the horizontal/fast scan direction parameter, and vertical/slow scan “above” and “below” parameters.

Knox used an exponential filter, with no justification, other than that it appeared to work on the input with which he tried it. An exponential filter fits Herloski’s results rather poorly.

### Correction

Once we have the full averaging kernel, we can reverse the process of integrating cavity effect.

Assuming the average reflectance (i.e.  $G^*R$ ) varies slowly enough, we may compute it by assuming it is locally constant. Were the reflectance constant, it would be correct to re-write Eq. (2) as:

$$M(x)=R(x)/(1-\alpha R(x)) \quad (3)$$

Solving for  $R(x)$ :

$$R^0(x) = M(x)/(1+\alpha M(x)) \quad (4)$$

This gives us a first approximation of the reflectance, which we have superscripted with a 0. From this approximation, we may compute an averaged reflectance  $\langle R^0(x) \rangle$  as  $G^*R^0$ . Now given an approximation  $R^i$  to the reflectance and the measured value at each pixel, we can solve Eq. (2) for an improved approximation to the reflectance:

$$R^{i+1}(x)=M(x)(1-\alpha G^*R^i(x)) \quad (5)$$

This may be repeated until the change in the image is small. The rate of convergence depends on  $\alpha$ . Generally the number of bits of precision increases by  $-\log_2\alpha$  at each iteration until it is limited by numerical precision or measurement noise.

### Determining the Parameters

There are six parameters: three geometry-independent parameters and the three Gaussian spread parameters of the filter: two in the slow scan direction and one in the fast scan direction. The geometry-independent parameters will be discussed first.

In a large white page, the reflectance will be the reflectance of white paper,  $R_w$ , while the measured value  $M$  will be a constant  $M_w$ . The reflectance being constant, the local average reflectance will be  $R_w$  as well, so:

$$M_w = R_w (1-\alpha R_w) \quad (6)$$

Actual scanner values are not exactly the same as apparent reflectances: the scanner reflectances have been scaled so that some arbitrary white value maps to the maximum representable.  $M = M_k + (M_w - M_k)S/255$ . (We assume linearity from a CCD based scanner). Given three or more values from large regions of constant colour, we obtain for each such region

$$M_k + S_i/255(M_w - M_k) = R_i/(1 - \alpha R_i) \quad (7)$$

or

$$(1 - S_i/255)M_k + S_i/255M_w = R_i/(1 - \alpha R_i) \quad (8)$$

For any fixed  $\alpha$ , this may be treated as a matrix system which may be solved using conventional (such as SVD) techniques.

Since  $\alpha$  is in the 0..1 range, its value may be found by computing the solution to Eq. (8) for three tentative values ( $\alpha_0, \alpha_1, \alpha_2$ ), and fitting a parabola to the residual at each of these values. If the parabola has a minimum in the range between the largest and smallest value of  $\alpha$ , that is the next trial location. It will either be between  $\alpha_0$  and  $\alpha_1$  or between  $\alpha_1$  and  $\alpha_2$ . In either case a new parabola may be fit using the new point and its two neighbours. At all times, an upper bound and a lower bound to the value of  $\alpha$  corresponding to the minimum residual should be retained. If the newly fit parabola has a minimum outside the bounds, then the midpoint between the two values of  $\alpha$  with the least residuals should be used as the new value. This should converge fairly rapidly, within 5-10 iterations.

### Computing the Gaussian Parameters

To find the filter shape parameters we print a solid black page, with a square block of nine white “patches” near the middle. We superimpose upon this an imaginary grid, whose grid resolution is the size of the white patches. In other words we have a page of many “patches”, all but nine of which are black. We further assume a super-grid of super-patches, the block of nine white patches composing one such super-patch. To eliminate errors due to printer non-uniformity, we measure the page using a densitometer at the low resolution grid, and then assume that the high resolution grid is piecewise constant. (We could also assume piecewise linear, within regions of constant input colour and interpolate). Scanning the page, and averaging over “patches” (the higher resolution ones), gives a set of scan values corresponding to the reflectances found using the densitometer. Filtering the (known) reflectance values and applying the expression of Eq. (1) gives an estimate of the measurements. Comparing the actual measurements to the estimated measurements provides a measure of the conversion quality for the set of parameters for the filter. Optimizing the conversion quality by varying each of the three of the filter parameters in turn gives the optimum set of parameters. The parameters may be independently optimized, however a far-from-optimum value for one of the parameters will make the result relatively insensitive to changes in another.

In addition to the black page with a group of white patches in the centre, we may print a white page with a group of black patches in the centre. This doubles the quantity of input data, and improves robustness. The different colour channels of the scanner should share the geometry-dependent values, but the white point and black point may differ.

### Results

The method was tested using prints made on a Phaser 7750 and scanned on a prototype scanner. The values of  $M_k, M_w$ , and  $\alpha$  were

computed from nine large patches on a single page approximately equally spaced in reflectance. Values of -0.01157, 1.21078 and 0.2177 respectively provided the best fit. The RMS error between the actual reflectance and the reflectance predicted with the model was 0.0026, while the largest error was -0.0043. For comparison, when the best linear fit was found to convert from scanner measurements to reflectance, the RMS error was 0.0173, and the largest was 0.031, an order of magnitude worse.

The optimized Gaussian filter spanned seven “pixels” wide by five “pixels” high. Here a “pixel” is actually an averaged 2.55 mm region. Using that filter on the black on white and white on black pages, and comparing the computed reflectances with the measured reflectances of the 81 patches in the centre (with the contrasting patches in the center of the 9 by 9 block), the residual error was 0.024 at the 95th percentile, 0.0089 RMS. For comparison, a single pixel filter was used (equivalent to computing  $M/(1+\alpha M)$  but not filtering and iterating), the error was 0.077 at the 95th percentile, and 0.026 RMS. Linearly scaling the measured values to obtain reflectances gave a 95th percentile error of 0.158, and 0.056 RMS. These results show that simply applying a tone curve to the measurements gives a little under half the improvement, the other half coming from filtering.

These values represent the worst case: one would expect that the arrangement of patches in an actual calibration page would be unlikely to create such large errors, however the largest error could easily be half as large as these.

### Printer Calibration

The proof of the pudding is in the tasting; in order to test the value of these techniques we characterized the scanner used in the above experiments over a region spanning the neutral axis, and then used the characterized scanner to measure prints for grey balancing the same printer.

Over all of the patches used in characterizing the scanner, the model (a 3D B-Spline) fit the data with an error of  $\Delta E = 0.9$  at the 95th percentile, 0.42 mean, max 1.8. This only tells us how well we fit the input data, and not how well intermediate colours would be predicted. These values are significantly better than the best reported value from the literature: Hardeberg’s 0.92 mean and 3.36 max.

As a more challenging test, the scanner was used to compute grey balance TRCs for the printer for which the scanner was characterized. That is, for each of many steps along the  $L^*$  axis the CMY yielding the colour closest to that  $L^*$  with  $a^*=b^*=0$  was found. These tone response curves are applied individually to the three separations. Then a series of “neutral” patches was printed and measured on a Gretag spectrophotometer, the figure of merit being the neutrality of these patches. The overall quality of the result depends on the stability of the printer (spatially and temporally), and the particular algorithm used to find the grey balance TRCs. Thus absolute values are less important than relative values. The patch size for printer calibration was approximately 2.55 mm and multiples thereof on a side; for scanner calibration the minimum patch size was 5.1 mm,\* but 2.55 mm subpatches were

used for calculating ICE correction. With neither ICE correction nor aligned averaging applied, using patches this small gave a very poor grey balance result:  $\Delta E=12.0$  at the 95<sup>th</sup> percentile, mean  $\Delta E=9.3$ . Applying only ICE correction brought it down to 4.0/1.8, an improvement of 66%/80%. Adding in aligned averaging reduced the 95<sup>th</sup> percentile error further, to 3.1, leaving the mean essentially unchanged (1.8). For comparison, the same printer was calibrated using the spectrophotometer, with results 3.2/1.5. These numbers do not differ meaningfully from the scanner results.

As only a single paper stock was involved, and grey balance TRCs do not involve K, the ability of the spectrophotometer to produce spectral data rather than  $L^*a^*b^*$  was no advantage. While the spectrophotometer-based grey balance measurements and computation involved spectral data, the results were equivalent.

## Conclusions

We have presented three techniques for reducing the noise in scans of halftoned patches: variable patch size, halftone-cell averaging and integrating cavity effect correction. Taken together, these allow the scanner to operate as an acceptable surrogate for a spectrophotometer, when calibrating a printer, with the advantage of dramatically reduced measurement time. In fact, the errors we measured were within the repeatability of the spectrophotometer: the scanner may even be better than standard against which we measured it.

These results are specific to calibration, a process normally controlled at the single-separation level, and in which the amount of black ink is (trivially) controlled. Results for a full characterization, which includes a variable amount of black ink substitution may be worse, although for a fixed black substitution strategy we would expect them to generalize. The main advantage with simple separation independent calibration is that the scanner needs to be characterized over a much smaller portion of colour space.

With these improvements to scanner-based printer patch measurements, we are now able to measure much larger numbers of patches in a reasonable amount of time. (An A3 sheet has room for 19,000 patches with no margin). This leads to new problems in effectively using this data to compute robust models of printer behaviour, in a reasonable amount of time. We expect to continue to improve our calibration results as we tackle these issues.

## References

- \* Characterizing the scanner requires measuring the same patches with the spectrophotometer, and its (circular) aperture is 4.5 mm, hence the larger patch size.
- 1. G. Sharma, S. Wang, D. Sidavanahalli and K. Knox "The impact of UCR on scanner calibration", in proc PICS Conf., pp. 121-124, Portland, OR, 1998.
- 2. V. Ostromoukhov, R.D. Hersch, C. Péraire, P. Emmel, I. Amidror, "Two approaches in scanner-printer calibration: colorimetric space-based vs. closed-loop" in proc. SPIE 2170, pp. 133-142 1994
- 3. J. Hardeberg, "Desktop Scanning to sRGB" in IS&T and SPIE's Device Independent Color, Color Hardcopy and Graphic Arts V, (San Jose, CA), Jan. 2000.
- 4. H. Haneishi, T. Hirao, A. Shimazu, and Y. Miyake, "Colorimetric precision in scanner calibration using matrices," in Proceedings of IS&T and SID's 3rd Color Imaging Conference: Color Science, Systems and Applications, pp. 106-108, (Scottsdale, Arizona), Nov. 1995.
- 5. A. R. Rao, "Color calibration of a colorimetric scanner using non-linear least squares," in Proc. IS&T's 1998 PICS Conference, (Portland, OR), May 1998.
- 6. J. Hardeberg, Acquisition and Reproduction of Colour Images: Colorimetric and Multispectral Approaches, Doctoral Dissertation, l'Ecole Nationale Supérieure des Télécommunications, Paris 1999.
- 7. H. Wach and E. Dowski, "Noise modeling for design and simulation of color imaging systems", in proc CIC 12, pp. 211-216, 2004
- 8. Eastman Kodak Corporation, "CCD Image Sensor Noise Sources", Application Note MDT-PS-0233, Rev. 2.1 Jan 2005. <http://www.kodak.com/global/en/digital/ccd/>
- 9. Frank Crow, "Summed-area tables for texture mapping." Proceedings of SIGGRAPH '84. In Computer Graphics 18, 3 (July 1984), pp. 207-212
- 10. Keith Knox, "Integrating cavity effect in scanners", proc. IS&T/OSA Conference on Optics and Imaging in the Information Age, Rochester 1996, pp. 156-158
- 11. R. Herloski, "Analytical Computation of Integrating Cavity Effect", in proc. SPIE 3482, International Optical Design Conference, 1998

## Author Biography

*R. Victor Klassen received a B.Sc. in Physics and a Ph.D. in Computer Graphics from the University of Waterloo, then moved to Xerox, where he is now a Principal Scientist. His current interests include physical limitations of color measurement and reproduction, and limitations of the visual system to discriminate colours. A member of ACM SIGGRAPH, IEEE-CS, and IS&T, he has served on the program committees of SIGGRAPH, Graphics Interface, and RIDT.*