

Colour Laser Scanner Characterisation by Enhanced LUT

Lindsay W. MacDonald, University College London (UCL), UK

Abstract

This study investigated how to improve the accuracy of colour characterisation for a three-colour laser scanner, implemented by a lookup table (LUT) with interpolation. The transfer function was trained on a huge number of real and synthetic reflectance spectra, refined through statistical analysis. The lookup table enabled a 'baseline' matrix fitting to be enhanced through local deformations of 3D colour space to give optimal colorimetric performance.

Background

The signals produced by digital image capture devices, usually encoded at each pixel position as red, green and blue (R,G,B), are device dependent. The relationship between the object's reflectance spectrum and each signal depends on the spectral sensitivity of the corresponding channel of the device. In a digital scanner, for example, this includes the integrated product of the power at each wavelength of the inbuilt light source, the reflectance of the surface, the transmittance of the colour filter, and the sensitivity of the sensor.

Input device characterisation is the process of establishing the relationship between the colour of the original object and signals generated by the device (Fig. 1). The former is expressed by device-independent values related to human vision, such as CIE $L^*a^*b^*$, and the latter by device-dependent R,G,B values. Simplistically, if the signal can be represented as a function of the colour: $\text{Signal} = f(\text{Colour})$ and if the function is invertible (i.e. continuous and single-valued) then the colour may be recovered as the inverse function: $\text{Colour} = f^{-1}(\text{Signal})$.



Figure 1. 'Black box' model of input device behaviour

Because there is generally no analytical function available, an approximation is frequently used, by fitting the data with a low-order polynomial through a regression procedure. Hong *et al* [1] tested six different orders of polynomial and found the best overall performance (i.e. lowest average error) with an 11-term polynomial, including a cubic term for RGB, over a set of 168 test samples. They also investigated the number of training samples needed, and concluded that 60 samples were sufficient.

The problem is that the transfer function f should be optimised over all possible reflectance spectra that the device could encounter. If the usage will always be limited to a specific population of materials or objects, for example photographic prints or paints or textiles made with known dyes, it may be sufficient to take a selection of test samples from that population. But for a digital camera or general purpose scanner that could be used for any material it is not sufficient to train it on a small number of samples with a limited range of reflectance spectra, e.g. by using a standard test target such as the GretagMacbeth DC chart, and then expect the transfer function to predict accurately the device performance on all possible spectra that might be encountered in practice.

Characterisation of colour laser scanner

This study aimed to characterise a large colour laser scanner (Fig. 2), used at UCL for 3D digitisation of heritage objects from Museum collections, such as Egyptian tablets. Because it employs three lasers to sample the object surface at three precise wavelengths (473, 532 and 635 nm), the colour response is highly metameric [2]. Its signals are much more sensitive to the form of the reflectance spectrum than a normal scanner or digital camera or the human eye, for which the channel spectral sensitivities are quite broad.



Figure 2. The Arius 3D colour laser scanner at UCL.

The objective was to use sets of reflectance spectra to determine the inverse transfer function f^{-1} , to transform the R,G,B signals from the scanner sensor directly into CIELAB colorimetric values. The reference tristimulus values X,Y,Z were calculated by multiplying the spectral reflectance distribution of each sample by the CIE 2° Standard Observer and the spectral power distribution of the D65 illuminant (Fig. 3). Knowing the spectral sensitivity of the scanner (in this case three delta functions) enabled the scanner R,G,B responses to be predicted, by sampling the reflectance spectrum at each of the three specific wavelengths. The coefficients of a 3x3 matrix were calculated by a regression procedure, enabling the corresponding X',Y',Z' to be predicted for each R,G,B . From the two sets of X,Y,Z the CIE $L^*a^*b^*$ values were calculated using the tristimulus values of D65 as the reference white.

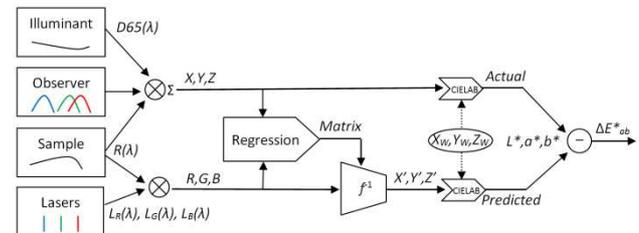


Figure 3. Procedure for calculating matrix coefficients and testing the accuracy of prediction.

A large dataset of 8,714 reflectance spectra was collated from readily available spectral reflectance measurement sets (Table 1) over the wavelength range 380 to 780 nm, and

interpolated to 1 nm intervals by the Matlab function `interp` with the cubic spline option. Where the original wavelength range was smaller, for example with no data provided below 400 nm or above 700 nm, the reflectance was set to zero. As a baseline for the transform, a regression procedure was employed to find coefficients c_{ij} that would minimise the error in the tristimulus X, Y, Z domain over the full dataset:

$$\begin{aligned} X_i &= c_{11}R_i + c_{12}G_i + c_{13}B_i \\ Y_i &= c_{21}R_i + c_{22}G_i + c_{23}B_i \\ Z_i &= c_{31}R_i + c_{32}G_i + c_{33}B_i \end{aligned} \quad (1)$$

where R_i, G_i, B_i are normalised signal values from the detectors in laser scanner for $i = 1..n$ and $n = 8174$. The system of equations can be written more compactly in matrix form as:

$$[X \ Y \ Z] = [R \ G \ B] \mathbf{M} \quad (2)$$

where the $n \times 1$ response vector $X = [X_1 \dots X_n]^T$ (and similar for Y and Z), the $n \times 1$ signal vector $R = [R_1 \dots R_n]$ (and similar for G and B) and the 3×3 matrix $\mathbf{M} = [c_{ij}]$ contains the coefficients. For the collated dataset of real samples the resulting matrix was:

$$\mathbf{M}_{\text{real}} = \begin{bmatrix} 48.61 & 28.84 & 4.93 \\ 34.96 & 68.68 & -6.69 \\ 10.18 & 1.81 & 107.02 \end{bmatrix} \quad (3)$$

The performance of this simple linear fit of laser R, G, B to tristimulus X, Y, Z was evaluated by converting both the reference colour and the estimated colour of each sample to $L^*a^*b^*$ and calculating the colour difference as ΔE^*_{ab} (Fig. 3).

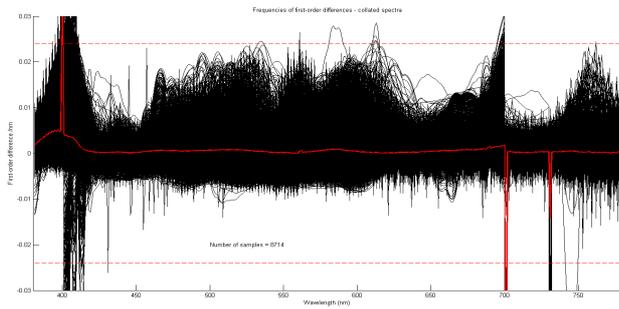


Figure 4. Slopes (first derivatives) of 8,714 real reflectance spectra, with mean (red). The spikes at 400, 700 and 730 nm are artifacts arising from discontinuities in some datasets that were undefined outside this range.

Synthetic spectra

To simulate the widest possible range of reflectance spectra that might be encountered by the scanner, a set of synthetic spectra was generated. These were constrained to be continuous, single-valued functions of wavelength at 1 nm

intervals over the range 380 to 780 nm, with a maximum slope (first derivative) not exceeding $\pm 0.024 \text{ nm}^{-1}$. The latter limit was derived from analysis of the collated spectra (Fig. 4), which showed that the first derivatives fall largely within the range from -0.01 to $+0.02 \text{ nm}^{-1}$. These values can be related to the colour gamut [3].

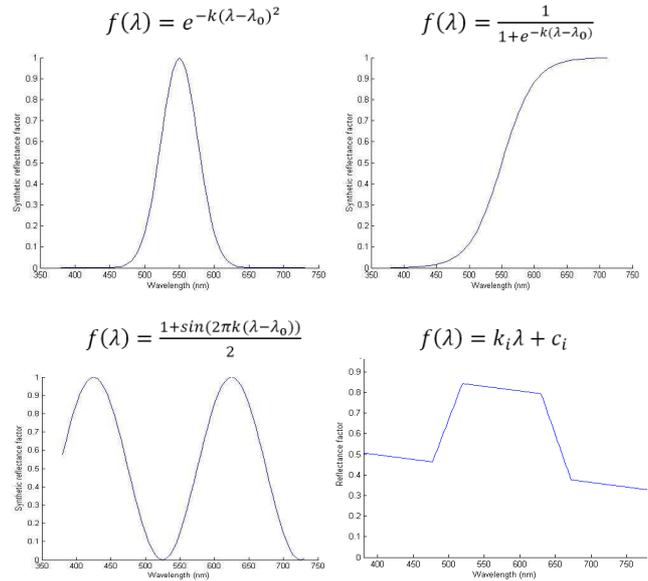


Figure 5. Four generating functions for synthetic spectra: (top left) Gaussian; (top right) logistic; (bottom left) sine; (bottom right) ramp.

The synthetic reflectance spectra were based on eight generating functions: Gaussian, inverted Gaussian, logistic, inverted logistic, sine, upward ramp, downward ramp, sum of three Gaussians, and piecewise linear (Fig. 5). The central wavelength λ_0 and width parameter k were randomised to produce a family of 100,000 curves for each function, giving 800,000 synthetic spectra in total. The central wavelengths ranged from 400 to 700 nm, and the full range of reflectance factors was utilised, from 0 to 1. Limits were set on the slope and width parameters to restrict the maximum slope of the curves in all cases within the range ± 0.024 per 1 nm interval.

$L^*a^*b^*$ coordinates of all real and synthetic spectra were calculated for the D65 illuminant and plotted in 3D (Fig. 6). The real samples fit within cuboidal bounds of approximately $10 < L^* < 95$ and $-60 < a^* < +70$ and $-55 < b^* < +105$. The synthetic spectra, however, fill a much larger volume of $0 < L^* < 100$ and $-130 < a^* < +135$ and $-125 < b^* < +140$. They are therefore much more demanding as a test of the robustness of the device characterisation.

Table 1. Collection of 8,714 reflectance spectra assembled from various sources.

Spectral reflectance dataset	Source	Surface type	Instrument	Geometry	Wavelength interval (nm)	Number of samples
NCS atlas	NCS	Matte	Gretag Spectrolino	45-0	10	1950
NCS atlas	NCS	Matte	Macbeth ColorEye 7000	d-8	10	1950
Munsell atlas	Joensuu	Gloss	Perkin-Elmer lambda 18	45-0	1	1600
Munsell atlas	Joensuu	Matte	Perkin-Elmer lambda 9	d-8	1	1269
Chinese NCS	Shanghai	Textiles	Datacolor 600	d-8	10	899
Natural materials	Derby	Leaves	Macbeth ColorEye 7000A	d-8	10	494
Natural materials	Joensuu	Plants	Acousto-Optic Tunable Filter	45-0	5	218
GretagMacbeth DC chart	LCC	Matte	X-Rite i1Pro	45-0	10	180
Berns-Taplin acrylic paints	RIT	Paints	X-Rite i1Pro	45-0	10	100
Traditional art pigments	Natl. Gallery	Paints	Monolight	45-0	2	64

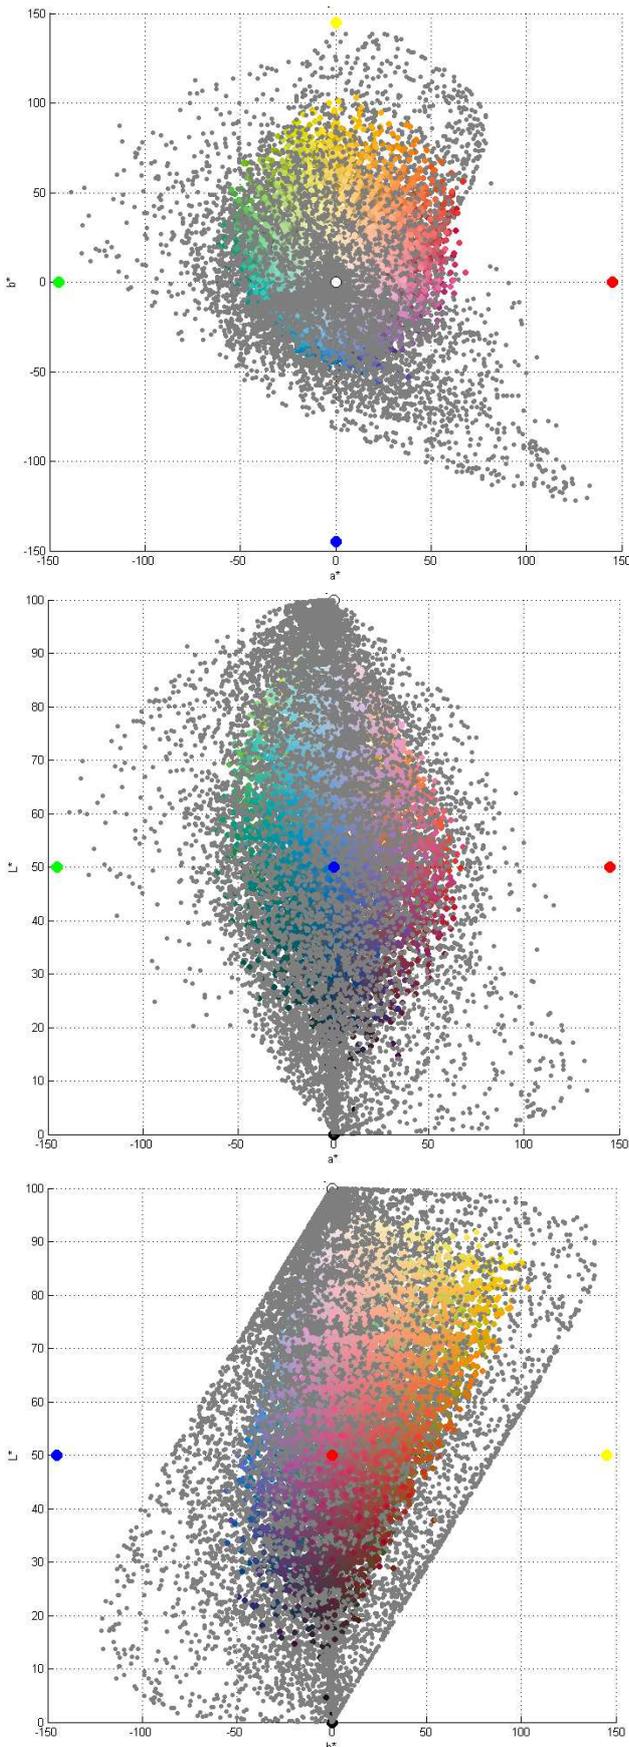


Figure 6. Colours of 8,714 real (coloured) and 20,000 synthetic (grey) spectra in CIELAB space for D65 in three cross-sections: (top) a^*b^* plane; (centre) L^*a^* plane; (bottom) L^*b^* plane.

Colour transformation by lookup table

The transformation from object colour to device RGB is conveniently implemented by a lookup table (LUT), with less computation and faster speed for real-time applications. The contents of the LUT can be pre-computed and also edited to include media-dependent corrections or operator preferences.

The first use of LUTs was in digital drum scanners for the digitising of transparency photographs for graphic arts and colour printing [4]. The logarithm of intensity of light in red, green and blue (R, G, B) channels transmitted through each point of the transparency was digitised to give density values D_R, D_G, D_B . These were used to address the correct cell of the LUT in which were stored corresponding values of print density or ink %dot in cyan, magenta, yellow and black (C, M, Y, K) [5].

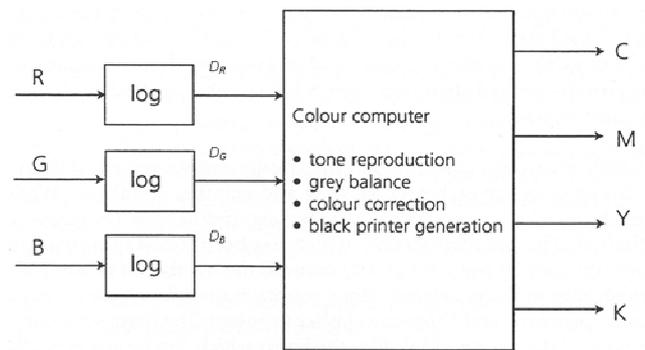


Figure 7. Signal processing by LUT in the Crosfield scanner, c.1975

For good quality picture reproduction, at least eight bits per channel are required, preferably 10 bits in a logarithmic density domain or 12+ bits in a linear domain. To implement this as a LUT with a separate cell for every possible combination of input values would require a prohibitively large amount of memory. For example the arrangement in Fig. 7 would need $2^{24} = 16$ million cells with four output bytes per cell (64 Mbyte) for 8-bit input, and $2^{30} = 1$ billion cells (4 Gbyte) for 10-bit input. The solution has been to separate the input signals into two parts: the most significant bits (msb) address a smaller LUT, and the least significant bits (lsb) are used to interpolate values from all surrounding points in the lattice. With a 4:4 allocation of 8 input bits, the LUT requires only $2^{12} = 4096$ cells (16 Kbyte). The simplest interpolation scheme is trilinear amongst all eight corner points of the cubelet within which the input value falls (Fig. 8), but other schemes are widely used [6].

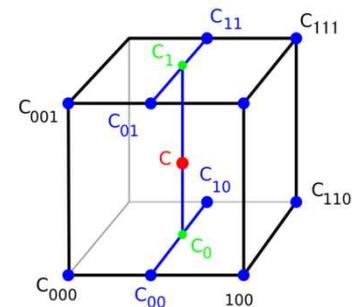


Figure 8. Trilinear interpolation

To implement the colour transformation for the colour laser scanner, a lookup table with $33 \times 33 \times 33 = 35,937$ cells was constructed, addressed by the 5 msb of each of the R, G, B laser signals. The 7 lsb of the 12-bit linear input signals were used to interpolate amongst the eight corner points of the cubelet. The cells of the cube were initially loaded with the $L^*a^*b^*$ values predicted by the matrix in Eq. 3 and standard conversion from X, Y, Z . This is similar to the arrangement used by Hung for the colorimetric calibration of imaging devices [7].

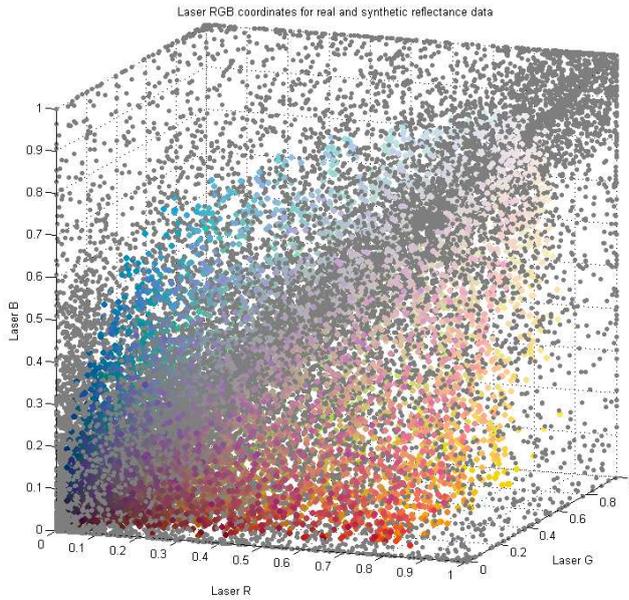


Figure 9. Distribution of all samples in the LUT in the RGB address space. The grey points represent 20,000 synthetic spectra.

Plotting all samples within the normalised cube (Fig. 9) reveals that the real samples fill only a limited volume around the grey axis up the long diagonal from the black corner to the white corner. The coordinates corresponding to the synthetic samples, while still clustered around the grey axis, are spread much more widely and fill all of the cube volume although they are sparse in some regions. This suggests a novel way of improving the accuracy of the lookup table, by loading it with the actual $L^*a^*b^*$ values corresponding to the spectra, instead of the approximate values predicted by the regression formula. The 808,714 spectra available, if they were uniformly spread throughout the RGB cube address space, should yield an average of over 22 values in each of the 35,937 cells. In fact the distribution turns out to be very uneven where some cells, especially dark colours near neutral, have hundreds of entries whereas others have none. For conservation of memory, the maximum number of entries in each cell was limited to 101, commencing with the real samples. Fig. 10 shows the scatter in such a distribution of $L^*a^*b^*$ values.

The median was taken of all values accumulated in each cell. Because the first entry in each cell is the predicted value (from the regression matrix), there is always guaranteed to be at least one entry, even in sparse regions with no corresponding values from spectra. Fig. 11 shows the resulting set of 33 values of L^*, a^*, b^* along one row of changing R and constant G and B . In this case the predicted value matches the actual value closely for L^* , but there are significant departures for a^* and b^* . In order to achieve better homogeneity of the function along all three dimensions of the lookup table, the data was filtered at every internal point by taking the median of the 5x5x5 surrounding cells.

This method proved to be very effective as a way of removing isolated values and reducing noise (Fig. 12). The resulting transfer function can be visualised effectively in 3D for planes through the LUT where one of the input values is held constant (Fig. 13). The local deformations in the 3D function loaded in the LUT enable the colour transformation to adapt to the most probable signal produced for each spectrum, far more sensitively than any polynomial fitted over the whole dataset.

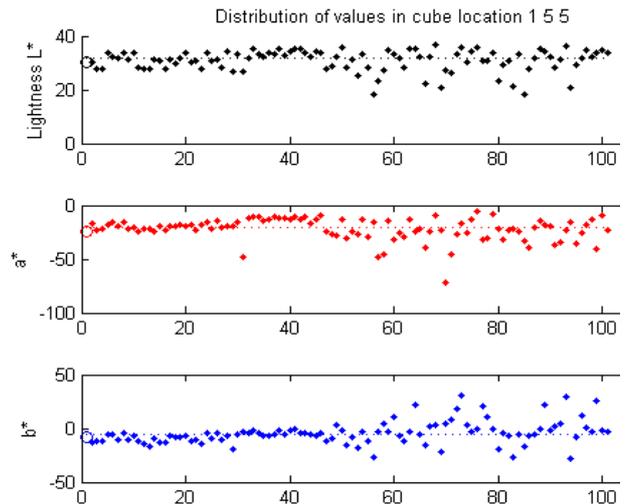


Figure 10. Distribution of $L^*a^*b^*$ values in one cube cell, all with the same R,G,B index values. The first entry (open circle) is the predicted value.

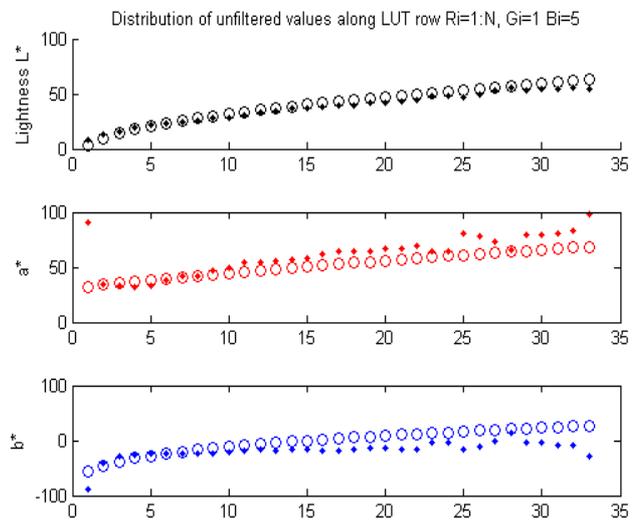


Figure 11. Values along one row of cube after taking the median in each cell (dots). The open circles are the values predicted by the matrix.

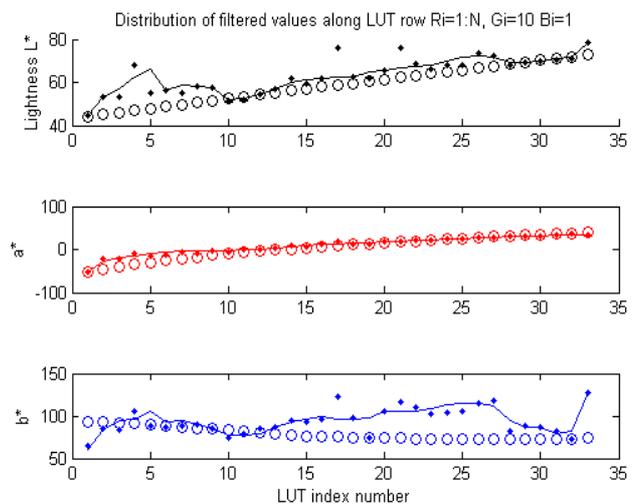


Figure 12. Values along one row of cube after applying the 5x5x5 median filter. The open circles are values predicted by the matrix, dots are unfiltered values and lines are filtered values.

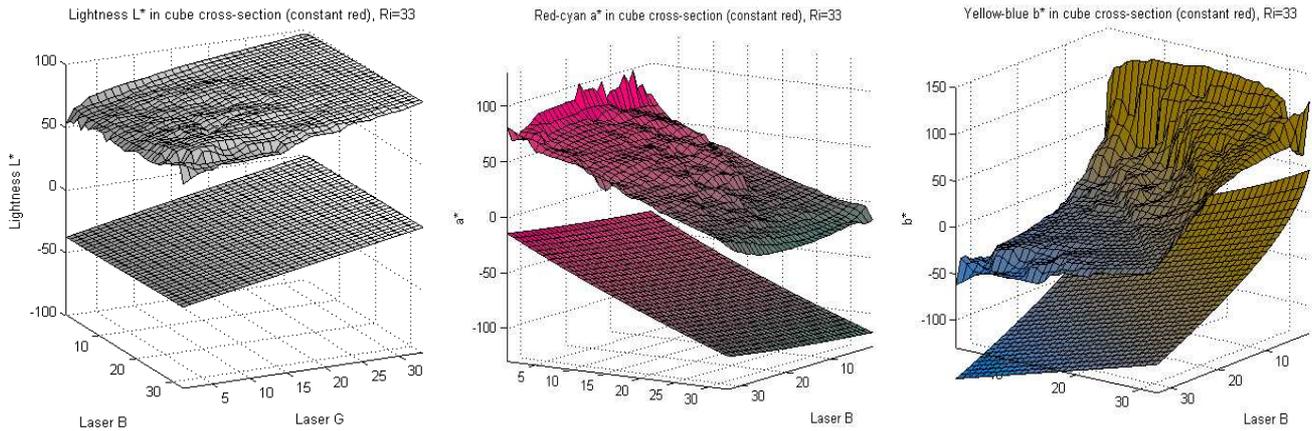


Figure 13. Values of L^* , a^* and b^* values for planes of constant $R=33$. The lower surface in each case shows the values predicted by the baseline matrix transform (displaced downwards for visualisation) and the upper surface shows the local deformations resulting from adaptation to the spectra.

Evaluation

Testing the performance of the simple 3×3 matrix M in Eq. (3) with the set of 8,714 real spectra gave the distribution shown in magenta in Fig. 14, with a median error $\Delta E_{ab}^* = 3.38$ (Table 2). The same matrix applied to the full set of 808,714 real and synthetic spectra resulted in a median error of 5.00 but a much longer tail in the distribution, with the 95th percentile increasing from 14.81 to 33.15. The maximum error rose to over 197, indicating that for a few spectra the monochromatic sampling by the laser scanner produces huge colorimetric errors.

The fitting procedure of Fig. 3 and Eq. (2) was repeated for the full dataset of 808,714 real and synthetic samples. The resulting matrix was:

$$M_{\text{full}} = \begin{bmatrix} 49.96 & 30.56 & 7.53 \\ 32.40 & 65.87 & -5.59 \\ 12.17 & 3.39 & 105.43 \end{bmatrix} \quad (4)$$

Comparison of this matrix with that fitted to the real samples only (Eq. 3) indicates close agreement in all coefficients. Testing the performance of this matrix on the real samples (Fig. 10) shows that it has a slightly worse performance (median error of 3.66 instead of 3.38), but an improved performance for the full dataset of combined real and synthetic samples (median error of 4.89 instead of 5.00), as expected.

Small further improvements to the performance of the regression could be made by introducing additional terms (for example by second- or third-order polynomials), but these are not justified by the nature of the linear system. The errors arise not from non-linearities but from the monochromatic sampling of the spectra which cannot satisfy the Luther-Ives condition.

Table 2. Performance of matrix and LUT on spectral datasets.

Method	#samples fitted to	#samples tested on	Med	Mean	95 th %ile	Max
Matrix	8174	8174	3.38	4.98	14.81	36.34
Matrix	8714	808714	5.00	9.87	33.15	197.22
Matrix	808714	8174	3.66	5.22	15.40	40.65
Matrix	808714	808174	4.89	9.67	33.55	180.69
LUT	808174	8174	3.38	4.74	13.22	26.79
LUT	808174	100000	4.61	8.34	28.80	175.18

Testing the performance of the lookup table method on the 8,174 real samples gave the error distribution shown in red in Fig. 14. Overall the LUT performed better on the set of spectra of real samples than a 3×3 matrix fitted to those samples, with

the same median error of 3.38 but a reduced 95th percentile value of 13.22. A similar improvement was achieved when testing the LUT on a large set of 100,000 spectra randomly sampled from the full set. Table 3 shows that the errors for the LUT are lower by every measure for both the real samples and the full dataset. The LUT method reduces errors by maximising the probability that output values correspond to the spectra most likely to be encountered in practice.

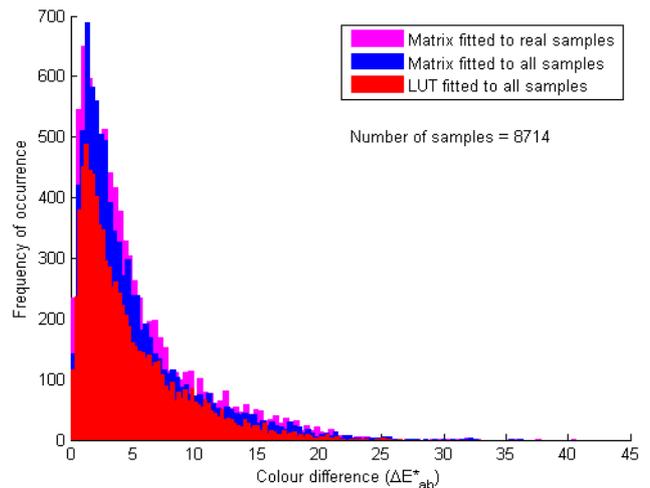


Figure 14. Distribution of colorimetric errors for 8,714 spectra of real samples for: (magenta) 3×3 matrix fitted to real samples; (blue) 3×3 matrix fitted to all samples; (red) LUT fitted to all samples.

Plotting the spectra of the 100 samples that produced the largest colorimetric errors (Fig. 15) shows the common characteristic of large changes in reflectance between the laser wavelengths. In the majority of spectra the reflectance factors, and hence the responses in the R, G, B channels of the scanner, would be well correlated. But for these spectra, the transition from a very small G response to a large R response, and vice versa, is relatively rare in the whole population of spectra. This low probability means that such spectra have little influence on the construction of the transfer function in the LUT.

In general spectra for which one of the R, G, B signals is zero are less well behaved than spectra corresponding to points near the centre. Plotting the differences in L^*, a^*, b^* values against R, G, B laser responses (Fig. 16) shows that the errors are well distributed around zero and tend to be larger for smaller signal values. The scatter is greater for the R channel than G and B .

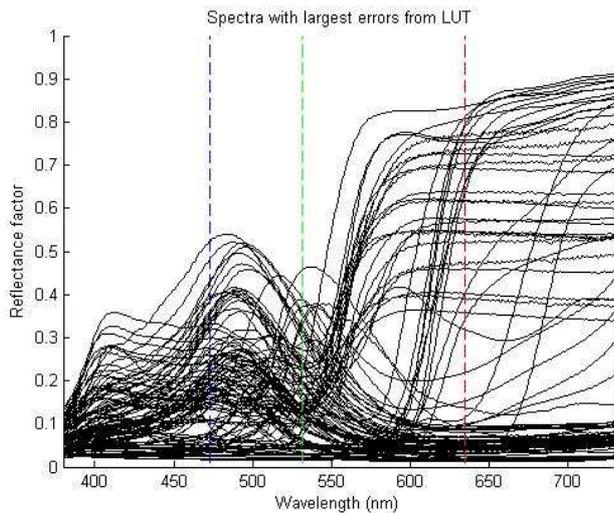


Figure 15. One hundred spectra of real samples producing the largest errors with the LUT technique. The dotted vertical lines are the wavelengths of the three lasers at 473 nm (blue), 532 nm (green) and 635 nm (red).

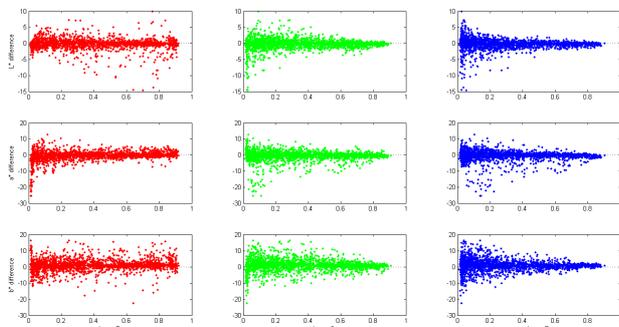


Figure 16. Colorimetric errors for L^* (top row), a^* (middle row) and b^* (bottom row) for normalised values of R,G,B laser response.

Plotting the cross-sectional planes of constant R , G and B in the LUT (Fig. 17) shows how the colour in each cell is subtly modified by the statistical influence of the spectra. For the purposes of visualisation the colours in this figure have been converted from CIE $L^*a^*b^*$ to sRGB, and out-of-gamut colours mapped onto the nearest point on the sRGB gamut boundary. The effect is most pronounced in the three planes for which one of the signals is zero, corresponding to the three faces of the cube intersecting at the black corner (Fig. 18).

Conclusion

In this study has been demonstrated a means of improving the colorimetric accuracy for characterisation of colour laser scanners over what can be achieved by the conventional fitting of a 3×3 matrix. The proposed method is novel because it uses the statistics of a huge population of reflectance spectra (both real spectra measured from samples and synthetic spectra generated from functions) to apply local adaptation (i.e. a localised deformation of the transfer function in the 3D LUT) to the prediction of a surface in colour space fitted by regression to the dataset. This gives optimal performance for device colour characterisation. The method is applicable to any type of input device, including scanners and cameras, for which the spectral sensitivity of each channel is known *a priori*. It could also be applied to multispectral imaging systems.

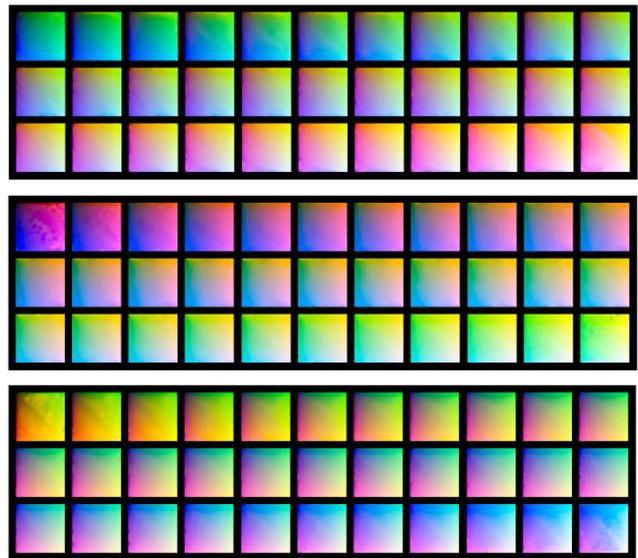


Figure 17. Cross-sections of the $33 \times 33 \times 33$ cube showing planes of constant red (top), green (middle) and blue (bottom) in sRGB colour space.

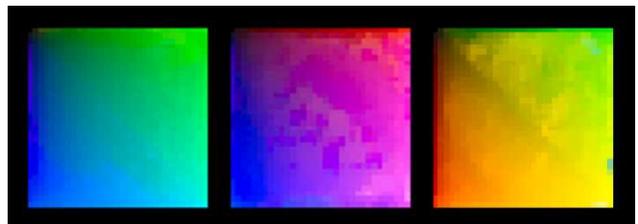


Figure 18. Cross-sections of the three faces of the cube for $R=0$, $G=0$ and $B=0$ showing the effect of the spectra.

References

- [1] Hong G, Luo MR and Rhodes PA (2001) A Study of Digital Camera Colorimetric Characterisation Based on Polynomial Modelling, *Color Res. Appl.*, **26**(1), 76-85.
- [2] MacDonald LW (2011) Choosing Optimal Wavelengths for Colour Laser Scanners, *Proc. 19th IS&T/SID Color Imaging Conf.*, San Jose.
- [3] Buchsbaum G and Gottschalk A (1984) Chromaticity coordinates of frequency-limited functions, *J. Opt. Soc. Am. A*, **1**(8):885-887.
- [4] Pugsley PC (1975) *Colour Correcting Image Reproducing Methods and Apparatus*, US Patent 3,893,166
- [5] Johnson AJ (1996) *Colour Management in Graphic Arts and Publishing*, PIRA Press, Leatherhead, UK.
- [6] Kang HR (1997) *Color Technology for Electronic Imaging Devices*, SPIE Press, Bellingham WA.
- [7] Hung PC (1993) Colorimetric calibration in electronic imaging devices using a lookup table model and interpolations, *J. Electronic Imaging*, **2**(1), 53-61.

Author Biography

Lindsay MacDonald is a Fellow of IS&T and has had twin careers in industrial R&D and academia as a colour engineer. He has edited eight books on colour image science and its applications in cultural heritage. He is a Consultant Editor for the IS&T book series on image science published by John Wiley. He was Co-Chair of the Color Imaging Conference in 1998 and is currently a member of the Executive Committee of the International Colour Association (AIC), Editor of the annual AIC Newsletter, and Co-Chair of the AIC Congress to be held in Newcastle UK in 2013.